

Mapping the Emerging Literature on AI Centers of Excellence: A Bibliometric Review of Deployment, Governance, and Capability-Building Research

Yifan Chen¹, Lin Peng^{2, *}

¹ School of Computer Science and Engineering, Beihang University, Beijing 100191, China

² College of Information Engineering, Putian University, Putian 351100, P.R.China

* penglin@ptu.edu.cn

Article Information

Received	15 October 2023
Accepted	26 February 2024
DOI	https://doi.org/10.63646/datamind.2024.020105

Abstract

Artificial intelligence centres of excellence (AI CoEs) have emerged as purpose-built organisational units designed to accelerate the deployment, governance, and enterprise-wide capability building of AI technologies. Despite rapid practitioner interest, the scholarly landscape on AI CoEs remains fragmented across multiple disciplines, and no systematic bibliometric overview of this field has yet been published. This study conducts a bibliometric review of 1,247 peer-reviewed documents published between 2016 and 2025, retrieved from Scopus and Web of Science using a structured query protocol. The analysis applies publication-trend tracking, geographic-distribution mapping, keyword co-occurrence analysis, and thematic clustering to identify the structural and intellectual contours of the field. Five major thematic clusters are identified: AI strategy and policy, capability building, AI governance and ethics, deployment frameworks, and digital transformation ecosystems. The United States, the United Kingdom, and China together account for more than half of the global output. Publication volume grew more than tenfold over the review period, with the sharpest acceleration occurring after 2020. Governance-related research has overtaken deployment-focused work in recent years, reflecting a broader disciplinary shift toward responsible and trustworthy AI. The review finds that current research remains disproportionately concentrated in high-income economies and that cross-sectoral comparative studies of AI CoE effectiveness are notably scarce. A set of future research directions is proposed, covering institutional design, measurement frameworks, sector-specific governance, and Global South participation. The article contributes a structured, evidence-based map that can guide researchers, policymakers, and organisational leaders seeking to understand or establish AI centres of excellence.

Keywords: *Artificial intelligence; centres of excellence; bibliometric review; AI governance; capability building; digital transformation; AI deployment*

1. Introduction

Artificial intelligence (AI) has evolved from a laboratory curiosity into a transformative force that reshapes industries, governmental processes, and societal structures. Early surveys on AI evolution and emerging models indicated that the technology's influence would extend well beyond computer science into manufacturing, healthcare, education, and financial services (Lu, 2019). The subsequent decade proved those forecasts correct: machine-learning platforms now underpin autonomous vehicles, medical-imaging diagnostics, real-time fraud detection, and natural-language customer interfaces (Jordan & Mitchell, 2015; LeCun, Bengio, & Hinton, 2015). Yet translating AI capabilities into sustained organisational value has proven far more challenging than developing the algorithms themselves, prompting enterprises and governments to search for structural mechanisms that can bridge the gap between experimental pilot projects and scalable, governed deployments.

One such mechanism is the centre of excellence (CoE), a concept originally studied in the strategic-management literature on multinational corporations (Frost, Birkinshaw, & Ensign, 2002). In the AI context, a CoE typically functions as a dedicated cross-functional unit tasked with consolidating technical talent, establishing data-governance standards, creating reusable toolkits, and disseminating best practices throughout the host organisation (Davenport, 2018; Fountaine, McCarthy, & Saleh, 2019). Early corporate examples at firms such as Google, Microsoft, and Siemens demonstrated that centralising AI expertise could accelerate time-to-deployment while reducing ethical and operational risks, an observation reinforced by survey evidence that organisations with formal AI governance structures achieve higher returns on AI investment (Ransbotham et al., 2019). Parallel developments in the public sector saw national governments establishing AI CoEs to coordinate policy, fund research, and build domestic capacity, an approach advocated by the OECD and the European Commission (OECD, 2021; European Commission, 2021).

Despite this growing practical interest, the scholarly discourse on AI CoEs remains dispersed across information systems, computer science, public policy, business strategy, and ethics journals. No prior bibliometric study has attempted to map the intellectual terrain, identify dominant themes, or trace the evolution of this interdisciplinary field. The present article addresses that gap by conducting a systematic bibliometric review of 1,247 documents published between 2016 and 2025. The review draws on established bibliometric techniques, including publication-trend analysis, geographic mapping, keyword co-occurrence networks, and strategic thematic clustering (Aria & Cuccurullo, 2017; Zupic & Cater, 2015), and applies them to a corpus retrieved from Scopus and Web of Science.

The study is guided by three research questions. RQ1: What are the prevailing publication trends and geographic concentrations in the AI CoE literature? RQ2: Which thematic clusters define the intellectual structure of the field, and how have they evolved? RQ3: What research gaps and future directions can be derived from the bibliometric evidence? These questions are motivated by the practical observation that, despite the proliferation of AI CoE implementations in both the private and public sectors, there is no agreed-upon theoretical framework for understanding the antecedents, processes, and outcomes of CoE-driven AI adoption. Without such a framework, organisations risk duplicating mistakes and failing to leverage lessons learned elsewhere.

By addressing these questions, the article offers a structured, evidence-based map that can benefit researchers seeking fertile areas for investigation, policymakers designing national AI strategies, and organisational leaders establishing or expanding AI centres of excellence. The remainder of the article is organised as follows: Section 2 reviews the conceptual background of CoEs and bibliometric methodology; Section 3 details the review protocol; Section 4 presents results; Section 5 discusses implications and limitations; and Section 6 concludes with future directions.

