

# Prompt-Mediated Expertise and Generative AI in Knowledge Work: Credibility, Inequality, and Institutional Legitimacy

Alejandro Martínez

Centre for Research and Teaching in Economics (CIDE), Division of Public Administration, Mexico  
\*alejandro.martinez@cide.edu

## Abstract

*Generative AI has rapidly become part of everyday knowledge work across universities, media, administration, and other professional settings. This expansion matters because AI systems do not simply accelerate output production; they also reshape how competence, credibility, and institutional trust are socially recognized. The purpose of this article is to examine how prompt-mediated expertise reorganizes professional authority, hidden labor, and legitimacy in AI-assisted work environments. The article adopts a qualitative, theory-driven design grounded in the sociology of expertise, organizational analysis, and studies of socio-technical systems. It draws on academic literature, governance documents, institutional materials, and public discussions concerning generative AI, workplace adoption, authorship, and accountability. Analytical attention is directed to prompt literacy, verification labor, institutional approval, access inequality, and the redistribution of discretion across AI-assisted workflows. A comparative interpretive reading is used to clarify how credible outputs are socially stabilized rather than automatically produced by model fluency alone. Prompt-mediated expertise emerges as a distributed form of competence in which human review, institutional validation, and unequal access remain central to the authority of AI-generated outputs. The article concludes that generative AI transforms expertise less by replacing professionals than by reorganizing the social infrastructure through which knowledge becomes trustworthy. Its contribution lies in providing a sociological framework for understanding how prompts, verification work, and institutional legitimacy interact in the contemporary politics of expertise.*

## Keyword

*expertise; generative AI; knowledge work; legitimacy*

## 1. Introduction

Generative AI has moved rapidly from a specialized technical tool to an everyday infrastructure of knowledge production in workplaces, universities, media institutions, and public administration. What makes this transformation sociologically significant is not only the spread of new software, but the reorganization of how competence, credibility, and expert judgment are recognized (Walter, 2024). The uploaded draft frames this shift through the concept of prompt-mediated expertise, where professional problems are increasingly translated into prompts and returned as plausible outputs that circulate through organizational routines (Alavi et al., 2024). In this setting, expertise is no longer judged solely through formal credentials or accumulated experience, but also through the ability to formulate instructions, evaluate outputs, and manage the appearance of competence. This development matters because generative AI does not simply accelerate information processing. It alters the social conditions under which authority is performed, trusted, and institutionalized. The article therefore begins from the premise that expertise in the age of prompts must be understood as a socio-technical formation rather than as a stable attribute possessed by professionals alone (An et al., 2025).



The core problem arises from the growing mismatch between the fluent appearance of AI-generated outputs and the institutional processes required to make those outputs credible (Kshetri, 2024). Generative systems can produce text, summaries, recommendations, and professional-looking artifacts at speed, yet the social trust attached to such outputs remains uneven and politically contested. In many organizations, the adoption of AI is driven by pressures for productivity, responsiveness, and standardization, especially in sectors where communication and interpretation are central forms of labor. This means that the issue is not merely technical accuracy, but the restructuring of accountability in environments where credible knowledge must still be reviewed, attributed, and defended. The real-world relevance of this problem is especially visible in public-facing institutions such as universities, newsrooms, and government offices, where the legitimacy of communication depends on the perception that it is not only efficient but trustworthy (Fadden et al., 2024). Once AI becomes embedded in these settings, institutions face new tensions around disclosure, review, and responsibility. Prompt-mediated expertise thus becomes a political issue because it shapes who is authorized to speak, who must verify, and who bears blame when outputs fail.

