

From Adoption to Mastery: A Conceptual Framework for AI Integration in Academic Research

Bogdan NICA¹

DOI: 10.24818/mer/2026.02-11

ABSTRACT

By 2025, over 80% of academic researchers report using generative artificial intelligence tools (Wiley, 2025; HEPI, 2025; UNESCO, 2025), yet the translation of this adoption into mature, methodologically grounded research practice has proceeded markedly slower. McKinsey's State of AI 2025 (n = 1,993 organisations) captures the structural manifestation at the cross-sectoral level: 88% of organisations use AI, but only 1% describe deployments as mature. This paper conceptualises the structural gap between adoption and mastery – the adoption-mastery gap – as a distinct object of study in research management. The paper conducts a scoping review applying the PRISMA Extension for Scoping Reviews (PRISMA-ScR; Tricco et al., 2018), aligned with the PRISMA 2020 reporting framework (Page et al., 2021), on a curated analytical corpus of 66 sources complemented by 18 foundational theoretical and methodological anchors (88 total bibliography entries). Sources were identified through triangulated multi-agent retrieval across three independent generative search interfaces, anchored in The Lens database with verification through Crossref records and indexing status confirmed across scholarly databases (OpenAlex, The Lens, Crossref, WoS and Scopus). The paper offers four contributions: a theoretical conceptualisation of the adoption-mastery gap with the augmentation trap as its causal mechanism; a descriptive taxonomy of five mature AI-assisted research practices; a four-lever institutional framework for doctoral schools anchored in research administration literature; and a theoretical hypothesis – to be tested in subsequent empirical work – of East European structural isolation as a regional driver of the gap.

KEYWORDS: adoption-mastery gap, augmentation trap, generative artificial intelligence, research management, scoping review

JEL CLASSIFICATION: I23, O33, M15

1. INTRODUCTION

The diffusion of generative artificial intelligence tools across the academic research workforce has reached high adoption rates that exceed those reported only a year earlier. Wiley (2025) documents adoption rising from 57% in 2024 to 84% in 2025, with 62% reporting direct use for research and publication tasks. The Higher Education Policy Institute (2025) finds that 92% of UK undergraduates and doctoral students use generative AI in some form. UNESCO's global survey of 400 institutions in 90 countries (UNESCO, 2025) confirms the breadth of the phenomenon, and Mohammadi et al. (2026), surveying academics in twenty countries, identify PhD students and early-career researchers as the heaviest adopters.

However, adoption does not translate uniformly into productive practice. McKinsey (2025; n=1,993 organisations, 105 countries) reports that 88% of organisations use AI, but only 1% describe deployments as mature. Hao et al. (2026), analysing 41.3 million scientific papers, find that AI-assisted researchers publish 3.02 times more articles and accumulate 4.84 times more

¹ Bucharest University of Economic Studies, email: bogdancristiannica@mail.com

citations – yet collectively cover 4.63% less topical ground and exhibit 22% lower inter-researcher engagement, a pattern the authors term "*lonely crowds*". The paradox – individual acceleration coexisting with collective narrowing – is developed as the Evans paradox in Section 4.4.

This paper conceptualises the structural gap between rapid adoption and slower development of mature research practices, and proposes an institutional response for universities and doctoral schools. The adoption-mastery gap is defined as the structural lag between the rate at which a research workforce takes up generative AI tools and the rate at which methodologically robust practices anchored in those tools become normalised within disciplinary communities. The gap is characterised in the converging literature not as transitional, but as a persistent feature of the diffusion process, and it admits institutional remedies.

Four contributions to research management are advanced in what follows. First, extending the assimilation-gap tradition (Fichman & Kemerer, 1999), it conceptualises the adoption-mastery gap as a distinct object of study in research management, identifying the augmentation trap – formalised by Caosun and Aral (2026) and empirically reinforced by Bastani et al. (2025), Rinta-Kahila et al. (2023), and Neshenko and Ryall (2026) – as its causal mechanism. Second, it develops a descriptive taxonomy of five mature AI-assisted research practices derived from systematic analysis of the empirical corpus. Third, it proposes a four-lever institutional framework for doctoral schools – infrastructure, structured training, community of practice, and declaration policy – anchored in research administration literature. Fourth, it advances a theoretical hypothesis, to be tested empirically, of East European structural isolation as a regional driver of the gap, drawing on regional innovation systems theory.

The paper is organised as follows. Section 2 reviews the empirical landscape, theoretical foundations, and documented practices and barriers. Section 3 presents the objectives, hypothesis, and methodology. Section 4 reports the bibliometric and content analysis results, develops the taxonomy and institutional framework, and discusses the East European hypothesis. Section 5 concludes with the contributions, managerial implications, limitations, and future research directions.

2. LITERATURE REVIEW

2.1 Empirical landscape of generative AI adoption in academic research

The empirical evidence on adoption converges on a single observation: a majority of the research workforce uses generative AI, and a majority of institutions do not yet provide adequate methodological scaffolding. Wiley (2025; $n = 2,430$) reports that 80% of adopters use general-purpose tools such as ChatGPT, only 25% have engaged with research-specific platforms, and 57% cite the absence of institutional guidelines as a primary barrier. Mohammadi et al. (2026), across twenty countries, find that dominant uses remain translation, proofreading, and literature review, with markedly less penetration into data analysis or methodological design.

The Digital Education Council (2024; $n = 3,839$ students, sixteen countries) reports 86% regular AI use; the parallel faculty survey (2025; $n = 1,681$ staff, 28 countries) reports only 6% considering institutional support sufficient. UNESCO (2025) documents the same dynamic across doctoral training programmes, and the European University Association (2026) finds that while more than two-thirds of universities report widespread AI use among doctoral

students, fewer than 40% have developed governance infrastructure. The comparison with neighbouring sectors – software engineering (Stack Overflow, 2024; GitClear, 2025), marketing (CoSchedule, 2025) – is developed in Table 1 (Section 4.2); the central distinction is that academic research lacks the fast feedback mechanisms (code reviews, campaign ROI) that enable real-time calibration of AI reliance.

Bianchini et al. (2025) identify *AI-rich social capital* – the network of colleagues and mentors with usable AI experience – as the most consistent facilitator of productive uptake: institutions whose senior researchers do not themselves practice mature AI use cannot develop junior researchers capable of doing so. Radu et al. (2024), in a bibliometric analysis of 1,028 publications on AI in competency-based education, document post-2017 acceleration but find that the literature's growth has not translated into pedagogical maturity. The picture across these sources is consistent: adoption is widespread, productivity translation is unequal and skill-dependent, and institutional scaffolding lags behind individual use.

2.2 Theoretical foundations

Four theoretical traditions provide the scaffolding for this paper's argument.

First, Rogers' diffusion of innovations framework (Rogers, 2003) characterises technology uptake through five attributes – relative advantage, compatibility, complexity, trialability, and observability. Singh and Strzelecki (2025) apply this framework to generative AI in higher education, finding that observability and compatibility with academic integrity norms are the most contested attributes.

Second, the Technology Acceptance Model (Davis, 1989) and its UTAUT successor (Venkatesh et al., 2003) decompose adoption intention into Performance Expectancy, Effort Expectancy, Social Influence, and Facilitating Conditions. Chen et al. (2024) find that Performance Expectancy and Facilitating Conditions are the dominant predictors for research scholars; Cao et al. (2026), combining PLS-SEM with TAM-UTAUT, demonstrate that AI anxiety significantly moderates acceptance. Popa et al. (2024) apply Necessary Conditions Analysis to AI competencies, distinguishing necessary from sufficient conditions for effective use – empirical grounding for the adoption-mastery distinction developed in Section 4.

The closest theoretical antecedent of the adoption-mastery gap is Fichman and Kemerer's (1999) assimilation gap, which formalises the divergence between cumulative acquisition and deployment curves for new technologies, with their earlier work (Fichman & Kemerer, 1997) attributing the gap to knowledge barriers. Cohen and Levinthal's (1990) absorptive capacity provides the cognitive mechanism by which the gap is bridged, and Zahra and George (2002) extend the construct by distinguishing potential from realised absorptive capacity – mapping directly onto the adoption-mastery distinction.

