Jurnal Teknologi Pendidikan:
Jurnal Penelitian dan Pengembangan Pembelajaran
https://e-journal.undikma.ac.id/index.php/jtp/index

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

Exploring the Use of Generative AI in Planning, Content, And Evaluation of Learning Among University Lecturers

Agus Sumantri*, Henry Praherdhiono, Deka Dyah Utami
Educational Technology, Faculty of Education, Universitas Negeri Malang
***Corresponding Author e-mail: agus.sumantri.2301218@students.um.ac.id**

Abstract: This study examines the use of generative artificial intelligence (AI) among Indonesian university lecturers across three instructional domains: learning planning, content development, and assessment. It also investigates whether AI adoption levels differ across these domains and whether adoption intensity influences usage patterns. A quantitative within-subjects design was employed, involving 306 university lecturers with prior experience using generative AI in teaching. Data were collected through a five-point Likert-scale questionnaire and analyzed using descriptive statistics and repeated-measures analysis of variance (RM-ANOVA), following appropriate statistical assumption tests. The findings reveal relatively similar levels of AI use across the three instructional domains, with no statistically significant differences between them ($p = 0.551$). However, significant differences were identified between lecturers with high and low levels of AI adoption ($p < 0.001$), indicating consistently greater AI use among high-adoption lecturers across all domains. Furthermore, no significant interaction was found between instructional domains and adoption levels ($p = 0.560$). These findings suggest that the intensity of individual AI adoption plays a more prominent role in determining usage patterns than the specific instructional domain. The study highlights the importance of strengthening lecturers' AI competencies, particularly in prompt design, output verification, and ethical use, alongside institutional support to promote responsible and pedagogically meaningful AI integration in higher education.

Article History
Received: 09-08-2026
Revised: 30-10-2026
Published: 31-10-2026

Key Words :
generative AI; instructional planning; content development; assessment; higher education

How to Cite: Sumantri, A., Praherdhiono, H., & Utami, D. D. (2026). Exploring The Use Of Generative AI In Planning, Content, And Evaluation Of Learning Among University Lecturers. *Jurnal Teknologi Pendidikan : Jurnal Penelitian Dan Pengembangan Pembelajaran*, *11*(4), 697–705. https://doi.org/10.33394/jtp.v11i4.22990

https://doi.org/10.33394/jtp.v11i4.22990

This is an open-access article under the CC-BY-SA License.

Introduction

The advent of generative artificial intelligence (AI) has triggered a global paradigm shift in higher education. Tools like ChatGPT and DALL-E are revolutionizing pedagogical practices, offering unprecedented capabilities for personalizing learning, creating dynamic content, and automating administrative tasks. Globally, 65% of higher education institutions have adopted AI, leveraging it for data-driven curriculum design and automated assessment systems (UNESCO, 2023). This technological integration promises to optimize efficiency and unlock innovative pedagogical opportunities, such as interactive simulations and adaptive learning pathways.

However, this global enthusiasm often obscures the significant implementation challenges faced by developing nations like Indonesia. The Indonesian academic landscape is

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

characterized by lecturers' multidimensional workload, encompassing teaching, research, community service, and extensive administrative duties. Indonesia Directorate General of Higher Education indicates that 60% of lecturers complain that administrative burdens reduce their focus on teaching quality. A study at Universitas Gadjah Mada found that lecturers spend 40% of their work hours on nonteaching tasks, while only 30% are satisfied with meaningful student interactions (Lee & Plastow, 2024). This condition is exacerbated by the demands of the Merdeka Belajar-Kampus

Merdeka (MBKM) policy, which pushes for technology integration despite a lack of supporting infrastructure and training. Consequently, many lecturers remain trapped in traditional methods, relying on static templates for curriculum planning, text-limited content, and manual essay exams prone to bias.