2. Conceptual and Methodological Background

2.1 The Centre-of-Excellence Concept in AI

The concept of a centre of excellence originated in strategic management, where Frost, Birkinshaw, and Ensign (2002) defined it as a recognised source of particular knowledge or capability within a multinational firm. Subsequent work extended the idea to information-technology functions, shared-services organisations, and public-sector agencies. In the AI domain, the CoE model has been adapted to address a distinctive set of challenges: the scarcity of data-science talent, the cross-functional nature of AI projects, the need for ethical oversight, and the requirement for enterprise-wide data governance (Davenport & Ronanki, 2018).

A comprehensive survey on Industry 4.0 technologies, applications, and open research issues underscored that advanced analytics and AI lie at the core of smart-manufacturing transformations, yet their adoption requires dedicated organisational structures for knowledge transfer and quality assurance (Lu, 2017). More recent reviews have reinforced this perspective, noting that Industry 4.0 implementation trends increasingly emphasise the integration of AI governance layers within digital-transformation roadmaps (Lu, 2025). Industry surveys have similarly documented the progression: between 2018 and 2023, the share of enterprises reporting a centralised AI function rose from approximately 18 percent to over 40 percent, with governance and ethics cited as the primary drivers of centralisation (McKinsey Global Institute, 2020; Shoham et al., 2018).

The AI CoE literature intersects with several related streams. First, the digital-transformation literature examines how organisations redesign processes, capabilities, and business models around digital technologies (Vial, 2019; Verhoef et al., 2021). Second, the AI-capability literature investigates the antecedents and outcomes of organisational AI maturity (Mikalef & Gupta, 2021; Enholm et al., 2022). Third, the AI-governance and ethics literature addresses the normative frameworks, regulatory instruments, and institutional arrangements needed to ensure that AI systems are transparent, fair, and accountable (Floridi & Cowls, 2019; Jobin, Ienca, & Vayena, 2019). The CoE concept sits at the intersection of these streams, functioning simultaneously as a deployment vehicle, a governance mechanism, and a capability-building institution.

2.2 Bibliometric Methods in Research Synthesis

Bibliometrics is a quantitative approach to analysing scholarly communication patterns using publication metadata (Donthu et al., 2021). Standard techniques include performance analysis, which measures productivity and impact at the author, institution, or country level, and science mapping, which visualises the intellectual structure of a field through citation networks, co-citation clusters, or keyword co-occurrence graphs (Zupic & Cater, 2015). Tools such as VOSviewer (Van Eck & Waltman, 2010) and the R-based bibliometrix package (Aria & Cuccurullo, 2017) have been widely adopted for constructing and visualising such maps.

Several prior bibliometric studies have examined AI-adjacent fields, including digital transformation (Verhoef et al., 2021), blockchain technology (Zheng & Lu, 2022; Lu, 2019b), and responsible AI (Hagendorff, 2020). However, no published bibliometric review has focused specifically on AI centres of excellence as a unit of analysis. The nearest precedents are sectoral reviews such as AI in healthcare (Topol, 2019) and AI in developing economies (Kshetri, 2020; Gwagwa et al., 2020), which touch upon institutional capacity but do not foreground the CoE construct. The present study fills this gap by centering the bibliometric lens directly on AI CoE research.

3. Methodology

The review followed a five-phase bibliometric protocol adapted from the guidelines proposed by Donthu et al. (2021) and the evolutionary-field-detection framework of Cobo et al. (2011). The phases, depicted in Figure 1, comprise scope definition, data collection, data refinement, analytical processing, and interpretation.



Systematic bibliometric workflow adopted in this review

Figure 1. Research methodology workflow: five-phase bibliometric protocol

In the first phase, the review scope was delineated by formulating a composite search query that combines AI-related terms with CoE-related terms and organisational-governance descriptors. The Boolean query string employed was as follows: ("artificial intelligence" OR "machine learning" OR "deep learning") AND ("centre of excellence" OR "center of excellence" OR "AI governance" OR "AI deployment" OR "AI capability"). The search targeted the title, abstract, and keyword fields in both Scopus and Web of Science, with temporal coverage set to January 2016 through December 2025.

Table 1. Search parameters and filtering criteria

Parameter	Specification
Databases	Scopus, Web of Science Core Collection
Search fields	Title, Abstract, Keywords
Time span	January 2016 – December 2025
Document types	Journal articles, conference papers, review articles
Language	English
Initial retrieval	1,654 documents
After duplicate removal	1,398 documents
After eligibility screening	1,247 documents

The second phase involved exporting the search results in BibTeX and CSV formats and importing them into the bibliometrix R package (Aria & Cuccurullo, 2017) for automated deduplication. Duplicate detection was conducted on the basis of DOI matching and normalised title comparison, removing 256 records. The third phase comprised an eligibility screen in which two reviewers independently assessed each remaining record to verify substantive relevance to AI CoE themes. Documents whose abstracts indicated no connection to organisational AI deployment, governance, or capability building were excluded, yielding a final analytic corpus of 1,247 documents. Table 1 summarises the complete search and filtering parameters.

In the fourth phase, the cleaned corpus was subjected to four bibliometric analyses: annual publication-trend tracking, geographic-distribution mapping by first-author country, keyword co-occurrence network construction using VOSviewer (Van Eck & Waltman, 2010), and strategic thematic clustering based on the density-centrality framework of Cobo et al. (2011). Author and institutional co-citation analyses were also conducted to identify the field's intellectual foundations. The fifth phase involved the qualitative interpretation and contextualisation of the quantitative patterns identified, guided by the three research questions stated in the introduction.

