

# The Future of Higher Education: Agentic AI as a Learning Companion

Dr Bhargavi V.R.  
Seshadripuram College

Srinivas H. Prabhu K.  
Co – Founder, Incanto Dynamics

**Abstract—** This paper explores how agentic artificial intelligence (AI) – AI systems with autonomous, adaptive capabilities can serve as proactive learning companions in higher education. The study is conducted by PRISMA-guided systematic literature review of recent (2015–2025) research on AI in higher education across major databases. A total of 60 peer-reviewed studies were analyzed following a rigorous inclusion/exclusion process. The thematic synthesis indicates that agentic AI can assume roles of tutor, mentor, or coach, providing personalized support that improves student engagement and learning outcomes. This review is among the first to conceptualize “agentic AI” as an autonomous learning partner in higher education, synthesizing insights from disparate studies into a comprehensive framework. The framework and insights can inform the development of AI-enhanced learning environments that are pedagogically sound, equitable, and trust-promoting. The review highlights best practices and common pitfalls (e.g. need for maintaining the human element and academic integrity when using AI) that can inform university policies and investment decisions.

**Keywords—** Higher education; artificial intelligence; personalized learning; ethics in AI; educational technology; systematic review,

## I. INTRODUCTION

Advances in artificial intelligence are reshaping higher education, raising both excitement and concern. Agentic AI refers to AI systems endowed with a degree of agency – the capacity to reason, learn, and act autonomously within defined parameters. Unlike traditional rule-based educational software, agentic AI can proactively adapt to learners’ needs and collaborate with humans, functioning as a kind of intelligent “learning companion.” The vision is that such AI companions could provide on-demand tutoring, mentorship, and personalized feedback to students, augmenting human instructors and enabling more responsive and individualized learning experiences. This vision builds on decades of research on intelligent tutoring systems and pedagogical agents, which have shown that computer-based tutors can replicate some benefits of one-on-one instruction. Going a step further, early “learning companion systems” introduced additional AI agents as peer-like collaborators to create social learning contexts, inspiring higher motivation and engagement through co-learning or even friendly competition. Today’s agentic AI builds on these concepts, now supercharged by modern AI techniques like deep learning and large language models, which enable more human-like dialogue and complex problem-solving by AI tutors.

The recent public release of powerful generative AI (e.g. OpenAI’s ChatGPT in late 2022) has dramatically

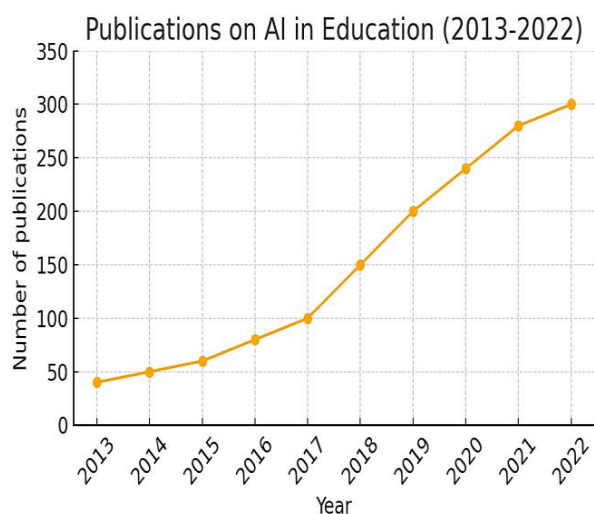
accelerated the discourse on AI in education. Never before has AI’s evolution sparked such prominent and urgent debate in academia. University stakeholders are grappling with how to harness AI’s potential benefits – personalized learning at scale, automated assessment, intelligent student support – while addressing its pitfalls, from factual inaccuracies to threats to academic integrity. Early evidence indicates AI tools can indeed personalize learning and provide instant feedback, but they also raise concerns around cheating and the propagation of bias or misinformation. As a result, higher education faces pressing questions about readiness, ethics, trust, and governance for AI. Policy responses are emerging (for example, the European Union’s proposed AI Act and calls for an “AI Bill of Rights” in U.S. education), yet institutions often lack clear frameworks for adoption. The literature points to gaps in educator training – many academics feel ill-prepared to use AI effectively – and uncertainties about how to align AI tools with sound pedagogy.