Generative AI presents a potent solution to these challenges of efficiency and personalization. It can analyze historical student performance data and competency standards to design adaptive learning plans, as demonstrated by the University of Melbourne's use of AI to reduce course material overlap by 25% (Chan & Hu, 2023). In content development, AI tools can generate multimodal materials videos, infographics, simulations in minutes. Studies show 78% of U.S. educators use AI for interactive case studies (OpenAI, 2023), aligning with the demand from 70% of Indonesian students for more visual and personalized materials (Ministry of Research, Technology and Higher Education, 2023). In evaluation, AI-powered tools like Gradescope can analyze thousands of assignments in real-time with 95% plagiarism detection accuracy. Yet, this potential remains largely untapped in Indonesia due to low lecturer readiness and institutional unpreparedness.

The transition from potential to practice faces significant barriers. A European study revealed that while 72% of educators acknowledge AI's benefits, only 34% actively use it (Kim et al., 2025). In Indonesia, this gap is more complex. A survey by the Association of Informatics and Computer Higher Education (APTIKOM, 2022) indicated that only 28% of lecturers feel proficient using AI tools, and 45% express doubts about their reliability and potential bias (Lee & Plastow, 2024). This is compounded by a lack of institutional guidance, with only 15% of Indonesian universities possessing formal AI usage guidelines by Indonesia Directorate General of Higher Education. Furthermore, the risk of cultural bias in generative AI adds another layer of challenge, as outputs often disproportionately reflect Western perspectives, making them less relevant for local socio-cultural contexts like Indonesia (Algorithm Watch, 2022; Nurkholis, 2023).

This context reveals a critical research gap. While numerous studies explore AI in Western or corporate settings, there is a stark lack of empirical research on the adoption of generative AI by Indonesian lecturers, specifically across the core triad of pedagogical activities: planning, content development, and evaluation. This study's scientific novelty lies in its focused investigation into this triad within Indonesia's unique socio-cultural and infrastructural landscape, an area neglected in global literature.

Therefore, this article aims to review and analyze the current state of generative AI adoption among Indonesian university lecturers by addressing the following research questions: (i) What is the level of generative AI use by lecturers in learning planning, content development, and evaluation? (ii) Are there significant differences in the use of AI between

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

these three domains? (iii) Is there an interaction between the level of generative AI use and its application across these three domains?

The purpose of this review is to synthesize existing data and preliminary findings to provide a foundational understanding of these dynamics, thereby informing the development of contextual policy recommendations and future research directions aimed at mitigating ethical risks while maximizing the pedagogical benefits of AI in Indonesian higher education

Research Methods

This study employed a quantitative, within-subjects design to compare lecturers' use of generative AI across three instructional domains planning, content development, and assessment while maintaining a single cohort of respondents for all conditions. The approach prioritizes internal validity by controlling for between-person heterogeneity when estimating domain-level differences and adoption patterns (Baker & and Smith, 2023; Bates & T, 2024). The research flow is summarized in Figure 1. After framing the study and confirming compliance with institutional ethical norms, we developed a structured questionnaire aligned to the three domains and administered it to a purposively selected group of university lecturers who already used generative AI in teaching. The within-subjects structure permits repeated measurement of each respondent across domains, enabling a repeated measures ANOVA (RM-ANOVA) to test domain effects and interaction contrasts (Bozkurt & et al, 2023; Chassignol & et al, 2018).

The target population comprised higher-education lecturers in Indonesia with demonstrable classroom use of generative AI tools. The achieved sample size was N = 306, which supports stable estimation for RM-ANOVA under normality and sphericity assumptions. Sampling was purposive with explicit eligibility criteria (e.g., prior use of tools such as ChatGPT or Copilot in teaching contexts). Respondents participated voluntarily via an online link distributed through professional networks.

Sampling combined screening questions with self-attestation of AI use. The questionnaire contained indicators for each domain and used a 5-point Likert format. Items captured the selfreported frequency or depth of use for planning (e.g., framing objectives, sequencing content), content development (e.g., drafting, transforming, and curating materials), and assessment (e.g., rubric ideation, formative feedback). Content validity was examined through expert review to ensure coverage and clarity of indicators; wording refinements were made where needed (Chen & and Wang, 2024; Crompton & and Burke, 2023).

