



Course-Specific AI TAs: A Platform Comparison for Higher Education

Joon S. Park

School of Information Studies, Syracuse University, Syracuse, NY, United States

Abstract

As generative AI tools increasingly spread across higher education and become more widely available, educators are confronted with a substantial decision: Which AI platform best facilitates the implementation of course-specific AI teaching assistants (AI TAs) in their courses? This study seeks to compare available options as of today by designing and evaluating three distinct course-specific AI TAs: OpenAI Custom GPT, Anthropic Claude Projects, and Google Gemini Gems, all implemented within the same college-level cybersecurity course. Each AI TA was meticulously configured with course-specific instructional objectives, lecture materials, a course syllabus, an academic schedule, and course-related resources. We assessed the following governance-relevant criteria: customization depth, ease of setup, course resource integration, interface & accessibility, multimodal capabilities, privacy & ethics controls, institutional cost, and pedagogical impact. At the time of this writing, our comparative analysis finds that Custom GPT stands out as the most fully featured no-code platform for a deeply customizable, tightly governed, shareable, course-specific AI TA deployment. Claude Project has closed the gap, particularly with the addition of Skills, though it remains more constrained in Project sharing across organizations/institutions and dynamic usage. Gemini Gems offers the strongest multimodal and Google-native integration but remains limited by upload, storage, and context-window constraints. Overall, this study offers guidance to instructors and institutions on deploying AI TAs that enhance learning experiences and outcomes while safeguarding academic integrity and promoting accessibility.

Keywords: *Course-specific AI Teaching Assistant (TA), AI-augmented Higher Education*

1. Introduction

Artificial intelligence (AI) is reshaping higher education by expanding the ways students seek explanations, practice skills, and receive feedback. Students now encounter general-purpose conversational systems as everyday academic tools, yet those systems are not designed primarily for course-grounded instruction, institutional governance, or academic integrity [1], [2]. For instructors, then, the central question is no longer whether an AI system can answer student questions. The more consequential question is whether it can function as a course-specific assistant that remains grounded in the syllabus, lecture content, and course policies; supports learning without producing graded solutions; and operates in ways that are consistent with institutional expectations for privacy, access, and oversight. These concerns have become more pressing as AI-supported tools move into mission-critical and educational settings, where system failures, hallucinations, and misuse can have consequences that extend well beyond convenience or productivity [3], [4], [5]. In higher education, the value of such tools depends not only on what they can produce, but also on whether they can be integrated into teaching in ways that are pedagogically sound, ethically defensible, and practically sustainable [6], [7], [8].

This paper examines those challenges by building and comparing three no-code, course-specific AI teaching assistants (TAs) for the same college course: OpenAI Custom GPTs, Anthropic Claude Projects with Skills, and Google Gemini Gems. All three platforms now allow instructors to define custom instructions and attach course-related materials, but they differ in meaningful ways with respect to persistence of course knowledge, sharing and access, multimodal support, and administrative controls. Our purpose is not to identify a single universal winner. Rather, we seek to clarify which trade-offs matter most when instructors and institutions select a platform for a course-specific AI TA.

This study builds directly on our earlier work on course-specific AI TAs and AI-augmented teaching [9], [10], [11]. In prior research, we developed and evaluated a course-grounded AI TA based on Custom GPT, with particular attention to implementation choices and effects on student learning [9]. Related work also examined broader pathways toward AI-augmented teaching in higher education and the practical design choices required for responsible instructional use [10]. More recently, the same design logic has been extended to other educational contexts, including medical education, suggesting



that course-specific AI TAs can be adapted across disciplinary settings while retaining the same underlying concerns about grounding, pedagogy, and governance [11]. The present study builds on that line of work by moving from a single platform case to a comparative, cross-platform analysis.