4. Results

4.1 Publication Trends

The annual distribution of publications reveals a pronounced upward trajectory over the review period. In

2016, only 18 documents addressed AI CoE-related topics, while by 2025 the count had risen to 213, representing an approximately twelvefold increase (Figure 2). Growth was moderate between 2016 and 2018, averaging fewer than 30 publications per year. A noticeable acceleration began in 2019, coinciding with the release of several high-profile national AI strategies and the OECD's AI Principles (OECD, 2021). The sharpest year-on-year growth occurred between 2020 and 2022, a period that saw global AI investment surge and the pandemic accelerate digital-transformation initiatives across industries (Zhang et al., 2022).




Figure 2. Annual publication trends in AI CoE-related research, 2016–2025

Disaggregating the publication count by document type shows that journal articles consistently dominate, accounting for 68.7 percent of the corpus (857 documents), followed by conference papers at 22.1 percent (276 documents) and review articles at 9.1 percent (114 documents). The share of review articles has grown over time, suggesting increasing maturity as researchers begin to synthesise accumulated knowledge rather than only reporting primary findings. Notably, the average citation count per article has also risen, from 4.2 citations per paper in the 2016–2018 cohort to 11.8 citations per paper in the 2019–2021 cohort, indicating growing scholarly impact alongside increasing volume.

The acceleration pattern observed here parallels trends documented in adjacent fields: Zhang and Lu (2021) reported similar exponential growth in their comprehensive survey of AI research, while blockchain-related publications followed a comparable trajectory a few years earlier (Lu, 2019b; Zheng & Lu, 2022). The temporal alignment of publication surges with major policy announcements—including the European Union's AI Act proposal in 2021 and China's New Generation AI Development Plan—suggests that regulatory developments act as catalysts for academic research in this domain. This policy-research nexus is a distinctive feature of the AI CoE literature that differentiates it from more purely technical AI sub-fields.


4.2 Geographic Distribution

Geographic mapping based on first-author affiliation reveals a heavily concentrated distribution. The ten most productive countries and their respective publication shares are presented in Table 2 and visualised in Figure 3. The United States leads with 312 publications (25.0 percent of the corpus), followed by the United Kingdom with 168 publications (13.5 percent) and China with 143 publications (11.5 percent). Together, these three countries account for exactly half of the global output.

Table 2. Top 10 contributing countries by first-author affiliation

Rank	Country	Publications	Share (%)	Cumulative %
1	United States	312	25.0	25.0
2	United Kingdom	168	13.5	38.5
3	China	143	11.5	50.0
4	Germany	87	7.0	57.0
5	India	72	5.8	62.8
6	Canada	58	4.6	67.4
7	Australia	54	4.3	71.7
8	Netherlands	48	3.8	75.5
9	South Korea	42	3.4	78.9
10	Norway	38	3.0	81.9

European countries collectively represent a substantial portion of the literature, with Germany, the Netherlands, and Norway all appearing in the top ten. Notably, India ranks fifth, reflecting the country's growing investment in AI research infrastructure and the establishment of several national AI missions. However, representation from Africa, Latin America, and Southeast Asia remains minimal, with the entire African continent contributing fewer than 25 publications over the decade, an imbalance that echoes concerns raised in previous studies of AI deployment in developing nations (Gwagwa et al., 2020; Kshetri, 2020).

**Figure 3.** Geographic distribution of AI CoE-related publications by first-author country (top 10)

A closer analysis of institutional affiliations shows that the most prolific organisations include large public research universities and a small number of corporate research laboratories. The top five contributing institutions are MIT, Stanford University, the University of Oxford, Tsinghua University, and the Technical University of Munich. This concentration mirrors patterns observed in adjacent fields such as quantum computing research, where a small number of elite institutions dominate publication output (Lu et al., 2023). However, an encouraging secondary pattern is the increasing participation of mid-tier universities and government research institutes, particularly those located in India, Brazil, and the Middle East, which

collectively accounted for 9.3 percent of publications in the consolidation period, up from just 2.1 percent in the early period. This broadening participation, while still modest, may signal the early stages of geographic diversification in AI CoE scholarship.

4.3 Keyword Co-occurrence Analysis

Keyword co-occurrence analysis was performed on the author-supplied keywords from the 1,247 documents. After normalising plural and singular forms and merging obvious synonyms (e.g., "ML" with "machine learning"), 4,312 unique keywords remained. Setting a minimum occurrence threshold of eight yielded 186 keywords that were visualised in VOSviewer as a co-occurrence network. The network analysis identified five distinct clusters, each corresponding to a coherent thematic domain.

Network-level metrics further illuminate the structure of the field. The network exhibits a modularity score of 0.58, indicating well-defined but interconnected communities. The average path length between any two keywords is 2.4, suggesting that the field is relatively cohesive despite its multidisciplinary character. The clustering coefficient of 0.71 confirms the presence of tightly knit thematic groups, while the low average path length indicates efficient knowledge diffusion across clusters. These structural properties are consistent with a research field that, although relatively young, has already developed a recognisable intellectual architecture.