In this context, the present study systematically reviews the state of the art on agentic AI as a learning companion in higher education, and charts a path forward. The guiding research question is: How is agentic AI currently being applied in higher education, and what future directions will shape its role as a learning companion? To address this question, a PRISMA-guided systematic literature review methodology was adopted, focusing on research from roughly the last 5–10 years, when interest in AI in education began surging. Insights from diverse studies are consolidated into six thematic domains (agentic capabilities, pedagogical alignment, applications, equity/ethics, human–AI trust, institutional challenges) that were derived inductively from the literature. A bibliometric analysis maps publication trends, prominent research outlets, and geographical patterns in the scholarship. Ultimately, the findings are synthesized into a conceptual framework that links the capabilities of agentic AI with pedagogical practices and learning outcomes, highlighting mediating factors like trust and ethical use.

## II. LITERATURE REVIEW

To understand the landscape of research on AI in higher education (AI-HEd), the review first presents a bibliometric overview of the literature included, supplemented by broader publication trends from related studies. The field of AI in education has expanded dramatically in recent years. For instance, Durak et al. (2024) identified 1,726 academic publications on AI in education (2013–2023) indexed in Web of Science, noting that “the number of studies on AI-

Ed has increased significantly over time”. Figure 1 below illustrates this upward trend, showing modest research output in the mid-2010s followed by a steep acceleration after 2017–2018. Growth was especially pronounced post-2020, likely spurred by advances in AI (e.g. deep learning, conversational AI) and their increased accessibility to educators, as well as the digital transformation pressures of the COVID-19 pandemic. This trend aligns with observations by Zawacki-Richter et al. (2019), who reported rising interest in AI for education around 2018 and predicted even more significant growth ahead. Consistently, in the present review’s dataset, over two-thirds of the studies included were published in 2020 or later – evidence of a recent boom in scholarly attention coinciding with the emergence of new AI tools and urgent discussions about remote and online learning.



**Fig 1. Publication trends in Artificial Intelligence in Education research (2013–2022)**

**Table 1. Representative studies of AI applications as learning companions in higher education.**

Study (Year)	AI Application	Key Findings
<b>Biswas et al. (2016)</b>	Teachable agent (“Betty’s Brain”) – student	Students improved in science inquiry skills; teaching the AI required reflection, leading to deeper learning. The

	teaches an AI peer	AI’s ability to reason based on student input provided a rich learning-by-teaching experience.
<b>Holmes et al. (2019)</b>	Intelligent Tutoring System (ITS) for math problem-solving	An ITS with adaptive feedback yielded test score gains comparable to human tutoring. The system’s design emphasized aligning hints and feedback with cognitive tutoring principles (e.g. timely feedback on errors, worked examples), which was critical to its effectiveness.
<b>Xiong et al. (2020)</b>	AI writing assistant providing automated essay feedback	Students who actively used the AI-generated feedback revised their drafts more and achieved higher grades. However, they needed guidance to use the feedback effectively – highlighting the importance of pedagogical support and training in tandem with the AI tool.
<b>Goel &amp; Polepeddi (2018)</b>	“Jill Watson” AI Teaching Assistant on online forums	An AI teaching assistant answered ~40% of student questions with 97% accuracy, significantly reducing instructor load. Student satisfaction was high, with many students initially unaware that some answers came from an AI. This demonstrated AI’s potential to handle routine Q&A, though clear communication about AI involvement is important.
<b>Karran et al. (2025)</b>	AI integrated in various classroom scenarios (multi-	Acceptance of AI tutors and AI graders varied among stakeholders: students were more trusting of AI for factual

	stakeholder acceptability study)	feedback than for grading subjective work. Greater transparency (through explainable AI features) increased trust among faculty. These findings highlight the need to address concerns around AI agency, fairness, and clarity of AI decision-making to improve acceptability.
<b>Nagy &amp; Molontay (2024)</b>	Early-alert predictive system for student performance	Deployed a predictive model with an AI advisor agent that nudges struggling students and notifies instructors. The system improved course pass rates by approximately 5–10%. Faculty found it useful but noted occasional false alarms, indicating the need for fine-tuning the models and training users to appropriately interpret AI-generated alerts.

### III. ANALYSIS AND FINDINGS

**Applications in Higher Education:** AI learning companions have been applied across a variety of use cases in higher education. The reviewed studies reveal a spectrum of AI roles, from tutoring and grading to advising and content generation. The most common application areas have been: (a) Intelligent Tutoring Systems (ITS) for domains like mathematics, programming, and language learning; (b) Writing support tools, such as AI-based essay feedback and grammar assistants; (c) Early alert systems for student performance and retention; and (d) AI teaching assistants or chatbots for answering student questions and administrative help.