The adoption-mastery gap extends the assimilation-gap construct on three axes not addressed in Fichman and Kemerer's original formulation: the unit of mastery is methodological rather than technical-deployment; the augmentation trap supplies a domain-specific causal mechanism producing stable rather than transitional divergence; and individual mastery trajectories can diverge from collective mastery trajectories – a possibility the assimilation-gap model does not anticipate and which the Evans paradox (Section 4.4) empirically demonstrates.

Burton-Jones and Straub (2006) reconceptualise system usage as a trajectory from vigilant engagement toward passive acceptance, anticipating the cognitive-offloading risk that contemporary capability-trap analyses formalise.

Third, Repenning and Sterman's capability traps framework (Repenning & Sterman, 2002; companion essay 2001) models how performance pressure produces short-term gains that mask capability erosion. Morton (2024) identifies three traps specific to generative AI; Brynjolfsson et al. (2018) provide the macroeconomic counterpart in the modern productivity paradox.

A notable recent contribution is Caosun and Aral (2026), whose formal dynamic model of the augmentation trap demonstrates that less-experienced workers deskill substantially while experienced workers retain capability, producing permanent skill divergence. Four independent lines of evidence converge on the same mechanism: Neshenko and Ryall (2026) establish the dynamic through a formal capability-trap model in Strategy Science; Rinta-Kahila et al. (2023) document vicious circles of skill erosion through a longitudinal case study in the Journal of the Association for Information Systems; Bastani et al. (2025) deliver experimental evidence in PNAS – students with unrestricted AI tutor access improved 48% during practice but underperformed controls by 17% after removal; and Gerlich (2025) demonstrates that frequent AI use correlates negatively with critical thinking, with cognitive offloading mediating the relationship.

Convergence across model-theoretic, longitudinal case-study, experimental, and correlational evidence strengthens the general cognitive-offloading claim while leaving exact transmission paths to doctoral-level expert knowledge work an open empirical question. The mechanism plausibly generalises – cognitive offloading operates independently of domain expertise level – but doctoral research involves longer feedback loops and higher-order methodological judgement, which may modulate both the speed and the reversibility of deskilling. The implications for doctoral training remain substantive: unchecked cognitive offloading may erode the methodological foundations on which future research capability depends.

Mollick's (2024) *co-intelligence* framework urges researchers to treat LLMs as fallible colleagues; Sakar and Nayak (2026) observe that it remains underspecified on institutional dynamics, motivating the institutional turn proposed here.

These traditions converge on a fourth: research administration and academic leadership. Trowler (2010) characterises the managerialist posture under which research administration operates; Whitchurch (2010) develops the concept of third-space professionals; Cooke (2001) supplies the regional innovation systems vocabulary for the East European hypothesis (Section 4.3).

2.3 Mature practices and barriers documented in the literature

A growing body of empirical work documents specific mature AI-assisted research practices – from large-language-model literature screening (Dai et al., 2025; Jayathilake et al., 2026) through prompt engineering as a documented capability (Anam, 2025; Susnjak, 2025) to multi-agent research systems (Lu et al., 2026; Guo et al., 2024; Pantiukhin et al., 2025; Sami et al., 2024). These practices are synthesised into the descriptive taxonomy reported in Section 4.2.

The practices are counterbalanced by barriers: Luo (2024) finds that policies at twenty world-leading universities concentrate on misconduct prevention rather than methodological

integration, with the chilling effect falling hardest on doctoral researchers. Bianchini et al. (2025) identify the social-capital mechanism by which barriers reproduce themselves.

2.4 Gaps in the literature

Four gaps emerge from this survey. First, the literature has not yet consolidated an integrated taxonomy of mature AI-assisted research practices. A contemporaneous scoping review by Raitskaya et al. (2025) maps the 2025 landscape of generative AI in scholarly writing, applying PRISMA-ScR – a parallel effort confirming the timeliness of scoping-review methodology but differing in scope: Raitskaya et al. map the landscape of generative AI in scholarly writing and publishing, whereas the present paper pivots the unit of analysis to institutional governance of research practice in doctoral schools, develops theory from the assimilation-gap tradition, and advances an explicit regional hypothesis for Eastern Europe that Raitskaya et al. do not address. Second, the literature has not yet developed an operationalised institutional framework for doctoral schools. Third, regional analyses for Eastern Europe remain sparse. Fourth, the adoption-mastery relationship has not been theorised as a coherent object of management research, although Dell'Acqua et al. (2023), Noy and Zhang (2023), Humlum and Vestergaard (2025), Brynjolfsson et al. (2025), and Mohammadi et al. (2026) provide empirical foundations.

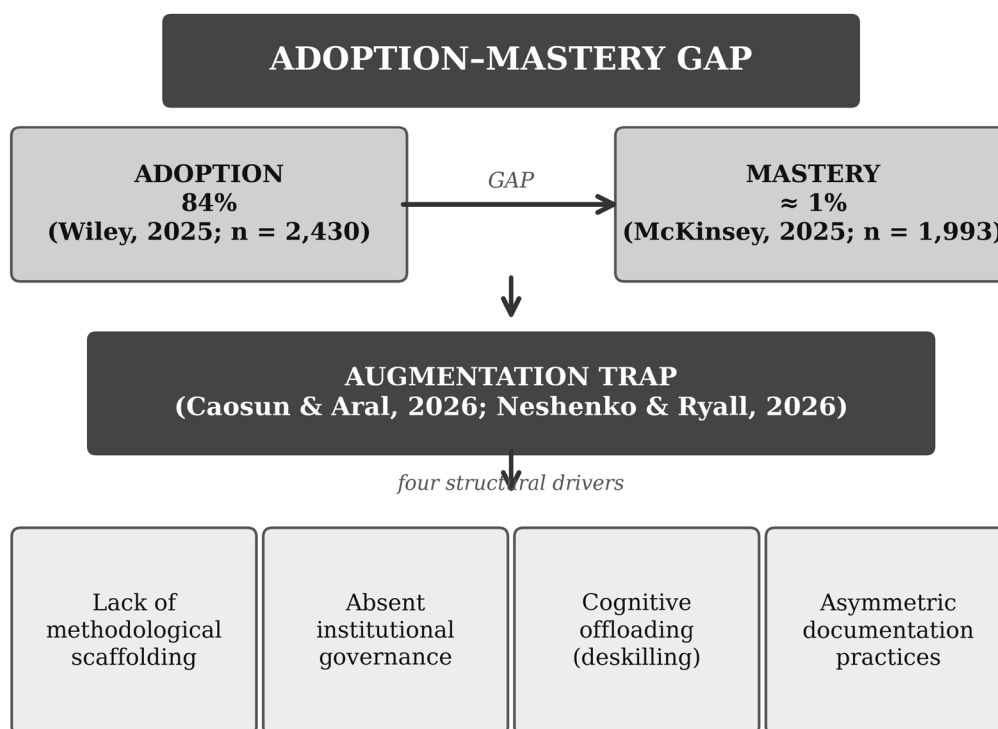


Figure 1. Conceptual model of the Adoption-Mastery Gap.

Source: Author's own elaboration.

3. OBJECTIVES AND METHODOLOGY OF RESEARCH

The *general objective* of this paper is to develop the conceptual framework of the adoption-mastery gap in academic research and to operationalise it for doctoral schools.

The *specific objectives* are:

1. To construct a descriptive taxonomy of mature AI-assisted research practices, derived from systematic analysis of the empirical corpus.
2. To derive an institutional framework on four levers – infrastructure, structured training, community of practice, declaration policy – anchored in research administration literature.
3. To position the East European academic context within the structural conditions that shape adoption-mastery dynamics.

The research is based on the following hypothesis: effective generative AI integration in academic research requires institutional mediation through specific organisational levers – not solely individual adoption – and the absence of such mediation produces a persistent and measurable adoption-mastery gap.