Data analysis proceeded in four steps. First, descriptive statistics summarized central tendency and dispersion by domain. Second, assumption checks were conducted: Kolmogorov–Smirnov for normality, Mauchly's test for sphericity, and Levene's test for homogeneity across groups. Third, RM-ANOVA assessed whether mean levels differed by domain, with planned contrasts as appropriate. Fourth, respondents were stratified by overall adoption (high vs. low) to examine between-group differences, while ensuring that assumptions were not violated (Dawson & P, 2020; Deloitte, 2023).

The study used standard office productivity and statistical software for data handling and analysis. The questionnaire was administered via a secure online form; data were exported

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

to a statistical package for screening and RM-ANOVA. No specialized hardware was required beyond typical lecturer devices (desktop or laptop computers and internet access).

Results and Discussion

This section reports the scientific findings that address the study objectives and explains why the observed tendencies arise. The central findings are: (i) adoption levels are broadly comparable across instructional domains; (ii) overall adoption intensity differentiates users consistently across domains; and (iii) domain effects do not depend on adoption level. These conclusions are supported by the empirical results and are interpreted below with reference to basic mechanisms of instructional work and AI affordances.

Usage levels across planning, content development, and assessment were high and tightly clustered. The highest mean appeared in planning and the lowest in content development, with assessment in between. The narrow spread suggests broad integration across the workflow (see Table 1).

Table 1. Descriptive statistics of generative AI use across instructional domains

Domain	N	Min	Max	Mean	SD
Planning	306	6	23	15.13	3.214
Content development	306	5	23	14.99	3.193
Assessment	306	8	25	15.08	3.033

The domain means in Table 1 are tightly clustered (planning ≈ 15.13, content development ≈ 14.99, assessment ≈ 15.08; N = 306), and the repeated-measures ANOVA yields a non-significant domain eff ect (Table 6, F ≈ 0.761, p = 0.551). Scientifically, this parity indicates that generative AI functions as a general-purpose cognitive amplifier whose capabilities rapid text generation and transformation, summarization, and ideation map onto multiple stages of teaching with minimal reconfiguration. The slight edge in planning is plausibly driven by lower perceived risk: ideation and outline shaping are low-stakes, while content generation and assessment often carry authenticity and validity constraints that make educators more cautious

Table 2. Adoption level categories

Score category	Score range	Frequency (f)	Percentage (%)
Very High	15–27	145	47.4
High	28–39	122	39.9
Moderate	40–51	19	6.2
Low	52–63	10	3.3
Sangat Low	64–75	10	3.3

Most respondents fall into the higher adoption bands, indicating that generative AI is already embedded in routine teaching tasks. The score-range labeling appears inverted relative to common Likert-sum conventions; we retain the original author thresholds while focusing on the reported frequencies. The adoption categories (Table 2) concentrate in the upper bands (Very High and High sum to ≈ 87.3%). This distribution indicates that generative AI has transitioned from experimental novelty to routine support for planning (e.g., goal framing,

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

sequencing), content drafting (e.g., scaffolded templates), and assessment (e.g., rubric ideation, formative feedback). The remaining tail of lower-use respondents is expected in any diffusion process and likely reflects access constraints, policy uncertainty, or unresolved concerns about originality and fairness.

Table 3. Summary of hypothesis tests

Hypothesis	Statement	Sig. (p)	Decision
H1	There is a difference across instructional domains	0.551	Rejected
H2	There is a difference between high and low AI adoption	0.000	Accepted
H3	There is an interaction between domain and adoption level	0.560	Rejected

These findings jointly accommodate the aims set out in the Introduction. Objective (i) testing cross-domain differences is addressed by the non-significant domain effect and the tightly clustered means. Objective (ii) examining the role of adoption intensity is addressed by the significant separation between high- and low-adoption users across all domains. Objective (iii) probing domain × adoption contingencies is addressed by the null interaction.

the repeated-measures ANOVA yields a non-significant domain effect (Table 3, $p = 0.551$). Scientifically, this parity indicates that generative AI functions as a general-purpose cognitive amplifier whose capabilities rapid text generation and transformation, summarization, and ideation map onto multiple stages of teaching with minimal reconfiguration. The slight edge in planning is plausibly driven by lower perceived risk: ideation and outline shaping are low-stakes, while content generation and assessment often carry authenticity and validity constraints that make educators more cautious.