Observations from practice motivate this comparison. The educational value of an AI TA depends on grounding in the course-specific context. A helpful course assistant should begin with the course's own materials, including the syllabus, lecture slides, lab guides, research projects, and other related academic resources, before relying on external resources. Furthermore, academic integrity cannot be simply assumed. Default assistants may be helpful yet still permissive, which means support patterns and coaching strategies must be deliberately designed and implemented in line with the instructor's teaching plan and philosophy. Additionally, institutional constraints, such as licensing, device configuration, privacy protections, ethical issues, and instructor oversight, shape whether an apparently useful system can be adopted fairly and responsibly across a course or program. Therefore, the evaluation focuses on several criteria that, in our judgment, best predict classroom usefulness and institutional adoptability: knowledge capacity and persistence; internal grounding and citation fidelity; external scholarly retrieval and citation completeness; academic integrity and ethics; mobility across devices; multimodal utility; and cost and usability within the institutional setting. We intentionally do not foreground timing comparisons or decontextualized accuracy leaderboards. In this course setting, those measures were less informative than the broader question of whether a platform could function as a dependable instructional assistant under realistic teaching constraints.

This study makes three contributions. First, it provides a replicable approach to evaluating course-specific AI teaching assistants under realistic instructional constraints. Second, it offers a platform-aware analysis that links observed behavior to governance choices, including what can be achieved through no-code configuration alone and when institutions may need to migrate toward SDK or LMS-based deployments for retrieval augmentation, logging, and learning analytics. Third, it offers practical guidance for instructors and institutions by identifying the conditions under which each platform is most useful and the trade-offs that accompany those choices. In this sense, the study moves the discussion from generic chat toward governed assistance: systems that begin with the course's own voice and values, support learning without requiring completion of graded work, and fit within the institutional fabric of higher education.

2. Background and Related Work

Prior work has established the promise of AI in higher education, while also warning that pedagogical value depends on context, governance, integrity safeguards, and institutional readiness [12]–[15]. More recent studies and policy analyses suggest that institutions are moving away from blanket prohibition and toward more explicit guidance on acceptable use, attribution, course-level expectations, and AI literacy [16]–[19]. Within that landscape, course-specific AI TAs have emerged as a particularly useful model because they can be anchored to local materials, shaped by instructor intent, and evaluated against concrete pedagogical goals rather than generic chatbot performance [9], [10], [11]. In this study, we examine three widely used AI platforms by configuring each with the same course materials and integrity policies and then comparing their pedagogically relevant trade-offs through hands-on experiments.

OpenAI's Custom GPT feature provides one of the widely recognized no-code pathways for building a course-specific assistant. Instructors can define persistent instructions, upload course materials as knowledge, and enable additional capabilities such as web search, data analysis, and connected tools. As of today, a GPT can include up to 20 uploaded knowledge files, with each file limited to 512 MB, making the platform attractive for courses that rely on extensive lecture decks, policies, lab guides, and reference materials. A major strength of Custom GPTs is controllability. The platform allows instructors to shape tone, refusal behavior, prompting logic, and task boundaries with relatively fine granularity, which is especially useful in courses where academic integrity expectations must be explicit. Custom GPT also offers broader sharing options than the other no-code platforms examined here, including direct link sharing and managed workspace distribution.

Anthropic's Claude Projects allows instructors to define project-level instructions and upload persistent files so that students remain inside a bounded course context rather than using a general chatbot with no local grounding. Furthermore, Skills now add reusable task behavior and can be shared or provisioned in some plan contexts, which materially strengthens customization beyond simple prompt writing. However, as of today, sharing is more constrained than in ChatGPT. Claude Projects can be shared organization-wide or invitation-based within the same organization, rather than openly distributable through a general-public-style URL. Therefore, if the instructor and students are not



affiliated with the same university, they cannot share Claude Projects for AI TAs. Usage limits are also dynamic and depend on factors such as load, conversation length, model use, and attachments. Even with those constraints, Claude remains a compelling option when instructors want a fast setup, a strong integrity posture, and course-grounded behavior without extensive prompt hardening.