This paper conducts a scoping review with multi-source bibliographic consolidation, applying the PRISMA Extension for Scoping Reviews (PRISMA-ScR; Tricco et al., 2018) as the primary framework for scoping-specific reporting, aligned with the PRISMA 2020 reporting structure (Page et al., 2021). The operational sub-sections – Data Collection, Data Pre-processing, Data Analysis, and Interpretation and Data Visualisation – correspond to the four PRISMA phases (Identification, Screening, Eligibility, Inclusion), as illustrated in Figure 2.

Data Collection. A systematic search was conducted between January and May 2026 using a triangulated multi-agent retrieval design: three independent generative search interfaces (Claude with web search, Perplexity Deep Research, Gemini Deep Research) were issued an identical structured query specification covering seven evidentiary priorities. Each agent searched its underlying indices (Crossref, Semantic Scholar, arXiv, OpenAlex, Google Scholar, PubMed, publisher repositories).

The LLM-mediated retrieval was complemented by a parallel Boolean search of The Lens scholarly database (lens.org; over 270 million scholarly works), combining ("artificial intelligence" OR "generative AI" OR "ChatGPT" OR "large language model") AND ("academic research" OR "scholarly" OR "researcher" OR "doctoral"), January 2023 to May 2026. The two streams operate under different retrieval paradigms – LLM-mediated retrieval applies internal relevance filtering and returns curated outputs (n = 520 cumulative across the three agents), while Boolean search returns syntactically matching records subject to author-applied filters on citation count, peer-review status, and publication venue (n = 1,021). After cross-source deduplication and consolidation, 150 unique high-confidence records were retained.

Methodological note on LLM-mediated search. The use of LLM interfaces for bibliographic identification introduces a recognised hallucination risk. Four procedural safeguards were applied: (1) every source was cross-validated against Crossref records or originating repositories; (2) the triangulated design served as cross-validation, with three candidates excluded after discrepancies indicative of metadata fabrication; (3) the corpus was extended through The Lens and Google Scholar searches for Eastern European sources, mitigating anglophone bias; (4) sources flagged as uncertain were reviewed individually against publishing platforms. LLM-mediated triangulation is treated as an efficient identification mechanism that requires the same epistemic discipline as any bibliographic source.

Data Preprocessing. Inclusion criteria comprised peer-reviewed journal articles, reputable institutional reports (Wiley, HEPI, UNESCO, Digital Education Council, OECD, EUA,

Eurostat, McKinsey, IMF), Crossref-indexed conference proceedings, refereed working papers (NBER, CESifo, HBS, Harvard, MIT), and curated practitioner sources with stable URLs. The excluded criteria comprised news aggregators, SEO content, social media posts, and sources without verifiable DOIs. Manual screening reduced 150 candidates to approximately 100; full-text assessment reduced the analytical corpus to 66 sources retained for bibliometric and content analysis, complemented by 18 foundational theoretical and methodological anchors: Davis (1989), Cohen and Levinthal (1990), Cooke (2001), Repenning and Serman (2001, 2002), Zahra and George (2002), Rogers (2003), Venkatesh et al. (2003), Burton-Jones and Straub (2006), Trowler (2010), Whitchurch (2010), Fichman and Kemerer (1997, 1999), Shamseer et al. (2015), Brynjolfsson et al. (2018), Tricco et al. (2018), Page et al. (2021), and Mollick (2024). The consolidated reference list contains 88 entries (66 analytical sources, 18 foundational anchors, and 4 methodological references).

Data Analysis. Bibliometric analysis was performed using the *openalexR* package for metadata retrieval and the *bibliometrix* R-package (Aria & Cuccurullo, 2017) for keyword extraction, co-occurrence network construction, and Louvain clustering – following the bibliometric methodology established in management reviews (Ciocoiu et al., 2024; Felea et al., 2026). Given the small corpus size ($N = 66$), the co-occurrence threshold was set at $df \geq 2$ (following Cobo et al., 2011), with a curated stoplist and a lexical thesaurus applied to reduce noise. Visualisations were produced in R (*ggraph*, *ggplot2*) with a greyscale palette and Times New Roman typography. The content analysis generated the taxonomy (Section 4.2) and framework (Section 4.5).

Interpretation and Data Visualisation. All DOIs were validated against Crossref and cross-verified for scholarly indexing across OpenAlex, The Lens, and Scopus. Visualisations include the PRISMA flow diagram, annual production, country production, and a keyword co-occurrence network, complemented by conceptual figures. Figure 2 summarises the funnelling process.

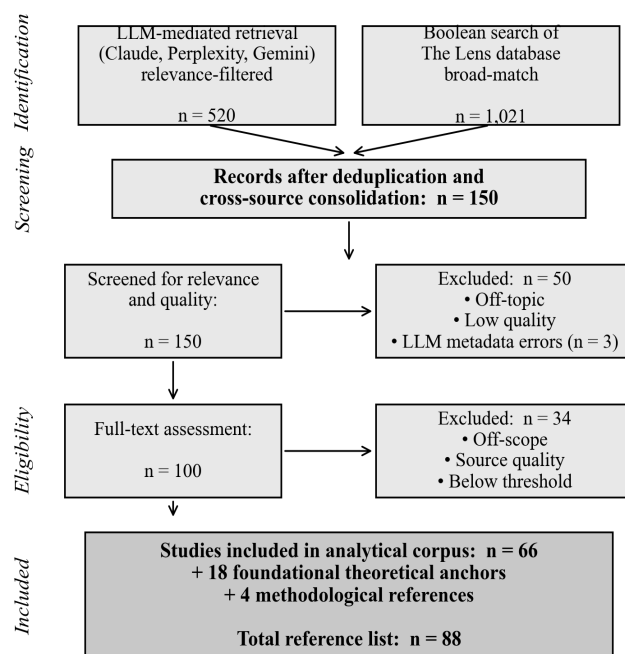


Figure 2. PRISMA-ScR Flow Diagram of the Scoping Review Process.

Source: Author's own elaboration based on data extracted from the consolidated multi-source research corpus.

4. RESULTS AND DISCUSSION

4.1 Bibliometric analysis

The bibliometric analysis of the retained analytical corpus reveals the structure of the field; the methodological template follows Ciocoiu et al. (2024) and Felea et al. (2026). Annual publication production (Figure 3) shows a marked inflection point in 2023 – the year of ChatGPT's broad release – followed by acceleration through 2026. Vorontsova et al. (2025) document the same trajectory using a Scopus-Web of Science PRISMA 2020 bibliometric workflow, confirming the suitability of the dual-database approach.

Figure 3. Annual Scientific Production on Generative AI in Academic Research, 2023–2026. *Source:* Author's own elaboration based on data extracted from the consolidated multi-source corpus via Crossref API metadata enrichment.

The keyword co-occurrence network (Figure 4) identifies four thematic clusters: a cognitive-behavioural cluster (cognitive offloading, critical thinking, AI literacy, AI anxiety), an adoption-tools cluster (adoption rates, AI adoption, AI tools, large language models, generative AI, machine learning), an institutional-educational cluster (education institutions, university students, student attitudes, technology acceptance, academic integrity), and a practice cluster (academic writing, bibliometric analysis, AI agents). Bridging keywords – adoption rates, AI tools, artificial intelligence – connect the clusters and supply the conceptual hooks for the descriptive taxonomy in Section 4.2.

4.2 A descriptive taxonomy of five mature AI-assisted research practices

From systematic content analysis of the empirical corpus, this paper derives a descriptive taxonomy of five mature AI-assisted research practices. The taxonomy is descriptive, not normative – the five categories represent operationally documented practices rather than ordered maturity stages. The categories were derived inductively from the empirical corpus, each requiring at least one peer-reviewed or refereed practitioner documentation source, and are grouped by operational prerequisites – from corpus-level batch processing requiring only a prompt protocol to multi-agent orchestration requiring deliberate software-engineering capability.