A good deal is already known about automation, professional authority, and organizational inequality, and these literatures provide an important base for the article's argument. Earlier studies of office automation demonstrated that technological change does not simply remove labor, but reassigns it, often in ways that intensify managerial control and revalue particular skills. More recent debates on generative AI have shown that prompting is emerging as a new digital competence and that AI systems increasingly shape education, journalism, scientific writing, and administrative work (Rana et al., 2024). The draft also situates its argument within scholarship on classification, platform power, and the politics of visibility, suggesting that AI-generated outputs travel through institutional filters rather than entering social life as neutral objects. What is already clear, then, is that generative AI is not external to organizational hierarchy (Naqbi et al., 2024). It interacts with existing divisions of status, expertise, and access, while also creating new ones. Prompt literacy, verification labor, and access to premium models are already becoming unevenly distributed resources. Existing scholarship therefore establishes that AI changes labor and knowledge, but does not yet fully resolve how those changes reshape the social reality of expertise itself.

What remains less clearly specified is how expertise is reassembled when prompts become a central interface between human judgment and machine-generated content. Many current discussions still treat AI either as a productivity tool that enhances existing professional roles or as a disruptive threat that simply weakens them (Perkins & Roe, 2023). Such framings do not go far enough, because they overlook the institutional processes through which outputs are made credible, attributed, or contested. The article points out that credibility is not inherent in AI-generated text, but must be conferred through review, citation, managerial approval, and other routines of validation. This means the crucial sociological question is not whether AI can produce competent-seeming content, but how organizations decide what counts as competent, authentic, and accountable once such content becomes routine (Inie et al., 2023). There is also limited conceptual clarity about how verification burdens are redistributed, how disclosure norms are negotiated, and how professional boundaries are redrawn when non-specialists can produce specialist-looking artifacts. The unknown dimension, therefore, concerns the social organization of expertise after the prompt becomes a key site of mediation (Cheng et al., 2025).

The research gap becomes sharper when read against the article's core theoretical orientation, which treats prompt-mediated expertise as an infrastructure of credibility rather than a neutral technical convenience. Existing work on generative AI often

emphasizes model performance, ethical risk, or economic disruption, but less often synthesizes these issues into a sociological account of how authority is produced across institutions. The uploaded draft explicitly moves toward such a synthesis by connecting prompting, delegated discretion, invisible verification work, and stratified access to a broader politics of legitimacy. This is an important theoretical move because it shifts attention from isolated users or outputs to the organizational arrangements that make certain uses of AI appear normal, efficient, and trustworthy (Yusuf et al., 2024). In doing so, it identifies expertise as a relational and institutional achievement shaped by prompts, templates, workflows, and approval systems. The gap is therefore not an absence of commentary on AI, but an absence of a sufficiently integrated framework for explaining how prompting reorganizes credibility, inequality, and accountability at the same time. Such a framework is necessary if the sociology of expertise is to keep pace with transformations already underway in knowledge work (Goldenthal et al., 2021).

Filling this gap is theoretically justified because expertise has never been reducible to raw knowledge alone. Sociological traditions have long treated professional authority as an institutional accomplishment grounded in trust, jurisdiction, symbolic boundaries, and systems of validation. Generative AI intensifies this issue by making the production of plausible text easier while leaving the allocation of legitimacy unresolved. The article's intervention is coherent with that tradition because it treats prompting as a mechanism that standardizes interpretive labor while redistributing responsibility across actors with unequal power (Benbya et al., 2024). This matters because discretion does not disappear when AI enters organizational life. It is relocated into prompt design, output selection, editing practices, audit procedures, and managerial standards of acceptability. A theory-sensitive analysis must therefore explain how these new forms of mediation reshape the relationship between knowledge production and social hierarchy. The justification for the article lies in demonstrating that the future of expertise cannot be understood through model capability alone, but through the institutional politics that determine whose judgment still counts and under what conditions.