Google's Gemini Gems lets instructors define a reusable assistant with custom instructions and attached materials, and Gemini is particularly effective when students need help interpreting screenshots, diagrams, figures, or mixed-media inputs. Gemini's practical constraints are different from those of ChatGPT and Claude. Google's current documentation emphasizes file and storage upload limits, as well as rolling usage limits, rather than presenting Gems as an unrestricted, persistent knowledge environment. Gemini Apps support file-backed interaction, but large or numerous uploads can exceed context windows or reduce response quality, and the platform explicitly warns that oversized uploads may cause missed details or weaker cross-document connections. Gemini also supports sharing, including viewer and editor roles, but attached files become visible to users with access, which raises governance questions when instructors work with copyrighted, licensed, or sensitive course materials. A further limitation for this use case is that Gems cannot currently be used with Gemini Live, so voice-based interaction does not extend to the custom course Gems. As with the other no-code platforms, built-in instructor-side logs are limited.

At the time of writing, several other popular AI platforms are relevant to higher education, including Grok (xAI), Microsoft Copilot, Mata AI, and others, but none currently offer the same combination of usability, course-specific grounding, and instructor-ready no-code deployment as Custom GPTs, Claude Projects, and Gemini Gems. We therefore exclude them from the comparative analysis in this study.

3. Building Course-Specific AI TAs (Same Knowledge, Different Platforms)

We built three course-specific AI TAs intended to serve the same instructional purpose across different platforms. The goal was not to optimize each platform separately, but to hold the teaching intent constant and evaluate the capabilities and constraints across Custom GPT, Claude Projects, and Gemini Gems. Our priorities were straightforward: make the setup replicable, keep the instructions clear, and avoid design choices that would create uneven access for students.

3.1 Study Setup and Key Constraint

We developed three course-specific AI TAs for the same course using the same course description, custom instructions, syllabus, lecture materials, and related resources. However, there is a critical platform-level constraint in one of the AI platforms. During our study period, the Gemini configuration we used limited the number of uploaded files in persistent use to 10, whereas Claude Projects and Custom GPTs could accommodate the full set of 21 course materials. As a result, the Gemini TA included only the syllabus and core lecture decks, whereas the Claude and Custom GPT versions also included lab manuals, external course resources, relevant research articles, and other academic resources. This directly affected how fully each TA could stay grounded in the course. In our setting, that difference became one of the most important practical limitations of Gemini.

3.2 Pedagogical Goals

The three TAs were built around the same instructional goals. First, they were expected to clarify concepts from lectures and labs. Second, they were expected to support problem-solving without giving away graded answers. Third, they were expected to direct students back to the relevant course materials rather than providing generic advice. Fourth, the TA was intended to function more like a guided study partner than a simple answer provider. That distinction has become increasingly important in recent work on generative AI in higher education, especially in discussions of student agency, tutoring, and ethical use [19], [20].

3.3 Standardized Instruction Block

Across all three platforms, the AI TAs were given the same operating logic. Accuracy was prioritized over speed. If a prompt was unclear, the AI TA was instructed to ask clarifying questions before answering. It was instructed to search the persistent course knowledge first, including the syllabus, lecture slides, lab materials, and related course materials, before producing a response. If that



knowledge was insufficient, it was instructed to request the missing information from external sources and cite them clearly. When students attached PDFs or images, the AI TA was expected to interpret them and explain what they showed. It was also encouraged to ask reflective follow-up questions, recommend credible external resources when appropriate, and defer questions about policy, schedule, and requirements to the syllabus and other course materials.