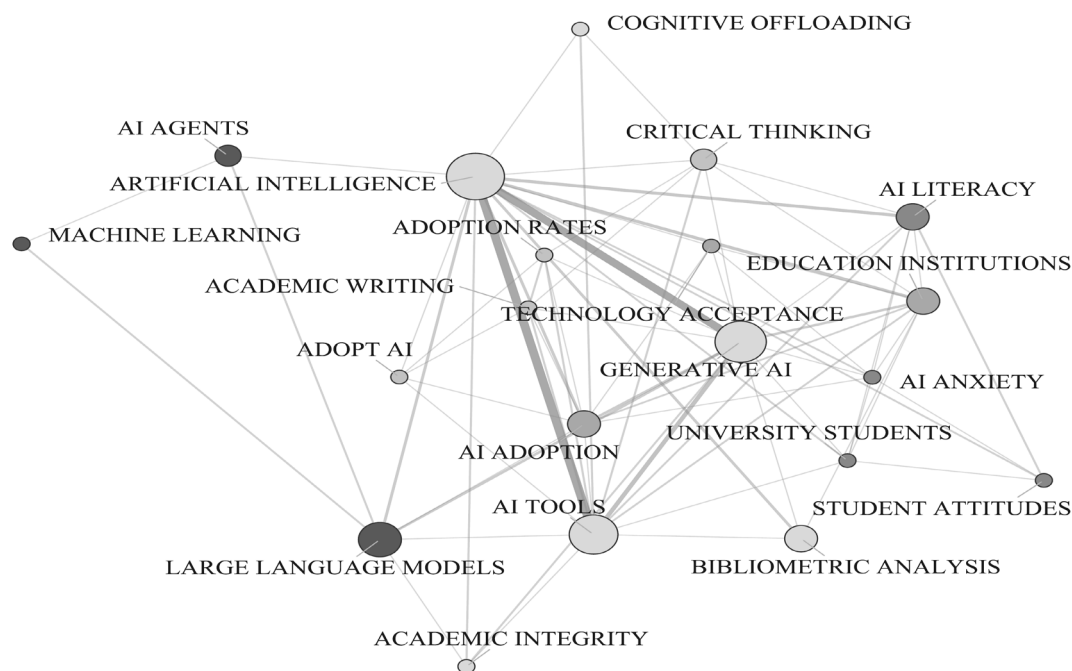


Figure 4. Keyword Co-occurrence Network.

Source: Author's own elaboration using bibliometrix (Aria & Cuccurullo, 2017) with term extraction from titles and abstracts retrieved via the OpenAlex API.

(a) Batch processing of research corpora. Dai et al. (2025) document large-scale systematic literature screening with sensitivity values between 0.77 and 0.99 across six thoracic surgery meta-analyses and a tenfold reduction in screening time. The prompt-engineering effect – modified prompts raising sensitivity from 0.87 to 0.98 – is documented evidence that the productive use of generative AI in evidence synthesis depends on technique, not access. Replicability conditions: a curated topic, a Boolean query against established databases, and a documented prompt iteration protocol.

(b) Local model deployment. Jayathilake et al. (2026) document Llama 3.3 deployed via Ollama for sentiment analysis of open-ended survey responses, achieving F1 scores around 0.97 – comparable to human coder agreement, GDPR-compliant, and runnable on consumer hardware. Replicability conditions: consumer-grade GPU access (≥ 16 GB VRAM), an open-source model with documented benchmark performance, and a validated prompt template.

(c) System prompt engineering as documented capability. Anam (2025; $n=243$) establishes prompt engineering competence – not adoption – as the operative predictor of productivity gain. Susnjak (2025) extends the methodological frontier through declarative prompt optimisation, in which prompts are compiled against test suites rather than hand-crafted. Replicability conditions: a defined task with measurable success criteria, a test corpus, and version-controlled prompt iteration.

(d) Custom research tools and minimum viable pipelines. Huang et al. (2025) have released *AI-Researcher*, an open-source end-to-end autonomous LLM agent pipeline mimicking the doctoral research lifecycle. Replicability conditions: Python and version-control fluency, access to cited models or open-source equivalents, and willingness to maintain the pipeline as a research artefact.

(e) Deep methodological assistance via multi-agent systems. Lu et al. (2026) demonstrate an end-to-end agentic pipeline – The AI Scientist – with one generated submission exceeding an ICLR workshop acceptance threshold. Guo et al. (2024) provide a refereed taxonomic survey of LLM-based multi-agent systems. Pantiukhin et al. (2025) extend the practice beyond computer science through PANGAEA GPT. Sami et al. (2024) document a multi-agent system automating the full systematic literature review pipeline. Practitioner documentation (Willison, 2025; Manandhar-Richardson, 2026; Maynard, 2026) supplies operational templates not yet entered into peer-reviewed literature. Replicability conditions: deliberate task decomposition, an agent platform with documented orchestration, and a verification layer subjecting outputs to human review.

Table 1 compares the academic research case with software development and marketing along three axes – adoption rate, productivity translation, and maturity indicators.

Table 1. Comparative Adoption, Productivity Translation, and Maturity across Three Sectors

Sector	Adoption rate (2024–2026)	Productivity translation	Maturity indicators
Software development	76% using or planning AI tools (Stack Overflow, 2024, n>65,000); rising to approximately 93% by 2025–2026 (GitClear, 2025)	Heterogeneous: RCT on 16 developers finds AI tools increased completion time by 19% (METR, 2025); telemetry from 22,000 developers shows 66% throughput gain alongside 28% more bugs, tripled incidents, tenfold code churn (Faros AI, 2026)	Code review processes; AI-assisted commit guidelines; emerging governance frameworks
Marketing	85% adoption, n>1,000 (CoSchedule, 2025)	Direct and measurable: campaign ROI and A/B testing provide rapid feedback	A/B testing protocols; brand voice governance; content audit workflows
Academic research	84% overall (Wiley, 2025, n=2,430); 62% for research tasks; 92% among UK students (HEPI, 2025, n=1,041)	Indirect, unequal across skill levels (Dell'Acqua et al., 2023; Humlum & Vestergaard, 2025; Brynjolfsson et al., 2025)	Largely absent – 19% of HEIs have formal AI policy (UNESCO, 2025); <40% of doctoral schools have governance (EUA, 2026); few training curricula

Source: Author's own elaboration based on Wiley (2025), HEPI (2025), Digital Education Council (2024, 2025), UNESCO (2025), EUA (2026), Stack Overflow (2024), GitClear (2025), METR (2025), Faros AI (2025, 2026), CoSchedule (2025), Dell'Acqua et al. (2023), Humlum and Vestergaard (2025), and Brynjolfsson et al. (2025).

The comparison clarifies the structural distinctiveness of the academic case: it lacks both the rapid-feedback mechanisms and the governance infrastructure that enable other sectors to translate adoption into measurable productivity. The taxonomy and cross-sector comparison together reveal that mature practices are distributed unevenly across institutional contexts.

4.3 The East European structural-isolation hypothesis

This paper advances a theoretical hypothesis – to be tested in subsequent empirical work – that the Eastern European academic context exhibits a distinctive structural-isolation pattern that widens the adoption-mastery gap regionally. The hypothesis is anchored in the regional innovation systems literature (Cooke, 2001).

The macroeconomic backdrop shows a marked asymmetry. Eurostat (2025; dataset *isoc_eb_ai*) reports enterprise-level AI adoption rates in Eastern European member states substantially below the EU average: Romania at 5.2%, Poland at 8.4%, Bulgaria at 8.5%, Hungary at 10.4%, and Serbia at 10.1%, against the EU-wide figure of 20.0%. By contrast, individual academic adoption reaches levels comparable to Western European norms – a pattern documented in Malița et al. (2025) on Romanian and Moldovan students ($n=189$) and in the larger regional surveys reviewed below. The asymmetry between low enterprise adoption and high individual academic adoption defines the structural-isolation dynamic.