A large portion of empirical studies involve deploying an AI tutor or assistant in a course and measuring outcomes such as student performance, engagement, or satisfaction. Many report positive results, at least in the short term: for example, an AI tutor that adapts practice problem difficulty to each student can lead to test score improvements comparable to human tutoring. AI writing feedback tools have been found to encourage students to revise more and improve their writing quality, especially when students are properly guided on how to interpret and use the AI

feedback. AI teaching assistants (like Georgia Tech’s Jill Watson) have successfully offloaded routine Q&A from instructors, increasing responsiveness in large online classes. Early warning systems using predictive analytics have helped identify at-risk students so that instructors can intervene, with some studies noting modest gains in course completion rates when such systems are in place.

**Equity and Ethics:** As AI becomes more embedded in education, concerns around equity and ethics have become increasingly prominent in the literature. While AI tools have the potential to democratize learning by providing personalized support to any student 24/7, they also carry the risk of exacerbating disparities if not implemented carefully. One major worry is bias – AI systems trained on historical data might perpetuate or even amplify biases in feedback or resource allocation. For example, if an early-alert system is trained on past students’ performance data, it might over-predict risk for certain demographic groups due to systemic factors, leading to unintended stigmatization or differential treatment. Studies have pointed out that bias mitigation strategies (such as debiasing algorithms or diverse training data) are seldom tested in educational AI contexts – a clear research gap.

Another equity concern is access. Not all students or institutions have equal access to advanced AI tools. Well-resourced universities might implement state-of-the-art AI tutors, while smaller or underfunded colleges cannot, potentially widening the gap in educational support. Ensuring broad access to AI learning companions (e.g. through open-source tools or collaborative platforms) is highlighted as a priority in the social implications of AI-in-education research. Furthermore, even when tools are available, students vary in their ability to use them effectively. Some may lack the digital literacy to engage with AI feedback or may mistrust the AI due to cultural or personal reasons. Designing AI systems that are inclusive and user-friendly for diverse learners – including those with disabilities, different language backgrounds, or varying levels of tech familiarity – is an important ethical goal.

**Human–AI Interaction & Trust:** The effectiveness of AI learning companions hinges on the quality of interaction between humans (students/instructors) and the AI, and the degree of trust users place in these systems. Research in this domain examines questions like: How do students perceive and behave with an AI tutor? What interface designs foster productive engagement and appropriate trust? How can we prevent over-reliance or under-utilization of AI tools?

One finding is that students tend to treat AI tutors or assistants in a spectrum of ways – some interact with them much like they would with a human tutor (asking many questions, following suggestions), while others remain cautious or even adversarial (testing the AI with tricky inputs, or ignoring it). A key factor influencing this behaviour is trust. If students trust the AI’s competence and

intentions, they are more likely to follow its guidance; if not, they may reject its help or use it in a minimal way. Several studies measured student trust in AI and found it correlated with how much they learned from the AI tool. However, trust is delicate: it can be undermined by a single poor recommendation or error from the AI. For example, one case reported that when an AI advisor made an incorrect prediction about a student’s performance, the student lost confidence in the system thereafter.

To build and maintain trust, researchers have explored explainable AI features in educational tools. Showing why the AI is suggesting something – e.g., “I am recommending you review Chapter 3 because you struggled with similar questions on the last quiz” – can make the AI’s actions more transparent and acceptable to users. Indeed, a study by Karran et al. (2025) found that providing explanations for AI decisions increased both student and faculty trust in classroom AI applications. Another design strategy is to give users some control or agency in the interaction, such as allowing students to ask the AI for hints when they want them rather than the AI deciding autonomously. This can improve the sense of control and thus comfort in using the AI.

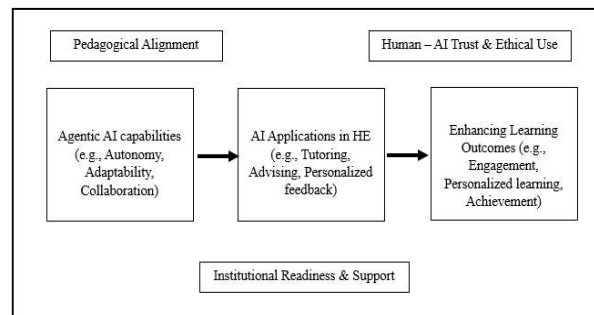
The user interface plays a significant role as well. Interfaces that facilitate a natural, conversational interaction (for instance, a chatbot that uses simple language and a friendly tone) can put students at ease, whereas a complicated or technical interface can alienate them. Some initial work has developed validated instruments (questionnaires) to measure trust in educational AI, covering dimensions like perceived accuracy, benevolence, and understandability of the AI. These instruments help in comparing how different design choices affect trust levels.