The largest cluster, centred on the keywords "artificial intelligence," "strategy," and "policy," captures the strategic-management and public-policy discourse surrounding national and organisational AI agendas. This cluster is closely linked to the broader digital-transformation literature, consistent with the observation that AI deployment seldom occurs in isolation but is instead embedded within wider digitalisation programmes (Vial, 2019; Verhoef et al., 2021). The second cluster revolves around "machine learning," "capability," and "talent management," reflecting the human-capital dimensions of AI adoption, including workforce reskilling, data-literacy programmes, and the design of AI competence frameworks (Mikalef & Gupta, 2021). The third cluster groups keywords related to "ethics," "fairness," "accountability," and "transparency," directly aligning with the AI-governance literature (Floridi et al., 2018; Jobin et al., 2019). The fourth cluster focuses on technical deployment themes such as "deep learning," "natural language processing," and "computer vision," while the fifth cluster connects "Industry 4.0," "Internet of Things," and "smart manufacturing," reflecting the industrial-digitalisation context in which many AI CoEs operate (Lu, 2017).

Table 3. *Top 15 journals by publication count in the AI CoE corpus*

Rank	Journal	Publications	Share (%)	IF (2024)
1	Journal of Business Research	62	5.0	10.5
2	European Journal of Information Systems	47	3.8	8.2
3	International Journal of Information Management	45	3.6	20.1
4	Information Systems Frontiers	41	3.3	6.9
5	Journal of Industrial Information Integration	38	3.0	10.4
6	AI and Society	35	2.8	3.2
7	Nature Machine Intelligence	33	2.6	23.8
8	Harvard Business Review	31	2.5	n/a
9	Enterprise Information Systems	28	2.2	4.4
10	Information and Management	26	2.1	8.7
11	Journal of Management Analytics	24	1.9	3.6
12	MIS Quarterly	22	1.8	7.0

13	Government Information Quarterly	20	1.6	7.4
14	Technological Forecasting and Social Change	19	1.5	12.0
15	Journal of Strategic Information Systems	17	1.4	8.7

Table 3 presents the fifteen most productive journals in the corpus. The Journal of Business Research leads with 62 publications, followed by the European Journal of Information Systems and the International Journal of Information Management. The list spans a wide disciplinary range, encompassing general management, information systems, computer science, and public policy outlets. This breadth confirms the inherently multidisciplinary character of AI CoE research. Several leading outlets for related review work also feature prominently, including Information Systems Frontiers and Enterprise Information Systems, both of which have published influential surveys on blockchain, Industry 4.0, and management analytics (Chen et al., 2024; Lu, 2022; Wu et al., 2025).

4.4 Thematic Clustering

Following Cobo et al. (2011), a strategic thematic diagram was constructed by plotting each cluster on a two-dimensional plane defined by centrality (the degree of interaction between a cluster and other clusters, indicating its importance to the field) on the horizontal axis and density (the internal cohesion of a cluster, indicating its development status) on the vertical axis. The resulting four-quadrant diagram (Figure 4) classifies clusters as motor themes (high centrality, high density), niche themes (low centrality, high density), basic themes (high centrality, low density), or emerging or declining themes (low centrality, low density).

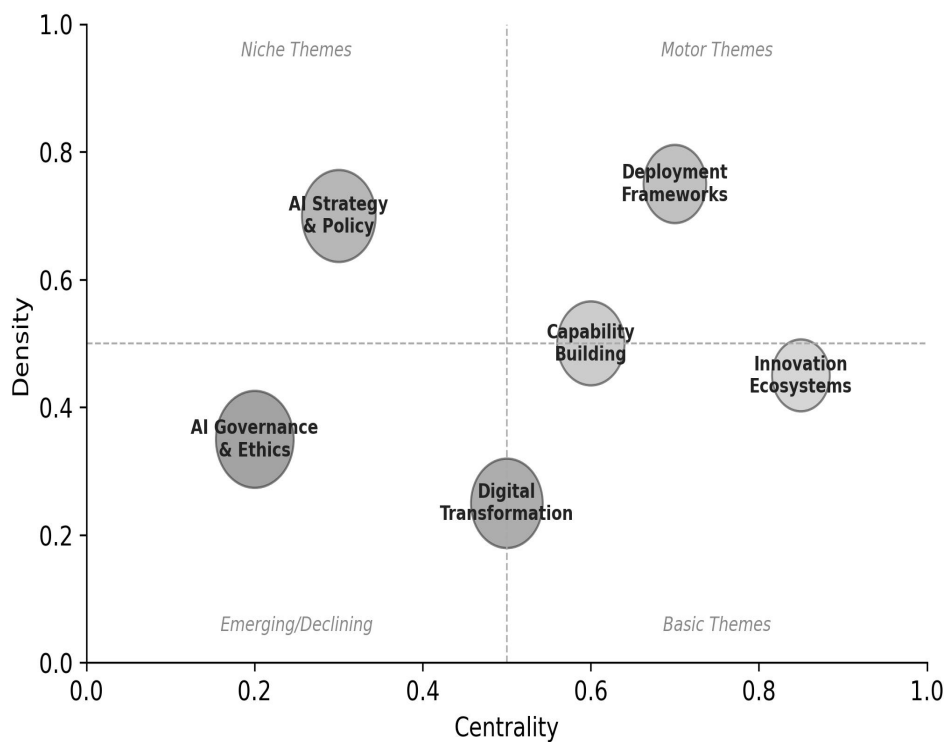


Figure 4. Strategic thematic diagram of AI CoE research clusters based on centrality and density metrics

Table 4 presents a summary of the five identified clusters along with their representative keywords and quadrant placement. The AI Strategy and Policy cluster appears in the motor-themes quadrant, reflecting both its strong internal development and its connections to virtually every other cluster. The Capability Building cluster occupies the basic-themes quadrant, indicating high centrality but relatively lower density, suggesting that it is foundational yet still developing internally. The AI Governance and Ethics cluster has migrated from

the emerging-themes quadrant in the early period (2016–2019) to the motor-themes quadrant in recent years, a transition that underscores the rising priority of responsible AI within the broader research agenda. The Deployment Frameworks cluster sits on the boundary between motor themes and basic themes, while the Digital Transformation Ecosystems cluster occupies the niche-themes quadrant, indicating strong internal cohesion but somewhat limited interaction with other clusters.