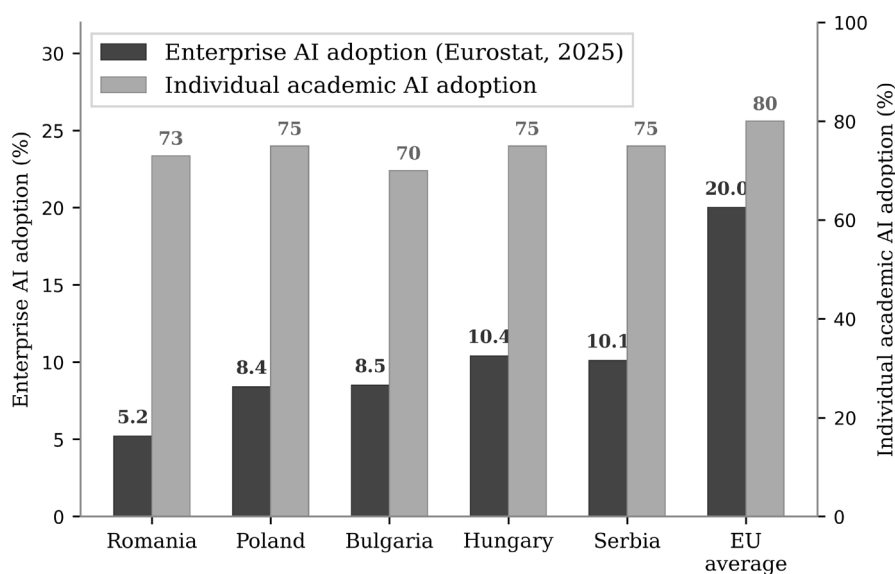


Figure 5. Eastern European Structural Disjunction: Enterprise vs Academic AI Adoption

Source: Author's own elaboration based on Eurostat (2025) and regional academic adoption surveys.

Three patterns emerge from the regional evidence.

First, country-level studies reveal a consistent picture of high individual adoption without institutional scaffolding: Galjak and Budić (2025) document a generational divide in Serbian academia ($n = 823$); Cernicova-Buca and Palea (2026) find that Romanian university leaders concentrate on anti-plagiarism risk mitigation rather than methodological redesign; Întorsureanu et al. (2024) identify two dimensions of LLM use among Romanian academics – structural (text production) and consultative (methodological deliberation) – with disclosure infrastructure emerging as the feasible response; Fekete (2026) finds that Hungarian AI literacy correlates with informal peer networks rather than formal instruction; and Rajki et al. (2025; $n=1,027$) confirm a gap between awareness and qualified use across Hungarian humanities, social sciences, and teacher education. Ștefan et al. (2025), employing PLS-SEM combined with artificial neural network analysis, demonstrate that social influence operates as an adoption driver only under supportive institutional conditions – a finding with direct regional implications.

Second, cross-national studies extend the pattern beyond individual countries. Skalka et al. (2025a; n = 1,195, IEEE Access) cover Slovakia, Czech Republic, Poland, Lithuania, and Ukraine; the companion dataset paper (Skalka et al., 2025b, Data in Brief) supports cross-national comparison. Singh and Strzelecki (2025; n=640 academics, ten Polish universities) provide a particularly rigorous study on Polish academic AI adoption. Lukianova et al. (2025; n=1,051) contribute a Polish-Ukrainian comparative study. Ravšelj et al. (2025; n = 23,218, 109 countries) supply the worldwide frame within which regional findings are anchored. Together with the Romanian, Serbian, and Hungarian evidence, the coverage spans eight CEE countries with over 27,000 respondents.

Third, the Czech Republic provides a contrasting case: centrally coordinated national initiatives – anchored in the European DigCompEdu framework and documented in the OECD (2025) review of AI adoption in education systems – aim to develop AI competence among teachers and integrate generative AI into higher education curricula, illustrating variation in regional policy maturity within the structural-isolation pattern.

The hypothesis, formally stated: the structural disjunction between low enterprise AI maturity and high individual academic adoption in Eastern European member states reduces cross-pollination between industry and academic research, generating a regionally widened adoption-mastery gap. The causal direction remains an open empirical question – high academic adoption may equally be a vanguard response to low industrial maturity. The hypothesis holds that the disjunction itself constrains cross-pollination dynamics. Empirical disentanglement through longitudinal studies is a priority for the future research agenda (Section 5).

4.4 The Evans paradox – individual acceleration, collective narrowing

A notable recent contribution to the empirical literature is Hao et al. (2026), whose analysis of 41.3 million scientific papers establishes the paradox at the centre of this paper's argument. Following the convention of naming a finding after its senior author, we propose the term Evans paradox, pending independent replication, to designate the asymmetry documented in Hao et al. (2026), in reference to James Evans (University of Chicago). At the individual level, AI-augmented researchers publish 3.02 times more papers, accumulate 4.84 times more citations, and achieve leadership 1.37 years earlier. At the collective level, the same population covers 4.63% less topical ground and exhibits 22% lower inter-researcher engagement – "lonely crowds".

The paradox is that conventional metrics – publication volume, citation accrual, time-to-leadership – capture individual private benefit while obscuring a collective public good: topical diversity and inter-researcher engagement. A discipline can simultaneously exhibit rising adoption, rising publication counts, rising citations – and falling intellectual diversity. The Evans paradox and the augmentation trap may represent two registers of the same underlying dynamic: at the individual level, experienced researchers retain capability while less-experienced researchers deskilling (Caosun & Aral, 2026); at the field level, the aggregate effect manifests as topical homogenisation and reduced inter-researcher engagement. Research management operates on both registers: protecting individual gains while creating institutional conditions that counter topical narrowing. The four-lever framework (Section 4.5) addresses both dimensions.

4.5 A four-lever institutional framework for universities

The taxonomy (Section 4.2) and the regional hypothesis (Section 4.3) converge on a managerial question: what can universities do to narrow the adoption-mastery gap? Drawing on Trowler's (2010) analysis and Whitchurch's (2010) third-space professionals, this paper proposes a four-lever institutional framework. Each lever creates conditions under which Social Influence – the UTAUT construct frequently identified as contested in academic adoption – can operate effectively. Ştefan et al. (2025) provide the empirical anchor: social influence functions as an adoption driver only under supportive institutional conditions. More broadly, the four levers map onto the UTAUT architecture: Infrastructure operationalises Facilitating Conditions; Structured Training reduces Effort Expectancy; Community of Practice channels Social Influence; and Declaration Policy anchors the normative environment within which Performance Expectancy is calibrated. Two contemporaneous frameworks reinforce the architecture: Smith et al. (2025) propose a multi-lever framework with structurally parallel components; Pinho et al. (2025) propose a Living GenAI Governance Model at macro, meso, and micro levels. The convergence across these Australian, Portuguese, and Eastern European analyses suggests that the four-lever architecture is consistent with patterns observed independently across national contexts.

Lever 1 – Infrastructure. Effective infrastructure provides institutional access to frontier LLM APIs alongside local compute capacity for sensitive data (GDPR, IRB constraints). The Jayathilake et al. (2026) case demonstrates operational feasibility at undergraduate and PhD-student scale. The European Data Protection Board (2025) guidance on AI privacy establishes conditions under which local deployment becomes the preferred architecture. Infrastructure cost appears to be receding as a binding constraint; procurement, data-governance compliance, and access protocols remain salient.

	L1 Infrastructure	L2 Structured Training	L3 Community of Practice	L4 Declaration Policy
Research Deans	● Procurement API access Local compute	● Hire instructors Approve curriculum Budget	□ Sponsor forums Set expectations	□ Endorse policy Monitor outcomes
PhD Programme Directors	□ Specify needs Co-design specs	● Curriculum design Mandatory modules	● Convene forums Link to milestones	□ Operationalise in programme
Research Administrators (third-space)	□ Operational mgmt & support	□ Logistics Scheduling	□ Coordination Documentation	● Template design Compliance audit

● Primary responsibility □ Supportive role

Figure 6. Four-Lever Institutional Framework for Doctoral Schools: Actors and Operational Responsibilities.