3.4 Optional SDK / External Tools

This study focuses on no-code classroom feasibility: how quickly an instructor can launch a usable AI TA, how much integrity support appears by default, and how much can be done without development support. Once an API layer is added, results begin to depend on engineering choices such as retrieval design, chunking strategy, tool orchestration, interface design, logging architecture, and caching. They can easily overshadow the out-of-the-box differences between platforms. For that reason, including SDK-based implementations in the present study would have made the findings harder to interpret. It would also have introduced additional data governance and research complexities, including decisions about storage, consent, retention, deletion, and institutional review. Therefore, we consider SDK and external tools as an important future path rather than part of the baseline comparison in this study.

4. Results and Comparative Analysis

Following the development of AI TAs for the same cybersecurity course across the three AI platforms described in the previous sections, we conducted a systematic evaluation and comparative analysis. This analysis examined multiple critical dimensions, including Customization Depth, Ease of Use, Course Resource Integration, Interface and Accessibility, Multimodal Capabilities, Privacy and Ethical Controls, Institutional Cost, and Pedagogical Impact. The results of this comparative evaluation are summarized in Table 1.

Table 1. Comparative Analysis of AI TA Platforms

Criteria	ChatGPT (Custom GPTs)	Claude (Projects + Skills)	Google (Gemini Gems)
Customization Depth	Very High — Persistent instructions, knowledge files, optional actions/apps, and version history.	High — Persistent project instructions/files; Skills extend customization but add governance and security caveats.	Moderate — Instructions, file-backed context, previewing, and sharing, but fewer explicit control layers.
Ease of Setup	Moderate — Streamlined builder, but strong performance still depends on careful prompt design and testing.	High — Simple workflow: create project, upload files, add instructions, and use persistently.	High — Fast setup, simple previewing, and optional AI-assisted instruction rewriting.
Course Resource Integration	Strong — Knowledge files are designed for persistent reference during conversation across common file types.	Strong, with format caveats — Supports multiple file types, but PDFs are best when visual elements matter; non-PDF documents are treated mainly as text.	Moderate — File-backed knowledge works well, but storage, upload, and large-file caveats are more prominent.
Interface & Accessibility	Strong — Private, link-based, workspace, and broader sharing options depending on plan and admin settings.	Moderate — Project sharing is mainly organization-wide or invite-only; Projects do not offer the same general link-sharing model as ChatGPT GPTs.	Strong — Available on the Web, mobile, and Workspace side panel; supports private, organization, link, or public sharing depending on account settings.
Multimodal Capabilities	Strong — Supports documents, spreadsheets, images, code, Web search, image generation, voice chat, and data analysis options.	Moderate — Good document/image support; visual interpretation is strongest in PDFs, weaker in non-PDF documents with embedded images.	Very Strong — Broad support for documents, spreadsheets, images, videos, code repositories, and data visualization.
Privacy & Ethics Controls	Strong — Workspace restrictions, compliance availability, and analytics in managed deployments.	Strong — Clear organizational boundaries for Projects and collaboration controls, but Skills introduce	Moderate — Admin controls and enterprise protections exist, but shared Gems expose instructions and



		prompt-injection and data-exfiltration risks.	uploaded files to authorized users.
Institutional Cost	Variable — Paid access and institutional governance options exist, but institutional licensing and entitlements must be verified locally.	Variable — Plan access is clear, but institutional licensing and entitlements must be verified locally.	Variable — Feature access depends heavily on Workspace and Google AI licensing; Institutional licensing and entitlements must be verified locally.
Pedagogical Impact	High — Best fit for tightly guided course assistants that require strong behavioral guardrails and durable reference to course materials.	High — Particularly strong for writing-intensive, discussion-based, and reflective learning support.	Moderate — Especially useful when multimodal course materials or Google Workspace integration are central.

As of April 2026 (at the time of this writing), our comparative analysis finds that Custom GPT stands out as the most fully featured no-code platform for a deeply customizable, tightly governed, shareable, course-specific AI TA deployment. Claude Project has closed the gap, particularly with the addition of Skills, though it remains more constrained in Project sharing across organizations/institutions and dynamic usage. Gemini Gems offers the strongest multimodal and Google-native integration but remains limited by upload, storage, and context-window constraints.