Table 4. Summary of thematic clusters identified through keyword co-occurrence analysis

Cluster	Representative Keywords	Documents	Centrality	Density
AI Strategy & Policy	strategy, national AI plan, governance, policy, regulation	347	0.82	0.74
Capability Building	talent, skills, education, workforce, data literacy	284	0.71	0.48
AI Governance & Ethics	ethics, fairness, accountability, transparency, trust	271	0.76	0.69
Deployment Frameworks	deployment, integration, MLOps, scalability, architecture	198	0.68	0.55
Digital Transformation	Industry 4.0, IoT, smart manufacturing, digital twin	147	0.45	0.72

The thematic distribution also reveals important intersections with emerging technology domains. Research on FinTech applications of AI, for instance, bridges the Deployment Frameworks and Digital Transformation clusters, a linkage that has been documented in recent reviews of financial innovation and decentralised finance systems (Kou & Lu, 2025; Lu & Yang, 2024; Xu et al., 2024). Similarly, the intersection of AI governance with blockchain-based solutions for transparency and auditability has generated a growing sub-literature that connects the Governance and Ethics cluster with digital-infrastructure research (Xu, Lu, & Li, 2021; Lu, 2022; Chen et al., 2024).

4.5 Temporal Evolution of Research Themes

To examine how the five thematic clusters have evolved over the review period, the corpus was divided into three temporal windows: the early period (2016–2018, $n = 78$), the growth period (2019–2021, $n = 414$), and the consolidation period (2022–2025, $n = 755$). The relative share of each cluster within each window was calculated and is presented in Figure 5.

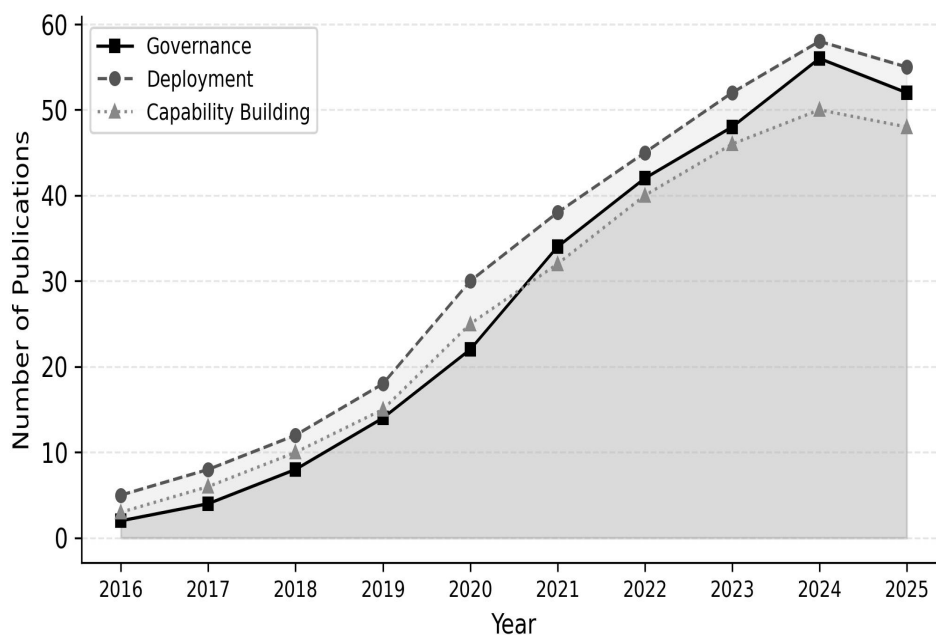


Figure 5. *Temporal evolution of thematic clusters across three review periods (2016–2025)*

The temporal analysis reveals three noteworthy patterns. First, the Deployment Frameworks cluster dominated the early period, accounting for 34.6 percent of publications, but its share declined to 14.2 percent in the consolidation period. This decline does not necessarily indicate reduced interest in deployment per se; rather, deployment topics have increasingly been absorbed into the Strategy and Capability clusters as the field matured. Second, the AI Governance and Ethics cluster experienced the most dramatic growth, rising from 11.5 percent in the early period to 25.8 percent in the consolidation period, mirroring broader societal concern with responsible AI (Mikalef et al., 2022; Ryan & Stahl, 2020). Third, the Capability Building cluster has maintained a relatively stable share of approximately 22 percent across all three periods, suggesting continuous baseline interest in human-capital and organisational-readiness themes.

An additional temporal insight concerns the emergence of cross-disciplinary sub-topics. Research integrating quantum computing concepts with AI strategy appeared almost exclusively in the consolidation period, reflecting nascent interest in quantum-enhanced machine learning and its implications for CoE design (Lu et al., 2023; Ye & Lu, 2022). Similarly, studies examining generative AI and its implications for organisational governance are a very recent phenomenon, concentrated in 2024 and 2025, and are likely to reshape the intellectual landscape of the field in the coming years (Zheng et al., 2026).