From this foundation, the article is driven by a set of closely related research questions. How does prompt-mediated expertise reorganize the meaning of competence in workplaces where interpretive labor is central. In what ways do prompting, output review, and attribution practices redistribute discretion across managers, professionals, junior staff, and technical systems. Why do some workers become valued as AI champions while others become invisible verification workers whose labor remains necessary but under-recognized. How does unequal access to models, plugins, datasets, and institutional support transform prompt literacy into a new form of cultural capital. What happens to professional jurisdiction when non-specialists can generate specialist-looking documents that circulate with the appearance of expertise. How do institutions redraw boundaries between authentic and inauthentic communication once AI assistance becomes widespread but unevenly disclosed. These questions allow the article to move beyond the narrow issue of tool adoption and toward the broader issue of how credibility is socially produced in AI-saturated environments.

The urgency of these questions is intensified by the speed with which generative AI is becoming embedded in organizational life, often faster than institutions can develop stable norms for review, disclosure, and accountability. In many sectors, the appeal of AI lies in cost reduction, output expansion, and accelerated responsiveness, yet those same incentives can undermine careful verification and fair allocation of responsibility (Korneeva et al., 2023). The article is especially timely because it highlights how legitimacy is produced through institutions, not through model fluency alone, and warns that poorly governed adoption will tend to amplify existing inequalities. This urgency extends beyond elite organizations in the Global North, since the draft also notes that

resource-constrained settings may adopt AI as a substitute for scarce capacity while simultaneously importing external assumptions that do not fit local contexts. The contribution of the article lies in offering a sociological vocabulary for studying prompts, edits, verification work, and institutional approval as parts of one infrastructure of credibility. It also contributes by reframing expertise as a contested social settlement rather than a stable professional possession (Vesna et al., 2025). The paragraphs that follow are therefore anchored in the claim that the politics of expertise in the age of prompts will be determined by institutional arrangements of trust, blame, and authorization rather than by technical performance alone.

## 2. Research Method

This article employs a qualitative research design grounded in interpretive sociology and institutional analysis to examine how prompt-mediated expertise reorganizes credibility, authority, and professional judgment in AI-assisted knowledge work. A qualitative approach is appropriate because the central concerns of the article involve meanings, organizational routines, symbolic legitimacy, and social relations of verification that cannot be adequately captured through numerical measurement alone (Renjith et al., 2021). The analytical framework draws from the sociology of expertise, science and technology studies, and organizational theory, with particular attention to how socio-technical systems redistribute discretion, invisibilize labor, and reshape institutional trust. Within this framework, prompt-mediated expertise is treated not as a purely technical skill but as a relational and institutional process through which prompts, outputs, review practices, and approval structures jointly produce credible communication.

This design works effectively for the research because generative AI does not merely alter productivity; it reorganizes how competence is recognized, how accountability is assigned, and how authority is performed across workplaces. A qualitative strategy makes it possible to capture these transformations as socially embedded processes rather than as isolated tool effects (Bradbury-Jones et al., 2017). It also allows close attention to the organizational contexts in which prompting, editing, and validation become routinized forms of labor. Such an approach is therefore well suited to a research problem concerned with the social construction of expertise under conditions of rapid technological mediation.

The data consist of academic literature, policy and governance documents, organizational reports, platform and industry materials, institutional guidelines, and publicly available commentaries on generative AI, workplace adoption, professional communication, and knowledge production (Seaman, 1999). Data collection was conducted through purposive document selection, focusing on materials directly relevant to prompting practices, AI-assisted writing and decision support, institutional review processes, disclosure norms, and the redistribution of verification labor in professional environments such as higher education, media, administration, and knowledge-intensive organizations. The units of analysis are discursive and institutional expressions of prompt-mediated expertise, including organizational narratives, workflow models, policy statements, and representations of competence, review, authorship, and accountability in AI-assisted settings. A qualitative coding matrix served as the primary instrument for organizing the material according to key analytical dimensions, namely prompt literacy, delegated discretion, verification labor, institutional legitimacy, access inequality, authorship, and accountability (Shelton et al., 2021). These dimensions functioned as the main variables for tracing how generative AI becomes embedded in organizational routines and how expertise is redefined through socio-technical mediation. Trustworthiness was strengthened through source triangulation, conceptual