5. Recommendations: When to Choose Which Platform (Decision Matrix)

Selecting among Custom GPTs, Claude Projects, and Gemini Gems depends less on which platform is “best” in the abstract than on which platform best fits the instructional scenario, course-specific support, governance requirements, and interaction mode of a given course. Accordingly, Table 2 presents scenario-based recommendations drawn from our comparative analysis and current vendor documentation. Rather than designating a universal winner, we align each platform with the use cases for which its documented capabilities are presently the best fit.

Table 2. Decision Matrix for Selecting a Course-Specific AI TA Platform

Instructional Scenario	Primary Choice	Notes
Fast, low-friction no-code launch with strong governance controls	Claude Projects	Claude Projects provide a simple no-code workflow: create a project, upload files, add project instructions, and keep chats scoped to that project. Skills can extend functionality, but they are optional and may require admin enablement or added security review.
Large, persistent course knowledge base (labs, policies, articles, slide decks)	Custom GPTs or Claude Projects	ChatGPT Custom GPTs support persistent knowledge files and no-code configuration. Claude Projects also provide a persistent project knowledge base and document automatic RAG (Retrieval-Augmented Generation) scaling for large project knowledge.
Heavy visual, screenshot, figure, and diagram workflows	Gemini Gems	Gemini is the strongest fit when course interaction depends heavily on multimodal inputs such as images, charts, screenshots, videos, spreadsheets, or code repositories.
Voice conversations with the TA	Custom GPTs	OpenAI explicitly documents voice conversations with GPTs. Claude also supports full voice conversations. Gemini Gems are not currently usable with Gemini Live.
Current-information retrieval, external links, and source-grounded answers	Custom GPTs or Claude Projects	ChatGPT Custom GPTs can enable Web search as a built-in capability. Claude documents live web search with citations and direct URL fetching when web search is enabled.
Auditability, analytics, and institutional oversight	Custom GPTs	OpenAI currently documents the clearest GPT-specific managed-workspace governance stack, including RBAC (Role-based Access Control), sharing controls, action-domain restrictions, compliance visibility, and GPT-level workspace analytics. Claude also now documents meaningful Team/Enterprise analytics and exports, but the GPT-specific governance model is most explicit in OpenAI’s Enterprise/EDU documentation.



Broad student-facing sharing by link or public discovery	Custom GPTs or Gemini Gems	ChatGPT supports private, workspace, link-based, and GPT Store/public sharing depending on plan and admin settings. Gemini supports private, organization, anyone-with-link, or public sharing depending on account type and admin controls. Claude Projects are primarily organization-visible or invite-based rather than generally public project links.
Instructor/TA collaboration inside an organization-scoped workspace	Claude Projects	Claude Projects are especially well suited when instructors and TAs need a shared, persistent workspace with project knowledge, project instructions, and organization-scoped visibility settings. Gemini Gems supports collaboration via the extended Google ecosystem.
Cost and usability under campus provisioning	Locally dependent; verify before recommending	This category remains institution specific. Platform availability, sharing permissions, model access, and governance features depend heavily on local licensing, admin settings, and procurement terms.

Taken together, the scenario matrix summarized in Table 2 suggests that no single AI platform can be regarded as uniformly optimal across instructional scenarios. Rather, each appears to support a distinct area of comparative advantages and is constrained by pedagogical purpose, interaction design, and governance requirements. In contexts where persistent course-specific materials, flexible student-facing availability, voice interaction, and the possibility of later migration to a more formally managed institutional environment are central, the OpenAI ecosystem appears especially well aligned. By contrast, Anthropic’s project-based workflow is particularly well suited to cases in which instructors prioritize rapid no-code setup, organization-scoped collaboration, and relatively bounded instructional use, although these advantages should be considered alongside its more limited sharing model and the additional governance considerations associated with enabling Skills. Google’s offering, in turn, is most compelling where instructional activity depends heavily on multimodal interaction, such as screenshots, figures, diagrams, charts, video, or other visually rich course materials, especially when those workflows are already embedded in Google’s ecosystem. More broadly, the matrix indicates that platform selection is most productively understood not as a simple comparison of model quality, but as a consideration of instructional and institutional fit. Accordingly, the relevant evaluative standard is less about which system performs best in the abstract than which one most appropriately matches the course’s knowledge structure, sharing requirements, interaction mode, and oversight context.