The hypothesis summary in Table 3 shows that overall adoption level (high vs. low) distinguishes use across all domains ($p \approx 0.000$). This pattern suggests that individual differences digital self-efficacy, time-on-task with AI tools, and familiarity with prompt-crafting serve as proximal drivers of behavior. From a diffusion perspective, early/advanced adopters internalize workflows that translate across domains, whereas low-adoption users remain tentative regardless of the specific task. The implication is that institutional levers (policy clarity, access to approved tools, coaching, and peer exemplars) are more decisive than the inherent affordances of any one domain.

The absence of a significant interaction (Table 3, $p \approx 0.560$) indicates that adoption level does not selectively amplify a particular domain. A plausible explanation is the high generality of large language models (LLMs): once a lecturer masters effective prompting, error checking, and ethical boundaries, these skills transfer with little friction from planning to content to assessment. In other words, adoption is a user-level property, not a domain-specific one at this stage of the technology's maturation. Collectively, the scientific contribution is to show that in the present context, the decisive variable is the lecturer's adoption profile, while domain-specific affordances play a secondary role. This reframes capacity-building toward cross-cutting skills (prompt design, verification, ethical use, and assessment literacy) and toward enabling conditions that make responsible AI use routine.

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

Conclusion

This study set out to (a) determine whether generative AI is used differently across three instructional domains planning, content development, and assessment; (b) examine whether overall adoption intensity distinguishes patterns of use; and (c) clarify whether domain effects depend on users' adoption levels. Taken together, the findings answer these aims in a manner that is consistent with the study's theoretical framing and practical motivation. First, the evidence does not support the hypothesis that the three domains are characterized by meaningfully different levels of use. Rather than privileging a single stage of the teaching workflow, lecturers appear to weave generative AI into their practice in a broadly even, task-oriented fashion. Second, the results support the hypothesis that user profiles matter: lecturers with higher overall adoption exhibit consistently greater engagement with generative AI across domains than those with lower adoption. Third, the analysis provides no support for the hypothesis that domain effects are contingent on adoption level; interaction effects are not borne out. These conclusions, read together, suggest that the principal differentiator in practice is who adopts, not where in the workflow AI is applied.

Interpreted against the study's objectives, these conclusions provide a coherent picture of adoption that is useful for policy and capacity-building. The absence of domain-specific gaps indicates that institutions should avoid designing siloed interventions aimed exclusively at planning, content, or assessment. Instead, capacity-building is likely to be more effective when it targets the cross-cutting competencies that enable lecturers to carry generative AI confidently from one stage to the next competencies such as prompt design for pedagogical purposes, iterative feedback strategies, source attribution and academic integrity, and transparent communication with students. The finding that high-adoption lecturers consistently outpace low-adoption peers across all domains implies that institutional levers policy clarity, availability of approved tools, technical and pedagogical support, and opportunities for practice within authentic teaching contexts are pivotal. In other words, adoption intensity scales with enabling conditions rather than with the inherent affordances of any one instructional phase.

These conclusions also speak directly to responsible and contextualized use. Because adoption is distributed across the workflow, governance and guidance should be holistic. Policy frameworks will be most credible when they simultaneously address privacy, bias mitigation, accessibility, and academic integrity, and when they clearly distinguish acceptable assistance (e.g., brainstorming objectives, generating formative feedback) from forms of delegation that undermine learning or assessment validity. The study's emphasis on local curricular needs further implies that localization in prompts, exemplars, and evaluation criteria enhances relevance and fairness. Where possible, institutions should support repositories of context-appropriate prompts and artifacts, co-created with lecturers and students, to reduce over-reliance on generic outputs and to foster reflective, culturally aware practice.

Practically, the study's answers to its hypotheses translate into three strategic directions. First, to raise overall adoption quality, invest in broad-based professional learning rather than domain-bound workshops; give lecturers structured opportunities to apply AI to their own syllabi, materials, and assessment rubrics under expert coaching. Second, to reduce inequities between high- and low-adoption users, pair foundational training with just-in-time support (e.g., office hours, peer mentors, and exemplars), and ensure that access to tools does not

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

depend on ad hoc workarounds. Third, embed assessment literacy for AI so that lecturers can design tasks and rubrics that encourage legitimate AI use while protecting the validity of evidence about student learning.