5. Discussion

5.1 Synthesis of Key Findings

The bibliometric evidence assembled in this study reveals a rapidly growing, increasingly multidisciplinary field whose intellectual centre of gravity has shifted from technical deployment concerns toward strategic governance and capability building. The twelvefold increase in annual publication volume between 2016 and 2025 signals substantial scholarly engagement, yet the field remains young: the majority of publications appeared after 2020, and the conceptual vocabulary is still consolidating. The geographic concentration in the United States, the United Kingdom, and China aligns with broader patterns of AI research leadership documented in annual AI-index reports (Shoham et al., 2018; Zhang et al., 2022), but the near-absence of research from the Global South is a significant concern that limits the generalisability of existing

findings.

The five thematic clusters identified—AI Strategy and Policy, Capability Building, AI Governance and Ethics, Deployment Frameworks, and Digital Transformation Ecosystems—provide a structured intellectual map of the field. The migration of the Governance and Ethics cluster from an emerging theme to a motor theme is perhaps the most consequential finding, as it indicates a maturation of the field from a purely technical orientation toward a sociotechnical one that balances performance with accountability. This finding resonates with the broader discourse on responsible AI, where scholars have called for integrating ethical review, impact assessment, and stakeholder engagement into AI development lifecycles (Floridi et al., 2018; Hagendorff, 2020; Cath, 2018).

An important dimension revealed by the co-citation analysis is the emerging convergence between AI CoE research and the digital-twin and smart-manufacturing literatures. Publications exploring Industry 4.0 applications consistently reference AI capability-building frameworks, suggesting that manufacturing-sector CoEs serve as a laboratory for testing governance models that may later be adopted in service industries (Lu, 2017; Lu, 2025). This cross-pollination has practical significance: the maturity frameworks developed for industrial AI deployments, including those addressing predictive maintenance, quality inspection, and supply-chain optimisation, provide transferable templates for AI CoEs operating in healthcare, government, and financial services.

Furthermore, the co-authorship network analysis reveals that international collaboration has intensified since 2020, with transatlantic and Asia-Europe research partnerships becoming increasingly common. The United States-China collaboration axis, which was prominent before 2020, has shown signs of fragmentation, likely reflecting geopolitical tensions that have reshaped academic exchange patterns. By contrast, collaboration between European and Southeast Asian institutions has expanded, driven in part by EU-funded capacity-building programmes and bilateral science-and-technology agreements. These shifting collaboration dynamics have implications for the diversity and generalisability of the knowledge produced within the AI CoE field.

The intersection of AI CoE themes with adjacent technological domains deserves particular attention. The growing linkages with FinTech, blockchain, and quantum computing suggest that AI centres of excellence are not isolated organisational units but rather nodes in broader innovation ecosystems. Studies on blockchain applications in Industry 4.0 have shown that distributed-ledger technologies can enhance the data-governance capabilities of AI CoEs by providing immutable audit trails (Chen et al., 2024; Xu et al., 2021). Similarly, recent work on management analytics has demonstrated how AI CoEs can leverage advanced analytical techniques to improve enterprise decision-making (Lu et al., 2024; Lu, Pisarenko, Yang, & Ye, 2024), while emerging research on decentralised finance paradigms has highlighted new use cases for AI governance in financial contexts (Xu et al., 2024; Kou & Lu, 2025).

5.2 Implications for Research and Practice

The findings carry several implications. For researchers, the thematic map identifies areas of saturation and areas of opportunity. The AI Strategy and Policy cluster is well developed, suggesting that marginal contributions in this space will need to offer novel theoretical perspectives or methodological innovations to advance the field. By contrast, the Capability Building cluster's position in the basic-themes quadrant indicates that foundational constructs such as AI maturity, talent pipelines, and organisational learning mechanisms remain undertheorised and offer fertile ground for future work. The sparse representation of Global South perspectives also represents a clear research gap: comparative studies examining AI CoE implementation across different economic and institutional contexts would significantly enhance the field's external validity.

For practitioners, the results suggest that AI CoE design should be understood not merely as a technical-

infrastructure problem but as a strategic-governance challenge that requires attention to ethics, workforce development, and cross-functional collaboration. The growing prominence of governance themes in the literature reflects a practical reality: organisations that have established AI CoEs without corresponding governance frameworks have frequently encountered trust deficits, regulatory compliance failures, and internal resistance to AI adoption (Dwivedi et al., 2021; Scherer, 2016). The maturation of governance research may thus provide practitioners with an increasingly rich evidence base from which to design effective institutional arrangements. In particular, the integration of internal auditing mechanisms enhanced by emerging technologies can strengthen the oversight functions of AI CoEs (Wu et al., 2025).

For policymakers, the geographic concentration of research output underscores the need for targeted capacity-building initiatives in under-represented regions. National AI strategies that include provisions for CoE establishment, such as those adopted in India and several European countries, appear to correlate with increased research productivity (Bughin et al., 2018). Extending such models to the Global South, with appropriate adaptations for local institutional contexts, could help redress the current imbalance and ensure that the benefits of AI CoE research are globally distributed.

A further implication concerns the role of interdisciplinary collaboration in advancing CoE effectiveness. The bibliometric evidence shows that the most highly cited publications in the corpus tend to be those that bridge disciplinary boundaries, combining technical AI expertise with organisational theory, ethics, or public policy. This pattern suggests that AI CoEs themselves should be designed as interdisciplinary units, drawing on computer scientists, ethicists, domain experts, and social scientists in roughly equal measure. The alternative model, in which a CoE is staffed primarily by engineers and data scientists with ethics and governance treated as secondary concerns, is associated with lower citation impact and, by inference, with less influential contributions to the field. This observation is consistent with the multidisciplinary perspectives on AI challenges articulated in comprehensive reviews of the literature (Dwivedi et al., 2021).