As a result, the matrix suggests three practical conclusions. First, Custom GPT is the strongest choice when the instructional priority is a tightly governed, shareable, knowledge-grounded course assistant that may later require stronger institutional oversight or analytics. Second, Claude Projects is particularly appropriate when instructors prioritize setup speed, rapid no-code deployment, and conservative instructional behavior, provided that its more limited Project-sharing model and dynamic usage constraints are acceptable within the same institution. Third, Gemini Gems is best understood not as the most general platform, but as the strongest option when the course design depends heavily on multimodal interaction, visual explanation, or Google Workspace alignment. Therefore, we emphasize that the selection of the AI platform should be treated as a pedagogical and governance decision based on the course objectives and related learning activities rather than merely a model performance.

6. Conclusions and Future Work

In this study, we compared three no-code AI platforms for course-specific AI TAs, using OpenAI Custom GPTs, Anthropic Claude Projects with Skills, and Google Gemini Gems for the same cybersecurity college-level course, and we evaluated them against the pedagogical and governance criteria. While the findings show that no single platform is uniformly superior across educational scenarios, each AI platform reflects distinct trade-offs and is effective in specific instructional conditions. Overall, as of today, Custom GPT offers the greatest behavioral management, customization details, accessibility, and the clearest path toward institutional oversight, making it the most suitable choice when governance and shareability are primary concerns. Claude Projects may provide the rapid no-code deployment and the strongest collaboration features, while its more restricted sharing model limits cross-institutional teaching, especially when the instructor and students are not within the same institution supported by the Claude service. Gemini Gems is the most capable platform for multimodal interaction, especially within the Google ecosystem, but its upload constraints (e.g., limited course knowledge space) and the exclusion of Gems from voice interaction hinder its usefulness as a general-purpose course-grounded



assistant. More broadly, these findings reinforce a conclusion from our earlier work on course-specific AI TAs [9], [10], [11]: AI platform selection is not merely a matter of model performance. It is a pedagogical and governance decision, and instructors are best served by anchoring it in the actual knowledge structure, interaction patterns, course-specific learning activities (e.g., labs, research projects, discussions, etc.), and oversight requirements of their course.

This study directly leads to several directions in future work. First, the accuracy and reliability of each AI platform's knowledge retrieval (i.e., response to queries) need to be formally evaluated. Future work should benchmark each platform against course-specific ground-truth queries to identify where retrieval quality degrades, where hallucinations occur, and how error rates differ across platforms. In a mission-critical educational context such as medicine and cybersecurity, this evaluation is critical. Second, scalability across AI platforms remains a challenge. As AI platforms update their capabilities, tools, storage limits, underlying models, and pricing structures, the trade-offs summarized in Tables 1 and 2 will shift, and the framework developed here should be re-evaluated at regular intervals to remain current. Third, instructor monitoring and auditability represent persistent gaps across all available no-code platforms. For instance, without access to query-response logs, instructors cannot verify if the AI TA behaved appropriately, identify boundary failures, or detect misuse. Therefore, future work should examine SDK-integrated deployments that enable logging, audit trails, and instructor-facing dashboards, building on the governance considerations we have identified here. Finally, extending the cross-platform comparison to additional disciplines, including medicine, law, arts, and engineering, would test whether the trade-offs observed in a cybersecurity context hold across domains with different precision requirements, content structures, and student interaction patterns, building on the cross-disciplinary AI TA framework.