Source: Author's own elaboration anchored in Trowler (2010) and Whitchurch (2010).

Lever 2 – Structured training. The training lever comprises two to three mandatory modules covering LLM foundations, prompt engineering (Anam, 2025; Susnjak, 2025), and local model deployment, targeting the competencies identified by Popa et al. (2024) as necessary for effective AI use. The OECD (2025) and Higher Education Authority (2025) provide international policy templates. Without structured training, the gap persists across doctoral cohorts.

Lever 3 – Community of practice. The augmentation-trap dynamic (Section 2.2) can be partially counteracted through structured forums – drawing on the community-of-practice model formalised by Wenger (1998) – in which mature practices are exchanged across generational lines, and doctoral candidates articulate and defend their methodological choices in peer settings. The model is established in software engineering (tech talks, post-mortems) and clinical medicine (case conferences); what is missing is deliberate institutional framing as professional development. Sakar and Nayak (2026) reinforce this: mastery requires structures that distribute responsibility and codify best practice.

Lever 4 – Declaration policy. Neither prohibition nor laissez-faire supports mature practice. Luo (2024) confirms that prohibition-oriented stances remain the default even at elite institutions. A mature institutional posture supports systematic disclosure with methodological detail. The Wiley (2025) finding that 73% of researchers expect publisher guidance signals that institutional disclosure leadership has been ceded to commercial actors. The University of Sydney's two-lane assessment framework (University of Sydney, 2024) illustrates one template for disclosure-based governance. Research deans and editorial boards remain the natural locus of disclosure leadership.

The four levers operate jointly: infrastructure without training yields surface-level adoption; training without community of practice yields cohort-level skill loss; community of practice without declaration policy yields normative fragmentation; declaration policy without infrastructure yields formal compliance without substantive practice. The framework lends itself to implementation within one academic year at the level of a single doctoral school (Whitchurch, 2010).

Discussion closure – the asymmetry of literature documentation. Across the corpus, disciplines with low entry barriers to AI use (computer science, computational social science) produce literature on practices; disciplines with high barriers (humanities, qualitative health research) produce literature on barriers. The two corpora rarely interact with each other. This structural feature reflects an unequal distribution of methodological resources and disciplinary norms. Systematic comparative analysis of cross-citation patterns is reserved for future work.

5. CONCLUSIONS

This paper has conceptualised the adoption-mastery gap as a distinct object of study in research management; the four contributions summarised below follow from this conceptualisation.

The first extends the assimilation-gap tradition (Fichman & Kemerer, 1997, 1999) to generative AI in academic research, characterising the gap as a persisting structural feature with the augmentation trap as its causal mechanism, anchored in convergent evidence from Caosun and Aral (2026), Neshenko and Ryall (2026), Rinta-Kahila et al. (2023), Bastani et al. (2025), and Gerlich (2025).

The second develops a descriptive taxonomy of five mature AI-assisted research practices derived from the empirical corpus (Section 4.2).

The third proposes a four-lever institutional framework for universities anchored in research administration literature (Trowler, 2010; Whitchurch, 2010) and reinforced by contemporaneous frameworks from Smith et al. (2025) and Pinho et al. (2025).

The fourth advances a theoretical hypothesis for future empirical testing of East European structural isolation, supported by regional evidence across eight CEE countries (Section 4.3) and benchmarked against the worldwide frame of Ravšelj et al. (2025).

MANAGERIAL IMPLICATIONS. The framework operationalises three actor roles. Research deans hold resource-allocation authority for infrastructure and structured training. PhD programme directors hold curricular and convening authority for training and community of practice. Research administrators – third-space professionals (Whitchurch, 2010) – hold operational responsibility for declaration policy, disclosure design, compliance monitoring, and institutional learning reports. A doctoral school implementing the framework can begin with Lever 4 (declaration policy – lowest implementation cost), proceed to Lever 2 (structured training, drawing on existing competency frameworks), build Lever 3 (community of practice, requiring convening authority rather than procurement), and address Lever 1 (infrastructure) as institutional resources permit.

LIMITATIONS. This paper develops a conceptual framework – it does not provide primary empirical validation on the East European population. The reliance on triangulated multi-agent LLM-mediated identification provides broader coverage at the cost of some uniformity in source quality control; the four safeguards (Section 3) mitigate but do not eliminate this risk. English-language sources predominate. The scoping review framing does not deliver PICO-style inferential precision.

FUTURE RESEARCH. Four lines follow. First, empirical validation of the structural-isolation hypothesis through a structured survey at ASE Bucharest and regional partners (planned 2026–2027). Second, systematic comparative analysis of cross-citation patterns between the practice and barrier corpora. Third, longitudinal monitoring of doctoral schools adopting the four-lever framework. Fourth, formalisation of the triangulated multi-agent retrieval methodology as a replicable scoping-review extension protocol.

The strategic question for research management is not whether to adopt generative artificial intelligence, but how to scaffold the institutional conditions under which adoption matures into methodologically grounded practice.

AI USE DECLARATION. In accordance with Bucharest University of Economic Studies Senate Resolution no. 38/25 March 2026 approving the *Guide on the Use of Generative Artificial Intelligence Tools*, the author declares the use of the following AI tools for triangulated bibliographic identification and pattern matching: Anthropic Claude Opus 4.7 [large language model] (May 2026; <https://claude.ai>); Perplexity Deep Research [multi-agent retrieval interface] (May 2026; <https://perplexity.ai>); Google Gemini Deep Research [large language model with retrieval] (May 2026; <https://gemini.google.com>). All bibliographic records were independently verified against Crossref, with the four procedural safeguards described in Section 3 applied. The author retains full responsibility for analytical content; no data fabrication, peer-review submission to AI tools, or plagiarism circumvention occurred.

REFERENCES

- Anam, R. K. (2025). Prompt engineering and the effectiveness of large language models in enhancing human productivity. *arXiv*. <https://doi.org/10.48550/arXiv.2507.18638>
- 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>
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- Bianchini, S., Müller, M., & Pelletier, P. (2025). Drivers and barriers of AI adoption and use in scientific research. *Technological Forecasting and Social Change*, 220, 124303. <https://doi.org/10.1016/j.techfore.2025.124303>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*. <https://doi.org/10.1093/qje/qjae044>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2018). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, 23-57. University of Chicago Press / NBER. <https://www.nber.org/chapters/c14007>
- Burton-Jones, A., & Straub, D. W. (2006). Reconceptualizing system usage: An approach and empirical test. *Information Systems Research*, 17(3), 228-246. <https://doi.org/10.1287/isre.1060.0096>
- Cao, K., Wang, P., & Jiang, X. (2026). AI anxiety and adoption intention in higher education: A PLS-SEM analysis combining TAM and UTAUT. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-35823-9>
- Caosun, M., & Aral, S. (2026). The augmentation trap: AI productivity and the cost of cognitive offloading. *arXiv*. <https://doi.org/10.48550/arXiv.2604.03501>
- Cernicova-Buca, M., & Palea, A. (2026). Leading through transformation: University responses to generative AI in Romanian higher education. *Professional Communication and Translation Studies*, 19, 44-55. <https://doi.org/10.59168/DDJK2862>
- Chen, G., Fan, J., & Azam, M. (2024). Exploring AI chatbot adoption among research scholars using UTAUT. *Journal of Librarianship and Information Science*. <https://doi.org/10.1177/09610006241269189>
- Ciociu, C. N., Radu, C., Colesca, S. E., & Prioteasa, A. (2024). Exploring the link between risk management and performance of MSMEs: A bibliometric review. *Journal of Economic Surveys*, 39(4), 1523-1552. <https://doi.org/10.1111/joes.12664>
- Cobo, M. J., López-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2011). An approach for detecting, quantifying, and visualizing the evolution of a research field: A practical application to the fuzzy sets theory field. *Journal of Informetrics*, 5(1), 146-166. <https://doi.org/10.1016/j.joi.2010.10.002>
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128-152. <https://doi.org/10.2307/2393553>
- Cooke, P. (2001). Regional innovation systems, clusters, and the knowledge economy. *Industrial and Corporate Change*, 10(4), 945-974. <https://doi.org/10.1093/icc/10.4.945>
- CoSchedule. (2025). *State of AI in marketing 2025*. CoSchedule.
- Dai, Z., Wang, F., Shen, C., Ji, Y., Li, Z., Wang, Y., & Pu, Q. (2025). Accuracy of large language models for literature screening in thoracic surgery: Diagnostic study. *Journal of Medical Internet Research*, 27, e67488. <https://doi.org/10.2196/67488>

- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319-340. <https://doi.org/10.2307/249008>
- Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality* (Harvard Business School Working Paper No. 24-013). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321
- Digital Education Council. (2024). *Global AI student survey 2024*. <https://www.digitaleducationcouncil.com/post/digital-education-council-global-ai-student-survey-2024>
- Digital Education Council. (2025). *What faculty want: Key results from the global AI faculty survey 2025*. <https://www.digitaleducationcouncil.com/post/what-faculty-want-key-results-from-the-global-ai-faculty-survey-2025>
- European Data Protection Board. (2025). *AI privacy risks and mitigations: Large language models (LLMs)*. <https://www.edpb.europa.eu/system/files/2025-04/ai-privacy-risks-and-mitigations-in-llms.pdf>
- European University Association. (2026). *Generative AI in European doctoral education: Adoption, governance, and institutional response*. EUA.
- Eurostat. (2025). *Enterprise artificial intelligence use dataset* [Data set]. https://ec.europa.eu/eurostat/databrowser/view/isoc_eb_ai/default/table
- Faros AI. (2025). *AI engineering impact report 2025: Telemetry from 10,000 developers*. Faros AI Research. <https://www.faros.ai/research>
- Faros AI. (2026). *AI acceleration whiplash: Telemetry from 22,000 developers across 4,000 teams*. Faros AI Research. <https://www.faros.ai/research/ai-acceleration-whiplash>
- Fekete, I. (2026). Between promise and practice: Bridging ethical artificial intelligence literacy gaps across students, educators, and policy. *Journal of University Teaching and Learning Practice*. Advance online publication. <https://doi.org/10.53761/9fr5vw20>
- Felea, M., Bran, I., Bujor, D., & Allahverdiyev, K. (2026). Exploring research trends in sustainable supply chain management: A bibliometric analysis. *Amfiteatru Economic*, 28(71), 88-106. <https://doi.org/10.24818/EA/2026/71/88>
- Fichman, R. G., & Kemerer, C. F. (1997). The assimilation of software process innovations: An organizational learning perspective. *Management Science*, 43(10), 1345-1363. <https://doi.org/10.1287/mnsc.43.10.1345>
- Fichman, R. G., & Kemerer, C. F. (1999). The illusory diffusion of innovation: An examination of assimilation gaps. *Information Systems Research*, 10(3), 255-275. <https://doi.org/10.1287/isre.10.3.255>
- Galjak, M., & Budić, M. (2025). Generational divide in AI adoption for academic writing: Evidence from Serbian social scientists. *Social Science Computer Review*. Advance online publication. <https://doi.org/10.1177/08944393251413796>
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), 6. <https://doi.org/10.3390/soc15010006>
- GitClear. (2025). *Adoption of developer AI tools 2022–2025*. GitClear Research. https://www.gitclear.com/research/developer_ai_assistant_adoption_by_year_with_ai_delegation_buckets
- Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). Large language model based multi-agents: A survey of progress and challenges. In K. Larson (Ed.), *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)* (pp. 8048-8057). International Joint Conferences on Artificial Intelligence Organization. <https://doi.org/10.24963/ijcai.2024/890>

- Hao, Q., Xu, F., Li, Y., & Evans, J. (2026). Artificial intelligence tools expand scientists' impact but contract science's focus. *Nature*, 649(8099), 1237-1243. <https://doi.org/10.1038/s41586-025-09922-y>
- Higher Education Authority. (2025). *National policy framework on generative AI in teaching and learning*. HEA Ireland. <https://hea.ie/2025/12/22/hea-publishes-national-policy-framework-on-generative-ai-in-teaching-and-learning/>
- Higher Education Policy Institute. (2025). *Student generative AI survey 2025* (HEPI Policy Note 61). <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2025/>
- Huang, C., et al. (2025). *AI-Researcher: An autonomous large-language-model agent for end-to-end PhD-level research* [Software repository]. GitHub. <https://github.com/HKUDS/AI-Researcher>
- Humlum, A., & Vestergaard, E. (2025). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1). <https://doi.org/10.1073/pnas.2414972121>
- Întorsureanu, I., Voicu-Dorobanțu, R., Nisioiu, C. F., & Ploae, C. (2024). Generative artificial intelligence and the academic integrity of graduation works in economics – Exploring perceptions of Romanian academia. *Economic Computation and Economic Cybernetics Studies and Research*, 58(2/2024), 132-147. <https://doi.org/10.24818/18423264/58.2.24.08>
- Jayathilake, H. D., Yoon, A., & Hong, M. (2026). Leveraging Local-LLM for sentiment analysis to enhance text data quality assessment in information science research. *The Electronic Library*. Advance online publication. <https://doi.org/10.1108/EL-08-2025-0354>
- Lu, C., Lu, C., Lange, R. T., Foerster, J., Clune, J., & Ha, D. (2026). Towards end-to-end automation of AI research. *Nature*. <https://doi.org/10.1038/s41586-026-10265-5> (Originally released as arXiv:2408.06292, 2024.)
- Lukianova, L., Symela, K., & Religa, J. (2025). Artificial intelligence in higher education: Scientific and educational convergence in the Polish-Ukrainian context. *The New Educational Review*, 81, 182-200. <https://doi.org/10.15804/tner.2025.81.3.12>
- Luo, J. (2024). A critical review of GenAI policies in higher education assessment: a call to reconsider the "originality" of students' work. *Assessment & Evaluation in Higher Education*, 49(5), 651-664. <https://doi.org/10.1080/02602938.2024.2309963>
- Malița, L., Dumbraveanu, R., Balmus, O., & Grosseck, G. (2025). Students' ethical reasoning about generative AI in higher education assessment: A cross-national survey study of Romania and Moldova. *Research and Education*, 13, art. 6. <https://doi.org/10.63467/red.13.2025.art6>
- Manandhar-Richardson, T. (2026, March 4). *Using Claude Code and Codex CLI for social science research: My learnings as they happen*. Evolving Impact. <https://evolvingimpact.wordpress.com/2026/03/04/using-claude-code-and-codex-cli-for-social-science-research-my-learnings-as-they-happen/>
- Maynard, A. (2026). *I cracked and wrote an academic paper with AI*. Future of Being Human. <https://www.futureofbeinghuman.com/p/i-cracked-and-wrote-an-academic-paper>
- McKinsey. (2025). *The state of AI: Global survey 2025*. McKinsey QuantumBlack. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- METR. (2025). *Measuring the impact of early-2025 AI on experienced open-source developers*. arXiv. <https://doi.org/10.48550/arXiv.2507.09089>
- Mohammadi, E., Thelwall, M., Cai, Y., Collier, T., Tahamtan, I., & Eftekhari, A. (2026). Is generative AI reshaping academic practices worldwide? *Information Processing & Management*, 63(1), 104350. <https://doi.org/10.1016/j.ipm.2025.104350>
- Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio/Penguin.