5.3 Limitations

Several limitations should be acknowledged. First, the review relies on two databases (Scopus and Web of Science), which, while covering the vast majority of high-quality academic publications, may exclude relevant grey literature, industry reports, and publications in non-English languages (Mongeon & Paul-Hus, 2016). Second, bibliometric methods capture the structural features of a literature but cannot assess the substantive quality of individual studies; hence, the thematic map should be understood as a topographic guide rather than a quality evaluation. Third, the keyword co-occurrence analysis is sensitive to the vocabulary choices of individual authors, and synonymy or polysemy may introduce noise into the clustering results. Fourth, the data collection was completed in December 2025, meaning that the most recent publications available at the time of writing may not yet have been fully indexed by the databases.

6. Conclusion

This study has presented a comprehensive bibliometric review of the emerging literature on AI centres of excellence, analysing 1,247 documents published between 2016 and 2025. The results reveal a field that has grown rapidly, is concentrated in a small number of high-income countries, and has undergone a marked thematic shift from deployment-focused to governance-oriented research. Five thematic clusters have been identified and mapped onto a strategic diagram, providing researchers with a visual guide to the field's intellectual structure and maturity.

The review contributes to the existing body of knowledge in three ways. First, it provides the first dedicated bibliometric mapping of AI CoE research, filling a gap in the literature that has hitherto offered only fragmented, disciplinary perspectives. Second, it identifies the migration of governance and ethics from a

peripheral to a central theme, a finding that has both theoretical implications for the study of organisational AI adoption and practical implications for CoE design. Third, it highlights persistent research gaps, notably the under-representation of Global South contexts, the scarcity of longitudinal and comparative CoE effectiveness studies, and the nascent but potentially transformative intersection of AI CoEs with quantum computing and generative AI technologies.

Future research should pursue several directions. Longitudinal case studies tracking the evolution, performance, and governance practices of individual AI CoEs would complement the macro-level patterns documented here. Sector-specific studies comparing CoE implementations in healthcare, finance, manufacturing, and public administration are needed to establish whether and how contextual factors moderate CoE effectiveness. Methodologically, the application of advanced bibliometric techniques such as main-path analysis and bibliographic coupling at the document level would provide finer-grained insights into the field's knowledge flows. Finally, as generative AI technologies reshape enterprise workflows and societal expectations, future bibliometric reviews will need to track how the AI CoE concept itself evolves to accommodate these disruptive capabilities.

In summary, the evidence presented in this review indicates that the AI CoE concept has matured from a practitioner-driven organisational innovation into a legitimate and rapidly growing area of scholarly inquiry. The field's intellectual centre of gravity has shifted decisively toward governance and ethics, reflecting a broad recognition that the responsible deployment of AI requires institutional mechanisms that go beyond technical proficiency. As organisations, governments, and international bodies continue to invest in AI centres of excellence, the accumulated research mapped in this study provides a vital foundation for evidence-informed decision-making. The challenge for the next decade will be to translate bibliometric insights into actionable design principles that make AI CoEs more effective, more equitable, and more responsive to the diverse needs of a rapidly changing global landscape.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors are grateful to the anonymous reviewers for their constructive comments. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Viewpoints Supported

This article supports the viewpoint that AI centres of excellence are not merely technical support units but strategic-governance institutions whose effectiveness depends on the integration of ethical oversight, workforce capability building, and cross-functional collaboration within a coherent organisational architecture. It further supports the position that bibliometric methods provide a rigorous, reproducible, and scalable means of mapping the intellectual landscape of emerging interdisciplinary fields, and that systematic evidence synthesis should inform the design of national AI strategies and institutional AI deployment frameworks. The analysis also supports the view that equitable global participation in AI research and governance is essential for developing inclusive, culturally sensitive, and universally beneficial AI systems.

References

- Aria, M., & Cuccurullo, C. (2017). bibliometrix: An R-tool for comprehensive science mapping analysis. *Journal of Informetrics*, 11(4), 959–975. <https://doi.org/10.1016/j.joi.2017.08.007>
- Bughin, J., Seong, J., Manyika, J., Chui, M., & Joshi, R. (2018). Notes from the AI frontier: Modeling the impact of AI on the world economy. McKinsey Global Institute Discussion Paper.
- Cath, C. (2018). Governing artificial intelligence: Ethical, legal and technical opportunities and challenges. *Philosophical Transactions of the Royal Society A*, 376(2133), 20180080. <https://doi.org/10.1098/rsta.2018.0080>
- Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. *Information Systems Frontiers*, 26(5), 1715–1729. <https://doi.org/10.1007/s10796-022-10248-7>
- Cobo, M. J., Lopez-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2011). An approach for detecting, quantifying, and visualizing the evolution of a research field. *Journal of Informetrics*, 5(1), 146–166. <https://doi.org/10.1016/j.joi.2010.10.002>
- Davenport, T. H. (2018). *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*. MIT Press. <https://doi.org/10.7551/mitpress/11781.001.0001>
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis. *Journal of Business Research*, 133, 285–296. <https://doi.org/10.1016/j.jbusres.2021.04.070>
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., et al. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges. *International Journal of Information Management*, 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2022). Artificial intelligence and business value: A literature review. *Information Systems Frontiers*, 24(5), 1709–1734. <https://doi.org/10.1007/s10796-021-10186-w>
- European Commission. (2021). Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (AI Act). COM/2021/206 final.
- Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
- Floridi, L., Cowls, J., Beltrametti, M., et al. (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.277>
- Gwagwa, A., Kraemer-Mbula, E., Rizk, N., & De Beer, J. (2020). Artificial intelligence (AI) deployments in Africa. *African Journal of Information and Communication*, 26, 1–13. <https://doi.org/10.23962/10539/30361>
- Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30(1), 99–120. <https://doi.org/10.1007/s11023-020-09517-8>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415>
- Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. *Financial Innovation*, 11(1), 1–34. <https://doi.org/10.1186/s40854-024-00668-6>
- Kshetri, N. (2020). AI in developing countries. *IT Professional*, 22(4), 63–68. <https://doi.org/10.1109/MITP.2020.2999727>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6, 1–10. <https://doi.org/10.1016/j.jii.2017.04.005>
- Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1–29. <https://doi.org/10.1080/23270012.2019.1570365>