REFERENCES

- [1] Bowen J., Watson C., "Teaching with AI: A Practical Guide to a New Era of Human Learning," Johns Hopkins University Press, Baltimore, 2024.
- [2] Selwyn N., "Should Robots Replace Teachers? AI and the Future of Education," Polity Press, Cambridge, 2019.
- [3] Park J. S., "AI safety under uncertainty: Hallucinations and unpredictable failures," in *Proceedings of the 25th European Conference on Cyber Warfare and Security (ECCWS)*, Rome, Italy, 2026.
- [4] Bhuiyan S., Park J. S., "Cybersecurity threats and mitigation strategies in AI applications," *Journal of the Colloquium for Information Systems Security Education*, vol. 12, no. 1, p. 7, 2025.
- [5] Simmons T., Park J. S., "Cybersecurity challenges and opportunities with generative AI," in *Proceedings of the 23rd European Conference on Cyber Warfare and Security (ECCWS)*, Jyväskylä, Finland, 2024.
- [6] Burgstahler S., "Universal Design in Higher Education: From Principles to Practice," Harvard Education Press, Cambridge, 2015.
- [7] Ferguson R., Brasher A., Clow D., Cooper A., Hillaire G., Mittelmeier J., Rienties B., Ullmann T., Vuorikari R., "Learning Analytics in Higher Education: A Review of UK and International Practice (Full Report)," Jisc, United Kingdom, 2016.
- [8] Khosravi H., Shum S. B., Chen G., Conati C., Tsai Y.-S., Kay J., Knight S., Gašević D., "Explainable Artificial Intelligence in Education," *Computers and Education: Artificial Intelligence*, Elsevier, vol. 3, art. 100074, 2022.
- [9] Park J. S., "Course-specific AI TA: Implementation and impact on student learning," in *Proceedings of the 7th International Workshop on Artificial Intelligence and Education (WAIE)*, Yokohama, Japan, 2025.
- [10] Park J. S., "Towards AI-augmented teaching for higher education," in *Proceedings of the 6th International Conference on Artificial Intelligence in Education Technology (AIET)*, Munich, Germany, 2025.
- [11] Park J. S., Gorman C., Kelly J., Alfahl Z., "Enhancing medical education through a course-specific AI teaching assistant," in *Proceedings of the 12th International Conference on Higher Education Advance (HEAd'26)*, Valencia, Spain, 2026.
- [12] Zawacki-Richter O., Marín V. I., Bond M., "Systematic review of research on artificial intelligence applications in higher education," *International Journal of Educational Technology in Higher Education*, 2019.



- [13] Crompton H., Burke D., "Artificial intelligence in higher education: The state of the field," *International Journal of Educational Technology in Higher Education*, 2023.
- [14] Crompton H., Burke D., "Artificial intelligence applications in higher education: Theories, ethics, and case studies for universities," Routledge, 2024.
- [15] Chan C. K. Y., "A comprehensive AI policy education framework for university teaching and learning," *International Journal of Educational Technology in Higher Education*, 2023.
- [16] Moorhouse B. L., et al., "Adoption of artificial intelligence in education: Policy and practice," *Computers and Education: Artificial Intelligence*, 2023.
- [17] Luo T., "A critical review of GenAI policies in higher education assessment: Trends, gaps, and future directions," *Assessment & Evaluation in Higher Education*, 2024.
- [18] An P., et al., "Investigating the higher education institutions' guidelines and policies regarding generative AI," *International Journal of Educational Technology in Higher Education*, 2025.
- [19] Hutson J., "From prohibition to preparation: Reframing academic integrity in the age of AI," *MRS Journal of Arts, Humanities and Literature*, 2025.
- [20] Darvishi A., Khosravi H., Sadiq S., Gašević D., and Siemens G., "Impact of AI assistance on student agency," *Computers & Education*, vol. 210, p. 104967, 2023.