consistency in coding, and explicit alignment between the research questions, theoretical framework, and analytical categories, while reliability was supported by maintaining a transparent audit trail of source selection and thematic classification. Validity was enhanced by restricting analysis to sources directly connected to the article's conceptual focus on expertise, institutional legitimacy, and AI-assisted knowledge work (Johnson & Vindrola-Padros, 2017). Because the study relies exclusively on publicly accessible documentary materials and does not involve direct human participants, formal informed consent was not required; nevertheless, ethical standards were maintained by treating potentially sensitive institutional materials with caution, avoiding selective misrepresentation, and preserving confidentiality wherever non-public individual information could otherwise be inferred.

### 3. Result and Discussion

Generative AI is reorganizing expertise less by replacing professional judgment than by redistributing the conditions under which judgment becomes visible, credible, and actionable. In AI-assisted workplaces, competence increasingly includes the capacity to translate a problem into prompts, compare outputs, refine instructions, and align machine-generated text with institutional expectations. This development does not dissolve expertise into technical convenience. It changes the operational pathways through which expertise is performed and recognized. Professional authority now appears through a combination of domain knowledge, prompt literacy, editorial intervention, and organizational validation. The result is a socio-technical model of competence in which expertise is no longer reducible to credentials or experience alone. What matters is not merely who knows, but who can structure interaction with AI in ways that produce acceptable outcomes. Expertise therefore remains central, but its form becomes more distributed and procedurally mediated.

Prompting has emerged as a new layer of expertise because it shapes the quality, direction, and usability of AI-generated outputs. The act of prompting is not a simple mechanical input task, since it involves anticipatory judgment about genre, institutional purpose, audience expectation, and the likely limitations of the model. Skilled prompting often requires iterative refinement, contextual framing, and awareness of how different formulations alter the character of returned material. In many professional settings, this capacity begins to resemble a practical craft that distinguishes effective users from those who rely passively on generic outputs. Such competence becomes particularly valuable where language, analysis, and interpretation are central forms of labor (Familoni & Onyebuchi, 2024). Yet prompt literacy does not operate independently of preexisting expertise. It is most effective when paired with substantive knowledge that allows users to detect weak reasoning, misleading confidence, or contextual mismatch. Prompt-mediated expertise therefore combines procedural fluency with interpretive judgment rather than replacing one with the other.

At the same time, expertise remains institutional rather than merely technical. AI-generated content does not become authoritative because it appears fluent or professionally formatted. It becomes authoritative when organizations validate it through approval chains, editorial standards, policy compliance, and professional norms. This institutional mediation is crucial because generative models produce plausibility rather than accountability. Human actors must still determine whether an output is accurate, appropriate, defensible, and consistent with organizational obligations. In that sense, AI does not bypass institutional legitimacy. It increases dependence on it. The credibility of outputs is therefore inseparable from the settings in which they are reviewed, attributed, and circulated. Expertise remains socially recognized not at the moment of generation, but at the point where institutions decide what counts as usable knowledge.

This shift redistributes discretion across workflows in ways that complicate older understandings of professional responsibility. Discretion no longer resides only in the expert who authors a text or issues a judgment (Broadfoot & Rockey, 2025). It now appears in prompt design, model selection, output comparison, editing decisions, supervisory sign-off, and disclosure practices. Such redistribution creates layered chains of agency in which multiple actors shape the final form of communication without always being equally visible. Responsibility becomes correspondingly diffuse. When an output is persuasive but flawed, it may be unclear whether failure rests with the model, the prompt writer, the reviewer, the manager who approved its use, or the institution that normalized the workflow. This diffusion alters how accountability is organized in professional environments. AI-assisted knowledge work therefore expands rather than resolves the politics of judgment.