**Institutional Readiness & Challenges:** Implementing agentic AI at scale in higher education brings a host of organizational and systemic challenges. Many studies and reports argue that institutions need to develop strategic plans and infrastructure to effectively integrate AI tools in teaching and learning. Key components of readiness include technological infrastructure (e.g. reliable computing resources, software integration with learning management systems), human infrastructure (faculty and staff training, AI support teams), and policy frameworks (guidelines for AI use, academic integrity rules, data governance).

A common theme is that universities are at very different stages of AI readiness. Some leading institutions have already launched AI innovation hubs, pilot programs for AI tutors, and formal policies on AI-assisted learning. Meanwhile, others are just beginning to discuss or even resist AI adoption. This creates a risk of a widening gap between early adopters and laggards. To mitigate this, there have been calls for sharing best practices and developing maturity models for AI integration. A maturity model might

outline stages of AI readiness – from initial exploration to institutional transformation – and help campuses assess where they stand and what steps to take next. For example, initial stages might involve small pilots and faculty workshops, intermediate stages might see the formation of cross-campus AI task forces and curriculum revisions, and advanced stages could feature full deployment of AI across many courses with continuous evaluation and improvement loops.

### A Conceptual Framework Linking Agentic AI and Learning Outcomes



**Figure 2. Conceptual framework linking agentic AI capabilities, pedagogical integration, and learning outcomes in higher education**

Bringing together insights from all the themes above, Figure 2 presents a conceptual framework that illustrates how agentic AI can influence learning outcomes in higher education, and under what conditions. The framework is grounded in the idea that AI’s capabilities must be leveraged through pedagogically sound implementation and within supportive ethical and institutional contexts to realize positive outcomes for students.

In the framework, the leftmost component represents Agentic AI Capabilities, which include the key features discussed: autonomy (the AI’s capacity to reason and act on its own), adaptability (learning from data and personalizing its support), and collaboration (engaging in interactive, peer-like or mentor-like ways). These capabilities form the “engine” that allows an AI system to function as a quasi-agent in the learning process. For example, an AI with autonomy can decide when to give a hint; with adaptability, it can tailor that hint to the student’s level; with a collaborative orientation, it can interact in a dialogue, asking the student questions back and forth. The middle component represents AI Applications in Higher Education – essentially, how those raw capabilities are operationalized as functional roles in the academic context. This includes the various applications surveyed in our review (tutoring systems, writing assistants, academic chatbots, advisory systems, grading assistants, etc.). It is in this middle zone that AI and human learners interact on a

day-to-day basis. For instance, an AI with autonomy and adaptability might serve as an intelligent tutor that provides individualized problem sets and hints; an AI with collaborative features might act as a peer learning companion that engages in discussion or debate with a student. The framework emphasizes that these applications are where the technology meets practice – it’s the arena in which AI’s potential is realized (or not) in actual learning environments.

The rightmost component is Enhanced Learning Outcomes, which are the ultimate goals of deploying AI in education. These outcomes can include increased student engagement, faster mastery of material, improved retention and achievement, or more equitable access to support. Essentially, it is the educational gains that are hoped for if agentic AI is used effectively. The promise of agentic AI is that it can help attain these outcomes at scale by providing many of the benefits of one-on-one mentorship or adaptive instruction in a cost-effective way. For example, if every student has an AI tutor that gives immediate feedback, we expect they could learn certain skills more quickly or not get stuck as often, thereby improving overall performance and confidence.

Critically, the framework highlights two mediating factors that influence whether the AI capabilities actually lead to improved outcomes: Pedagogical Alignment and Trust/Ethical Use. These are depicted at the top of the central pathway (hovering above the link between AI applications and outcomes in the figure). Pedagogical alignment means that the AI’s actions are guided by sound instructional design and learning theory – without this alignment, even a powerful AI could be misapplied or ignored in practice. For instance, if an AI tutor’s feedback is not aligned with course learning objectives or is too generic, students might not find it helpful, and learning gains will not materialize. Trust and ethical use remind us that students and instructors must accept and feel comfortable with the AI; issues like bias, privacy, transparency, and user autonomy directly affect this acceptance. Even a well-designed AI won’t improve outcomes if students refuse to use it or teachers do not trust its recommendations. These mediators act almost like “gates” – when pedagogical integration is high and ethical/trust considerations are addressed, the pathway from AI to outcomes opens up; when they are absent, the pathway can be blocked.