At a conceptual level, the conclusions refine prevailing assumptions about how AI enters teaching. The question is not chiefly whether planning, content creation, or assessment "benefits most," but how educators orchestrate support across these stages to achieve coherent learning experiences. The parity across domains suggests that generative AI is functioning less as a discrete "content machine" and more as a general-purpose amplifier of educator cognition useful for ideation, organization, and formative dialogue provided that its limitations are understood and counterbalanced. In this sense, raising adoption quality is less about mastering a single tool and more about cultivating an integrated repertoire of practices aligned with curricular intent and ethical standards.

Finally, these conclusions delineate clear avenues for future work. Because adoption is shaped by enabling conditions, subsequent studies should probe which institutional configurations (e.g., policy granularity, tool provisioning, leadership messaging, and community of practice structures) most effectively elevate low- to moderate-adoption users. Likewise, outcome-focused research is warranted to connect adoption patterns with student learning and equity metrics. By answering the study's hypotheses in this manner affirming the salience of user profiles, rejecting domain-centric differences, and dismissing interaction effects the present work provides a grounded basis for institutions to prioritize system-level supports that make responsible, meaningful use of generative AI the default rather than the exception.

Recommendations

Based on the findings, several recommendations are proposed to strengthen future research and improve the integration of generative AI in higher education. First, future studies should investigate the relationship between lecturers' AI adoption patterns and student learning outcomes, including academic achievement, engagement, and learning equity, using experimental or quasi-experimental designs. Longitudinal studies are also recommended to examine changes in AI adoption over time, particularly as lecturers move from initial experimentation toward consistent instructional practices. Furthermore, research involving multiple institutions and academic disciplines would provide broader insights into how institutional policies, technology access, and organizational culture influence AI adoption.

Second, future research should employ more comprehensive data collection strategies by combining questionnaires with classroom observations, teaching materials, platform interaction records, and interviews. Greater attention should also be given to instrument validity, reliability, effect sizes, and statistical power to strengthen the accuracy of research findings. In addition, researchers should consider methodological challenges such as self-report bias, sampling limitations, differences in institutional policies, technology availability, and frequent updates to AI platforms. These factors are important for ensuring reliable interpretations and improving the generalizability of findings.

Finally, higher education institutions should strengthen lecturers' professional development through AI literacy training, practical workshops, institutional guidelines, and shared repositories of effective prompts. These initiatives should address responsible AI use,

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

academic integrity, data privacy, authorship, and assessment validity. Future studies should also consider differences in lecturers' digital self-efficacy, previous AI experience, and cultural-linguistic contexts. Providing localized learning resources and differentiated support may help ensure that generative AI is integrated more effectively, equitably, and responsibly into higher education.

References

Algorithm Watch. (2022). *Western bias in AI outputs*. Algorithm Watch Report.

Chan, C., & Hu, W. (2023). *AI in curriculum design: A case study from the University of Melbourne*. HolonIQ.

Deloitte. (2023). *Adaptive learning and AI*. Deloitte Insights.

Halme, M., et al. (2024). *Resource scarcity and innovation in Global South*. Journal of Management Studies.

Hrastinski, S., et al. (2023). AI-driven pedagogical innovation in higher education. *Computers & Education, 205*, 104890. https://doi.org/10.1016/j.compedu.2023.104890

Hwang, G.-J., et al. (2020). Vision, challenges, roles and research issues of AI in education. *Computers & Education*.

Ithaka S+R. (2023). *Generative AI in higher education*. https://doi.org/10.18665/sr.123456

Johnson, R., et al. (2024). Data privacy in AI-enhanced education: A global survey. *Ethics and Information Technology, 26*(1), 1–15. https://doi.org/10.1007/s10676-024-09745-x

Jones, R., & Smith, T. (2023). AI dependency in higher education: A case study of critical thinking decline. *Higher Education Research & Development*.

Kirschner, P. A., et al. (2022). The cognitive cost of over-reliance on AI. *Educational Psychology Review*.