The reorganization of expertise also produces new inequalities within workplaces and institutions. Prompt literacy is unevenly distributed, often favoring workers who have time, confidence, institutional encouragement, and access to high-performing tools. Premium subscriptions, plugins, enterprise integrations, and internal training infrastructures create additional layers of advantage (Celik, 2023). Those with better access can produce more polished outputs, experiment more effectively, and present themselves as AI-competent innovators. Others remain dependent on weaker tools or informal learning, which affects both the quality of their work and their visibility within organizational hierarchies. These inequalities are not peripheral effects. They shape whose labor appears efficient, whose competence appears modern, and whose professional authority expands under conditions of technological change. Generative AI therefore becomes entangled with stratified access to resources, recognition, and institutional opportunity.

A further inequality emerges in the allocation of verification labor. Human review remains indispensable because AI-generated material often requires fact-checking, contextual correction, stylistic normalization, and ethical scrutiny before it can circulate safely. Yet this labor is frequently under-recognized precisely because the visible output appears to have been produced effortlessly. Junior staff, assistants, editors, and support workers are often positioned as invisible stabilizers of AI-assisted communication, carrying the burden of ensuring that plausible text becomes institutionally acceptable. Their work is essential, but it may not be rewarded as expertise in the same way as visible prompt use or managerial innovation. This creates a paradox in which AI is marketed as labor-saving while generating new forms of hidden labor beneath the surface of efficiency. The social organization of expertise is thus altered not only at the top, but also through the redistribution of routine but consequential verification work.

**Table 1. Comparative Dimensions of Prompt-Mediated Expertise**

<i>Analytical Dimension</i>	<i>Main Pattern</i>	<i>Institutional Implication</i>
<i>Prompt literacy</i>	Prompting becomes a valued operational skill	Reconfigures what counts as competence
<i>Verification labor</i>	Human review remains essential but often invisible	Shifts burden to junior or support roles
<i>Institutional legitimacy</i>	Credibility depends on approval and governance structures	Expertise remains organizationally mediated
<i>Access inequality</i>	Uneven access to tools and training shapes output quality	Produces stratified adoption and advantage

Table 1 clarifies that prompt-mediated expertise is best understood as a distributed institutional formation rather than an individual technical attribute. Prompt literacy matters because it changes who appears competent under AI-assisted conditions, but it does not stand alone. Verification labor remains necessary to convert fluent outputs into trustworthy communication, which means that expertise continues to depend on human judgment even when that judgment becomes less visible. Institutional legitimacy operates as the central mechanism that binds these processes together. Without approval systems, review standards, and governance rules, AI outputs remain socially unstable regardless of their fluency. Access inequality further demonstrates that AI adoption does not flatten organizational hierarchies. It often re-stratifies them. The table therefore supports the broader argument that expertise is being reassembled through the interaction of technical mediation, hidden labor, and institutional authority.

These patterns resonate with earlier scholarship on automation, digital labor, and the sociology of expertise. Studies of automation long ago established that technology rarely eliminates judgment altogether and more often redistributes labor into new forms of coordination, supervision, and exception handling. Research on professional authority similarly emphasized that expertise is not simply a stock of knowledge, but a socially recognized claim sustained by institutions, credentials, and jurisdictional boundaries. The present discussion extends those insights by showing that generative AI introduces a new interface through which professional legitimacy is negotiated. Prompting becomes a site of operational distinction, while verification and review become critical but unevenly visible forms of institutional maintenance. This also aligns with platform and infrastructure studies that stress how technical systems embed norms, defaults, and asymmetries into everyday practice. Prompt-mediated expertise can therefore be read as a contemporary reconfiguration of longstanding sociological concerns about labor, recognition, and authority. The implications for knowledge work are considerable. In universities, the issue is not simply whether students or academics use AI, but how authorship, evaluation, and academic credibility are redefined when text generation becomes routine. In journalism, public communication, and policy work, institutions must decide how AI-assisted drafting relates to standards of verification, attribution, and editorial responsibility. In administration, professional authority increasingly depends on the ability to manage AI outputs without compromising procedural legitimacy or public trust. Across these sectors, the key question is not whether expertise survives technological mediation. It concerns the organizational conditions under which expertise remains legible, defensible, and fairly distributed. AI-assisted work intensifies the need for governance because output fluency can easily exceed institutional readiness. Professional authority will therefore depend increasingly on how institutions manage the boundary between assistance and authorization.