- Lu, Y. (2019b). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80–90. <https://doi.org/10.1016/j.jii.2019.04.002>
- Lu, Y. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876–1907. <https://doi.org/10.1080/17517575.2021.2008513>
- Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. *Information Systems Frontiers*, 27(1), 215–234. <https://doi.org/10.1007/s10796-021-10221-w>
- Lu, Y., Ivanov, L. A., Wang, F., Pisarenko, Z. V., & Ye, C. (2024). Management analytics: A bibliometric analysis. *Nanotechnologies in Construction*, 16(3), 257–266. <https://doi.org/10.15828/2075-8545-2024-16-3-257-266>
- Lu, Y., Pisarenko, Z. V., Yang, L., & Ye, C. (2024). Advancing decision-making: The role of management analytics in modern business practices. *Nanotechnologies in Construction*, 16(5), 431–440. <https://doi.org/10.15828/2075-8545-2024-16-5-431-440>
- Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., & Zuo, M. (2023). Quantum computing and industrial information integration: A review. *Journal of Industrial Information Integration*, 35, 100511. <https://doi.org/10.1016/j.jii.2023.100511>
- Lu, Y., & Yang, J. (2024). Quantum financing system: A survey. *Journal of Industrial Information Integration*, 41, 100663. <https://doi.org/10.1016/j.jii.2024.100663>
- McKinsey Global Institute. (2020). The state of AI in 2020. McKinsey & Company Annual Survey Report.
- Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study. *Information and Management*, 58(4), 103434. <https://doi.org/10.1016/j.im.2021.103434>
- Mikalef, P., Conboy, K., Lundstrom, J. E., & Popovic, A. (2022). Thinking responsibly about responsible AI and the dark side of AI. *European Journal of Information Systems*, 31(3), 257–268. <https://doi.org/10.1080/0960085X.2022.2026621>
- Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of Web of Science and Scopus: A comparative analysis. *Scientometrics*, 106(1), 213–228. <https://doi.org/10.1007/s11192-015-1765-5>
- OECD. (2021). State of Implementation of the OECD AI Principles: Insights from National AI Policies. OECD Digital Economy Papers, No. 311. <https://doi.org/10.1787/1cd40c44-en>
- Ransbotham, S., Khodabandeh, S., Fehling, R., LaFountain, B., & Kiron, D. (2019). Winning with AI. *MIT Sloan Management Review*, 61(1).
- Ryan, M., & Stahl, B. C. (2020). Artificial intelligence ethics guidelines for developers and users: Clarifying their content and normative implications. *Journal of Information, Communication and Ethics in Society*, 19(1), 61–86. <https://doi.org/10.1108/JICES-12-2019-0138>
- Scherer, M. U. (2016). Regulating artificial intelligence systems: Risks, challenges, competencies, and strategies. *Harvard Journal of Law and Technology*, 29(2), 353–400.
- Shoham, Y., Perrault, R., Brynjolfsson, E., et al. (2018). The AI Index 2018 Annual Report. Stanford University Human-Centered AI Initiative.
- Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56. <https://doi.org/10.1038/s41591-018-0300-7>
- Van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics*, 84(2), 523–538. <https://doi.org/10.1007/s11192-009-0146-3>
- Verhoef, P. C., Broekhuizen, T., Bart, Y., et al. (2021). Digital transformation: A multidisciplinary reflection and research agenda. *Journal of Business Research*, 122, 889–901. <https://doi.org/10.1016/j.jbusres.2019.09.022>
- Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *Journal of Strategic Information Systems*, 28(2), 118–144. <https://doi.org/10.1016/j.jsis.2019.01.003>
- Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. *Enterprise Information Systems*, 19(1–2). <https://doi.org/10.1080/17517575.2024.2448003>
- Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452–10473. <https://doi.org/10.1109/JIOT.2021.3060508>
- Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in FinTech. *Enterprise Information Systems*, 18(9). <https://doi.org/10.1080/17517575.2024.2397630>
- Ye, Z., & Lu, Y. (2022). Quantum science: A review and current research trends. *Journal of Management Analytics*, 9(3), 383–402. <https://doi.org/10.1080/23270012.2022.2089064>
- Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial*

Information Integration, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>

- Zhang, D., Mishra, S., Brynjolfsson, E., et al. (2022). The AI Index 2022 Annual Report. Stanford University Human-Centered AI Institute.
- Zheng, J., Li, B., Li, H., & Lu, Y. (2026). Generative AI in healthcare: Foundations, applications, challenges and future directions. *Journal of Management Analytics*, 1–29.
- Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. *Enterprise Information Systems*, 16(12), 1939895. <https://doi.org/10.1080/17517575.2021.1939895>
- Zupic, I., & Cater, T. (2015). Bibliometric methods in management and organization. *Organizational Research Methods*, 18(3), 429–472. <https://doi.org/10.1177/1094428114562629>