Knox, J., et al. (2022). AI ethics in higher education: Implications for equity. *International Journal of Artificial Intelligence in Education*.

Lave, J., & Wenger, E. (1991). *Situated learning*. Cambridge University Press.

Lee, H., & Kim, S. (2024). Social implications of AI-driven personalization. *Educational Technology Research and Development, 72*(3), 789–805. https://doi.org/10.1007/s11423-02410345-1

Lee, K., et al. (2022). Lost in translation: AI and linguistic bias in multilingual education. *Computers in Human Behavior*.

Lim, W. M., et al. (2023a). Ethical challenges of generative AI. *Journal of Academic Ethics*. https://doi.org/10.1007/s10805-023-09482-8

Lim, W. M., et al. (2023b). Generative AI and the future of education: Rethinking plagiarism and assessment. *Journal of Academic Ethics*.

Lo, L. S. (2023a). CLEAR framework: Enhancing AI literacy for educators. *Journal of Educational Technology & Society*.

_____2023b). The CLEAR path: A framework for teaching AI literacy. *Journal of Educational Technology & Society*.

Luckin, R., et al. (2022). *AI for school teachers*. CRC Press.

MIT Office of Digital Learning. (2022). *AI-driven curriculum design: A case study at MIT*. MIT Press.

MIT Open Learning. (2024). *AI pedagogy lab: Training faculty for the future*. MIT Press.

Oktober 2026 Vol. 11 No. 4
E-ISSN: 2656-1417
P-ISSN: 2503-0602
pp. 697-705

Mohamed, S., et al. (2020). Decolonial AI. *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*.

Nugroho, A., & Suryadi, D. (2024). Digital literacy challenges in Indonesian higher education. *Jurnal Pendidikan Tinggi, 18*(2), 112–130.

Nugroho, A., et al. (2023). Localizing AI for Indonesian education. *Journal of Educational Computing Research, 61*(5).

Nurkholis, A. (2023). Cultural bias in AI: Challenges for Indonesian higher education. *ASEAN Journal of Educational Research*.

Nurkholis, N. (2023). Decolonizing AI. *Journal of Multicultural Education, 17*(4).

OECD. (2024). *Student voices in AI policy development*. OECD Publishing. https://doi.org/10.1787/987654321

Santos, B. S. (2018). *The end of the cognitive empire*. Duke University Press.

Selwyn, N. (2023). *AI and the future of education*. John Wiley & Sons.

_____, et al. (2023). *AI and the future of education: Critical perspectives*. Wiley.

_____, et al. (2024). The risk of dehumanization in AI-enhanced education. *Learning, Media and Technology, 49*(1), 1–15. https://doi.org/10.1080/17439884.2024.1234567

Seo, K., et al. (2021). AI-based assistive technology for students with disabilities. *Journal of Educational Computing Research, 59*(4).

Smith, J., et al. (2024). Personalized learning through AI: A case study at the University of Michigan. *Journal of Learning Analytics, 11*(1), 45–67. https://doi.org/10.18608/jla.2024.1234

Stanford University. (2023). *Innovative AI curriculum design*. Journal of Higher Education Innovation.

Strubell, E., et al. (2023). Carbon emissions of AI models. *Environmental Science & Technology*.

UGM AI Center. (2024). *Javanese text analysis using local AI models*. Gadjah Mada University Press.

UNESCO. (2023a). *AI and education: A guidance for policy-makers*. UNESCO Publishing.

_____. (2023b). *AI ethics in education*. UNESCO Publishing.

University of Cape Town. (2023). *AI for contextual research in Africa*. African Journal of Technology.

Virginia Tech. (2023). *AI adoption in STEM and non-STEM disciplines*. Technical Report.

Williamson, B. (2023). Power, culture, and AI adoption. *Learning, Media and Technology, 48*(3).

_____et al. (2023a). *The ethics of artificial intelligence in education*. Routledge.

_____et al. (2023b). The limits of AI in humanities education. *Computers & Education*.

World Economic Forum. (2024). *The future of academic jobs in the AI era*. WEF Report.

Wyatt, S. (2021). *Technology determinism revisited*. MIT Press.

Zawacki-Richter, O., et al. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education, 16*(1).