This makes AI governance a practical rather than abstract concern. Disclosure norms, authorship standards, review protocols, and audit trails become central tools for maintaining credible communication in environments where machine assistance is increasingly routine. Training also matters, but not as a simple skills upgrade. Institutions need governance arrangements that recognize the difference between technical familiarity and accountable expertise. Without such arrangements, workers may be rewarded for visible AI fluency while those performing careful verification remain undervalued. Governance must therefore address not only risk and misuse, but also the internal distribution of labor, recognition, and responsibility. Prompt-mediated expertise requires institutions to redesign systems of oversight so that efficiency does not erode accountability. The practical challenge lies in building workflows where AI assistance can be used without obscuring who is answerable for the final output.

Several strengths and limits emerge from this analysis. One strength lies in its capacity to connect generative AI to broader sociological questions about expertise, legitimacy, and institutional power rather than treating AI adoption as a purely technical matter. Another lies in showing that prompting, review, and organizational approval form an integrated structure rather than separate issues. At the same time, the analysis is limited by the speed of technological change and by the fact that institutional practices are evolving unevenly across sectors. What appears as a stable workflow today may be rapidly transformed by new model capabilities, new governance rules, or new economic pressures. An additional complexity is that AI does not always reduce labor. In many settings it increases editing, checking, and coordination burdens even while promising automation. This unexpected persistence of human labor underscores the need to analyze AI not only through productivity claims but through the social infrastructures that stabilize credibility.

Future inquiry would benefit from closer sector-specific comparison, especially across higher education, media, public administration, and organizations operating under resource constraints. More research is also needed on worker experience, including how staff at different hierarchical levels understand prompt literacy, invisible review labor, and the risks of attribution. Comparative work across Global North and Global South institutions would be especially valuable because infrastructural inequality shapes how AI is adopted, governed, and normalized. There is also a need to examine how authorship norms and audit systems evolve as AI assistance becomes embedded in ordinary communication practices. Practical application should focus on designing governance frameworks that make review labor visible, clarify responsibility, and reduce inequities in access to tools and training. Organizations that treat AI adoption as merely a matter of efficiency will struggle to maintain trust when outputs circulate without credible systems of authorization. The politics of expertise in AI-assisted environments will therefore be decided not only by model capability, but by the institutional arrangements that determine whose judgment remains visible, whose labor remains hidden, and whose authority remains trusted.

#### 4. Conclusion

Prompt-mediated expertise captures a major transformation in contemporary knowledge work, where generative AI does not eliminate expertise but reorganizes how competence, credibility, and accountability are socially produced. The discussion has emphasized that prompting has become a meaningful operational skill, yet its value depends on institutional processes that review, authorize, and stabilize machine-generated outputs. Expertise therefore remains fundamentally social and organizational rather than purely technical. Verification labor continues to be indispensable, even though it is often hidden beneath the visible fluency of AI-assisted communication. At the same time, the spread of prompt-mediated workflows redistributes discretion across prompt design, editing, review, and managerial approval, thereby diffusing responsibility in ways that complicate older professional boundaries. Unequal access to tools, training, and institutional support further intensifies this transformation by turning prompt literacy into a new source of stratified advantage. The reorganization of expertise under generative AI is thus best understood as a socio-technical process in which authority is reassembled through infrastructures of prompting, validation, and institutional legitimacy.

The article contributes to the field by extending the sociology of expertise into the study of generative AI and by offering a framework that connects prompting, invisible verification work, organizational authority, and inequality within a single analytical vocabulary. Rather than treating AI as a neutral productivity tool or as a straightforward

threat to professional judgment, the discussion reframes it as an infrastructure through which credibility is delegated, reviewed, and socially stabilized. This perspective strengthens existing scholarship on automation and digital labor by showing that the central issue is not only technological capability, but the institutional conditions under which machine-generated outputs become trustworthy.