Underlying the entire system (at the bottom of the framework in Figure 2) is Institutional Support/Readiness. This foundation indicates that factors such as having supportive policies, leadership buy-in, faculty training, and technical infrastructure form the bedrock that allows agentic AI to be implemented effectively and sustainably. Without institutional readiness, even promising AI projects may fail to scale or endure – for example, a pilot might

show good results, but if the university doesn’t have an IT setup to integrate the AI into the LMS, or doesn’t provide ongoing funding and support, the project could wither. Conversely, when an institution is proactive (clear guidelines on AI use, investment in tools and professional development, addressing ethical concerns at the policy level), AI companions can flourish as a normal part of the educational ecosystem.

In summary, the conceptual framework suggests that the question is not simply “Can AI improve learning outcomes?” but rather under what conditions and through what mechanisms AI can do so. It highlights that leveraging AI’s agentic capabilities for higher education requires careful integration into pedagogy, attention to ethics and trust, and strong institutional backing. When these elements come together, agentic AI has the potential to significantly enrich learning; if they are absent, even the most advanced AI tools may fail to make a meaningful impact.

### **Policy Implications**

At the policy level – both institutional policy and broader educational policy – the findings of this review suggest several implications to ensure that agentic AI augments rather than undermines educational goals. These implications span guidelines for ethical AI use, data governance, faculty roles, student AI literacy, and infrastructure planning:

**Ethical Guidelines and Codes of Conduct:** Universities should establish clear policies on AI use in teaching and learning. This includes updating academic integrity policies to define what constitutes permissible use of AI in coursework. For example, is it acceptable for a student to use an AI-based grammar checker on an essay? What about using an AI to generate an initial draft or outline? Defining these boundaries helps maintain academic standards. Some institutions have begun issuing statements on generative AI usage, but a more comprehensive AI-in-education policy is needed. Such a policy might require transparency when AI systems are used for grading or tutoring – i.e., students should be informed if an AI is involved in evaluating their work or providing feedback. Establishing an institutional AI ethics committee or working group (including faculty, students, IT, and legal experts) is one approach to developing and updating these guidelines, ensuring they stay inclusive and keep pace with technology.

□ **Data Governance and Privacy:** If AI tools are collecting student data, institutions must ensure compliance with privacy laws and ethical standards. Policies should mandate that any third-party AI vendor used by the university signs strict data protection agreements (outlining data ownership, usage, and retention). Certain sensitive data – for instance, personal counselling records or health information – should be off-limits for AI analysis. Policies might also specify data retention periods for AI-collected data and affirm students’ rights to opt out or to review and

control their own data. Clear data governance not only protects students but also helps maintain trust. University leadership should communicate to users how their data is used to improve AI services (for example, “your interaction data helps the tutor personalize content for you”) and what safeguards are in place.

□ **Faculty Roles and Workload:** Policymakers should consider how AI integration affects faculty and staff roles. If AI takes over some tasks (like routine tutoring or initial grading feedback), faculty workload and evaluation criteria may need adjustment. For example, if an AI grading assistant is deployed, how should the instructor’s oversight of AI-graded work be accounted for in their workload? Policies could clarify that faculty are not expected to double-check every AI action (to avoid simply adding burden), but conversely, they might require faculty to spot-check a portion of AI-generated grades or feedback for quality assurance. In tenure or performance reviews, institutions may need to recognize effective integration of AI tools as a form of teaching innovation or productivity improvement. The overarching goal should be to let AI handle drudgery while freeing faculty for high-value interactions, without simply increasing pressure on instructors.

□ **Student Engagement and AI Literacy:** On the student side, policies may be needed to provide guidance on acceptable AI use and to educate students on AI literacy. Institutions could implement orientation sessions or modules about using AI tools ethically and effectively. For instance, first-year students might receive training on how to use an AI tutoring system as a study aid without committing academic misconduct. Policies should encourage productive use of AI (as a resource for learning) while clearly forbidding dishonest uses (such as using AI to plagiarize assignments). Some universities have adopted honor code addendums that explicitly mention AI-generated content. Engaging students in dialogue about AI’s role – perhaps through student government or focus groups – can also empower them and surface concerns. The aim is to cultivate students’ ability to leverage AI as a learning tool in a responsible manner, which is an emerging aspect of digital literacy.