- Morton, J. (2024). *Generative AI adoption and three traps for organizational agility*. California Management Review Insights. <https://cmr.berkeley.edu/2024/03/generative-ai-adoption-and-three-traps-for-organizational-agility/>
- Neshenko, N., & Ryall, M. D. (2026). When artificial intelligence does strategy: Learning, good times, lock-in, and human-driven strategic renewal. *Strategy Science*. Advance online publication. <https://doi.org/10.1287/stsc.2025.0448>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4375283>
- OECD (2025). *AI adoption in the education system*. Organisation for Economic Co-operation and Development. https://www.oecd.org/en/publications/ai-adoption-in-the-education-system_69bd0a4a-en.html
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Pantiukhin, D., Shapkin, B., Kuznetsov, I., Jost, A. A., & Koldunov, N. (2025). Accelerating earth science discovery via multi-agent LLM systems. *Frontiers in Artificial Intelligence*, 8, 1674927. <https://doi.org/10.3389/frai.2025.1674927>
- Pinho, I., Costa, A. P., & Pinho, C. (2025). Generative AI governance model in educational research. *Frontiers in Education*, 10, 1594343. <https://doi.org/10.3389/educ.2025.1594343>
- Popa, I., Cioc, M. M., Breazu, A., & Popa, C. F. (2024). Identifying sufficient and necessary competencies in the effective use of artificial intelligence technologies. *Amfiteatru Economic*, 26(65), 33–52. <https://doi.org/10.24818/EA/2024/65/33>
- Radu, C., Ciocoiu, C. N., Veith, C., & Dobrea, R. C. (2024). Artificial intelligence and competency-based education: A bibliometric analysis. *Amfiteatru Economic*, 26(65), 220-240. <https://doi.org/10.24818/EA/2024/65/220>
- Raitskaya, L., Tikhonova, E., & Mekeko, N. (2025). The 2025 landscape of generative AI in scholarly writing and publishing: A scoping review of uses and ethical approaches. *Journal of Language and Education*.
- Rajki, Z., Dringó-Horváth, I., & Nagy, J. T. (2025). Artificial intelligence in higher education: Students' artificial intelligence use and its influencing factors. *Journal of University Teaching and Learning Practice*, 22(2), 1015. <https://doi.org/10.53761/j0rebh67>
- Ravšelj, D., Keržič, D., Tomaževič, N., Umek, L., Brezovar, N., Andronic, O., et al. (2025). Higher education students' perceptions of ChatGPT: A global study of early reactions. *PLOS One*, 20(2), e0315011. <https://doi.org/10.1371/journal.pone.0315011>
- Repenning, N. P., & Serman, J. D. (2001). Nobody ever gets credit for fixing problems that never happened. *California Management Review*, 43(4), 64-88. <https://doi.org/10.2307/41166101>
- Repenning, N. P., & Serman, J. D. (2002). Capability traps and self-confirming attribution errors in the dynamics of process improvement. *Administrative Science Quarterly*, 47(2), 265-295. <https://doi.org/10.2307/3094806>
- Rinta-Kahila, T., Penttinen, E., Salovaara, A., Soliman, W., & Ruissalo, J. (2023). The vicious circles of skill erosion: A case study of cognitive automation. *Journal of the Association for Information Systems*, 24(5), 1378-1412. <https://doi.org/10.17705/1jais.00829>
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Sakar, R., & Nayak, R. (2026). Co-intelligence and the reconfiguration of human work and learning: A critical review of Ethan Mollick's *Living and working with AI*. *Journal of Engineering Education Transformations*, 36(1). <https://doi.org/10.55613/jeet.v36i1.223>

- Sami, A. M., Rasheed, Z., Kemell, K.-K., Waseem, M., Kilamo, T., Saari, M., Nguyen Duc, A., Systä, K., & Abrahamsson, P. (2024). *System for systematic literature review using multiple AI agents: Concept and an empirical evaluation* (arXiv:2403.08399). <https://doi.org/10.48550/arXiv.2403.08399>
- Shamseer, L., Moher, D., Clarke, M., Ghersi, D., Liberati, A., Petticrew, M., Shekelle, P., & Stewart, L. A. (2015). Preferred reporting items for systematic review and meta-analysis protocols (PRISMA-P) 2015: Elaboration and explanation. *BMJ*, *349*, g7647. <https://doi.org/10.1136/bmj.g7647>
- Singh, M., & Strzelecki, A. (2025). Diffusion of innovations in generative AI adoption in higher education. *Education and Information Technologies*, *31*(2), 621-645. <https://doi.org/10.1007/s10639-025-13835-8>
- Skalka, J., Przybyła-Kasperek, M., Smyrnova-Trybulska, E., Klimeš, C., Farana, R., Dagienė, V., & Dolgopolas, V. (2025a). Artificial intelligence literacy structure and the factors influencing student attitudes and readiness in Central Europe universities. *IEEE Access*, *13*, 93235-93258. <https://doi.org/10.1109/ACCESS.2025.3573575>
- Skalka, J., Przybyła-Kasperek, M., Smyrnova-Trybulska, E., Klimeš, C., Farana, R., Dagienė, V., & Dolgopolas, V. (2025b). Cross-national survey data on student attitudes toward artificial intelligence. *Data in Brief*, *62*, 112022. <https://doi.org/10.1016/j.dib.2025.112022>
- Smith, S., Tate, M., Freeman, K., Walsh, A., Ballsun-Stanton, B., Hooper, M., & Lane, M. (2025). A university framework for the responsible use of generative AI in research. *Journal of Higher Education Policy and Management*. Advance online publication. <https://doi.org/10.1080/1360080X.2025.2509187>
- Stack Overflow. (2024). *2024 developer survey*. Stack Overflow. <https://survey.stackoverflow.co/2024/>
- Ștefan, S. C., Popa, I., & Breazu, A. (2025). A new framework for the artificial intelligence entrepreneurship ecosystem. *Journal of Innovation & Knowledge*, *10*(4), 100850. <https://doi.org/10.1016/j.jik.2025.100850>
- Susnjak, T. (2025). Compiling prompts, not crafting them: A reproducible workflow for AI-assisted evidence synthesis. *arXiv*. <https://doi.org/10.48550/arXiv.2509.00038>
- Tricco, A. C., Lillie, E., Zarin, W., O'Brien, K. K., Colquhoun, H., Levac, D., Moher, D., Peters, M. D. J., Horsley, T., Weeks, L., Hempel, S., Akl, E. A., Chang, C., McGowan, J., Stewart, L., Hartling, L., Aldcroft, A., Wilson, M. G., Garritty, C., ... Straus, S. E. (2018). PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and explanation. *Annals of Internal Medicine*, *169*(7), 467-473. <https://doi.org/10.7326/M18-0850>
- Trowler, P. (2010). UK higher education: Captured by new managerialist ideology? In *Higher Education Dynamics*. https://doi.org/10.1007/978-90-481-9163-5_10
- UNESCO. (2025). *Global survey on generative AI in higher education: 400 institutions across 90 countries*. United Nations Educational, Scientific and Cultural Organization. <https://www.unesco.org/en/articles/unesco-survey-two-thirds-higher-education-institutions-have-or-are-developing-guidance-ai-use>
- University of Sydney. (2024). *Aligning our assessments to the age of generative AI*. Teaching@Sydney. <https://educational-innovation.sydney.edu.au/teaching@sydney/aligning-our-assessments-to-the-age-of-generative-ai/>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, *27*(3), 425-478. <https://doi.org/10.2307/30036540>
- Vorontsova, A., Tarasenko, S., Duranowski, W., Durasiewicz, A., Soss, J., & Bilovol, A. (2025). A bibliometric analysis of the economic effects of using artificial intelligence and ChatGPT tools in higher education institutions. *Problems and Perspectives in Management*, *23*(1), 101-114. [https://doi.org/10.21511/ppm.23\(1\).2025.08](https://doi.org/10.21511/ppm.23(1).2025.08)

- Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.
- Whitchurch, C. (2010). Optimising the potential of third space professionals in higher education. *Zeitschrift für Hochschulentwicklung*, 5(4). <https://doi.org/10.3217/zfhe-5-04/02>
- Wiley. (2025). *ExplanAltions: Insights from a global survey of 2,430 researchers on generative AI use in academic publishing and research*. John Wiley & Sons.
- Willison, S. (2025, June 14). *Multi-agent research system*. <https://simonwillison.net/2025/Jun/14/multi-agent-research-system/>
- Zahra, S. A., & George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review*, 27(2), 185-203. <https://doi.org/10.5465/amr.2002.6587995>