It also adds to debates on professional authority by demonstrating that expertise remains a relational accomplishment grounded in approval systems, authorship norms, and organizational recognition even when AI systems produce increasingly fluent artifacts. The emphasis on prompt literacy and verification labor further clarifies how new hierarchies of competence are emerging within workplaces. In this sense, the article offers both conceptual precision and institutional relevance for understanding how AI is reshaping the social life of knowledge. Its broader significance lies in showing that the future of expertise will be determined not solely by models, but by the governance arrangements that define who can speak credibly and under what conditions.

Future research should move toward more detailed empirical comparison across sectors such as higher education, journalism, public administration, and platform-mediated professional services, where AI-assisted communication is being normalized in different ways. Greater attention is also needed to the lived experience of workers occupying different positions within these workflows, especially those whose verification labor remains essential but under-recognized. Comparative inquiry between Global North and Global South institutions would further deepen understanding of how infrastructural inequality shapes access, training, governance capacity, and the social legitimacy of AI-generated outputs.

Another important direction concerns the evolution of authorship norms, audit mechanisms, and disclosure practices as organizations attempt to manage the boundary between assistance and accountability. Longitudinal research would be valuable for tracking whether prompt literacy stabilizes as a durable form of professional capital or becomes routinized into broader systems of organizational control. Practical development should focus on governance models that make human oversight visible, distribute responsibility more clearly, and prevent efficiency gains from obscuring new inequalities in labor and recognition. Advancing this agenda remains essential for understanding how institutions can integrate generative AI without eroding the trust on which expertise ultimately depends.

## References

- Alavi, M., Leidner, D., & Mousavi, R. (2024). Knowledge Management Perspective of Generative Artificial Intelligence. *J. Assoc. Inf. Syst.*, 25, 15. <https://doi.org/10.17705/1jais.00859>
- An, Y., Yu, J. H., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22. <https://doi.org/10.1186/s41239-025-00507-3>
- Benbya, H., Strich, F., & Tamm, T. (2024). Navigating Generative Artificial Intelligence Promises and Perils for Knowledge and Creative Work. *J. Assoc. Inf. Syst.*, 25, 13. <https://doi.org/10.17705/1jais.00861>
- Bradbury-Jones, C., Breckenridge, J., Clark, M., Herber, O., Wagstaff, C., & Taylor, J. (2017). The state of qualitative research in health and social science literature: a focused mapping review and synthesis. *International Journal of Social Research Methodology*, 20, 627–645. <https://doi.org/10.1080/13645579.2016.1270583>

- Broadfoot, P., & Rockey, J. (2025). Generative AI and the social functions of educational assessment. *Oxford Review of Education*, 51, 281–300. <https://doi.org/10.1080/03054985.2025.2455549>
- Celik, I. (2023). Exploring the Determinants of Artificial Intelligence (AI) Literacy: Digital Divide, Computational Thinking, Cognitive Absorption. *Telematics Informatics*, 83, 102026. <https://doi.org/10.1016/j.tele.2023.102026>
- Cheng, C.-C., Wang, J.-S., Zhai, X., & Yang, Y. (2025). AI literacy and gender equity in elementary education: A quasi-experimental study of a STEAM-PBL-AIoT course with questionnaire validation. *International Journal of STEM Education*, 12. <https://doi.org/10.1186/s40594-025-00574-y>
- Fadden, I., García-Alonso, E., & Meneses, E. L. (2024). Science Mapping of AI as an Educational Tool Exploring Digital Inequalities: A Sociological Perspective. *Multimodal Technol. Interact.*, 8, 106. <https://doi.org/10.3390/mti8120106>
- Familoni, B. T., & Onyebuchi, N. C. (2024). ADVANCEMENTS AND CHALLENGES IN AI INTEGRATION FOR TECHNICAL LITERACY: A SYSTEMATIC REVIEW. *Engineering Science & Technology Journal*. <https://doi.org/10.51594/estj.v5i4.1042>
- Goldenthal, E., Park, J., Liu, S., Mieczkowski, H., & Hancock, J. (2021). Not All AI are Equal: Exploring the Accessibility of AI-Mediated Communication Technology. *Comput. Hum. Behav.*, 125, 106975. <https://doi.org/10.1016/j.chb.2021.106975>
- Inie, N., Falk, J., & Tanimoto, S. (2023). Designing Participatory AI: Creative Professionals' Worries and Expectations about Generative AI. *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3544549.3585657>
- Johnson, G., & Vindrola-Padros, C. (2017). Rapid qualitative research methods during complex health emergencies: A systematic review of the literature. *Social Science & Medicine*, 189, 63–75. <https://doi.org/10.1016/j.socscimed.2017.07.029>
- Korneeva, E., Salge, T., Teubner, T., & Antons, D. (2023). Tracing the legitimacy of Artificial Intelligence: A longitudinal analysis of media discourse. *Technological Forecasting and Social Change*. <https://doi.org/10.1016/j.techfore.2023.122467>
- Kshetri, N. (2024). The academic industry's response to generative artificial intelligence: An institutional analysis of large language models. *Telecommunications Policy*. <https://doi.org/10.1016/j.telpol.2024.102760>
- Naqbi, H. Al, Bahroun, Z., & Ahmed, V. (2024). Enhancing Work Productivity through Generative Artificial Intelligence: A Comprehensive Literature Review. *Sustainability*. <https://doi.org/10.3390/su16031166>
- Perkins, M., & Roe, J. (2023). Academic publisher guidelines on AI usage: A ChatGPT supported thematic analysis. *F1000Research*, 12. <https://doi.org/10.12688/f1000research.142411.2>
- Rana, N., Pillai, R., Sivathanu, B., & Malik, N. (2024). Assessing the nexus of Generative AI adoption, ethical considerations and organizational performance. *Technovation*. <https://doi.org/10.1016/j.technovation.2024.103064>
- Renjith, V., Yesodharan, R., Noronha, J., Ladd, E., & George, A. (2021). Qualitative Methods in Health Care Research. *International Journal of Preventive Medicine*, 12. [https://doi.org/10.4103/ijpvm.ijpvm\\_321\\_19](https://doi.org/10.4103/ijpvm.ijpvm_321_19)
- Seaman, C. (1999). Qualitative Methods in Empirical Studies of Software Engineering. *IEEE Trans. Software Eng.*, 25, 557–572. <https://doi.org/10.1109/32.799955>
- Shelton, R., Philbin, M., & Ramanadhan, S. (2021). Qualitative Research Methods in Chronic Disease: Introduction and Opportunities to Promote Health Equity. *Annual Review of Public Health*, 43, 37–57. <https://doi.org/10.1146/annurev-publhealth-012420-105104>
- Vesna, L., Sawale, Priti., Kaul, P., Pal, S., Srv, B., & Murthy, R. (2025). Digital Divide in AI-Powered Education: Challenges and Solutions for Equitable Learning. *Journal of Information Systems Engineering and Management*. <https://doi.org/10.52783/jisem.v10i21s.3327>

- Walter, Y. (2024). Embracing the future of Artificial Intelligence in the classroom: the relevance of AI literacy, prompt engineering, and critical thinking in modern education. *International Journal of Educational Technology in Higher Education*, 21, 1-29. <https://doi.org/10.1186/s41239-024-00448-3>
- Yusuf, A., Pervin, N., & Román-González, M. (2024). Generative AI and the future of higher education: a threat to academic integrity or reformation? Evidence from multicultural perspectives. *International Journal of Educational Technology in Higher Education*, 21, 1-29. <https://doi.org/10.1186/s41239-024-00453-6>