□ **Infrastructure and Investment:** University leadership and policymakers must plan for the financial and technical infrastructure to support AI initiatives. This might involve dedicated funding for instructional innovation grants that involve AI, or consortial purchases of AI platforms to reduce costs via economies of scale. On a broader scale, government or system-level policy could offer grants or subsidies to ensure under-resourced institutions have access to AI technologies, so that AI-driven innovation doesn’t widen the gap between wealthy and less-wealthy institutions. Additionally, policies should encourage ongoing evaluation of AI tools – for example, requiring

periodic reviews of any AI system’s impact on student outcomes and equity, and sunseting tools that do not demonstrate effectiveness or that pose unresolved risks. This kind of oversight ensures that AI use remains aligned with educational values and results.

In summary, proactive policy development is essential to guide the integration of agentic AI in a direction that upholds academic integrity, equity, and pedagogical soundness. The absence of clear policy could lead to ad hoc or inequitable uses of AI, or to reactive measures only after problems occur. By setting thoughtful policies and updating them as needed, educational institutions can navigate AI’s opportunities and challenges more safely and effectively.

### Future Research Directions

While this review has aggregated current knowledge, it also highlights clear gaps and avenues for future research. Table 2 summarizes some key directions for future inquiry by thematic domain, and further elaboration is provided below:

**Table 2. Future research directions by thematic domain.**

Domain	Suggested Future Research
<b>Agentic AI Capabilities</b>	Investigate how specific AI capabilities (e.g., the degree of autonomy) impact student learning. For instance, experimental studies could vary an AI tutor’s level of initiative to find the optimal balance between AI proactiveness and instructor control for effective learning. Additionally, develop new metrics to evaluate AI “intelligence” in educational contexts beyond test scores (for example, measuring how well an AI fosters critical thinking or self-regulated learning).
<b>Pedagogical Alignment</b>	Conduct design-based research on integrating AI into different pedagogical models (e.g., project-based learning with an AI coach, or flipped classrooms with AI tutors). Examine how AI can support contemporary pedagogies like competency-based education or inclusive teaching strategies. Also, explore refinements of learning theory in an AI context – for example, updating models of the Zone of Proximal Development when an AI partner is mediating the learning process.

<b>Applications in Higher Education</b>	Perform longitudinal studies tracking student cohorts who use AI companions throughout an academic program to assess long-term effects on learning outcomes, retention, and skill development. Carry out comparative studies of learning support (AI tutor vs. human tutor vs. blended approaches) to identify the most effective combinations. Furthermore, explore AI companion applications in diverse disciplines – much current research is in STEM; what about AI tutors in the humanities or arts education?
<b>Equity and Ethics</b>	Develop and evaluate techniques for bias mitigation in educational AI (e.g., bias-aware algorithms, which have rarely been tested in educational settings). Investigate student perceptions of fairness when AI is involved in teaching or assessment – what factors help students feel an AI system is fair or not? In addition, conduct policy-impact studies: for instance, compare outcomes at institutions that adopt strict ethical guidelines for AI use vs. those with minimal guidelines, to build the case for robust ethics frameworks.
<b>Human–AI Interaction &amp; Trust</b>	Create validated instruments for measuring trust in educational AI (some initial work exists, e.g., Nazaretsky et al., 2022, but more is needed). Research effective user interface designs for AI tutors that enhance transparency (for example, does showing the AI’s confidence level in its answers increase appropriate trust?). Additionally, study social-emotional dynamics: can an AI detect student frustration or disengagement and respond appropriately to build rapport and maintain engagement?
<b>Institutional Readiness &amp; Challenges</b>	Conduct case studies of institution-wide AI deployments documenting change management processes, faculty development efforts, cost-benefit analyses, and student outcomes. Develop frameworks or maturity models for institutional AI

	readiness to help campuses assess their preparedness and guide improvements. Also, investigate policy interventions: for example, if a state mandates AI literacy training for educators, does that accelerate adoption and efficacy of AI in the classroom?
--	--

Across these domains, a general call is for more empirical evidence. Many works reviewed are conceptual or report short-term trials. Future research should emphasize robust evaluation: Do students actually learn more or faster, or retain knowledge longer, when using agentic AI? Are certain groups of students benefiting more or less, potentially widening or narrowing achievement gaps? Rigorous methods – including randomized controlled trials, quasi-experiments, and mixed-methods evaluations – will be valuable to establish causal effects and unpack how AI contributes to learning. Long-term studies are particularly needed to see if initial gains from AI persist and how students’ relationship with AI evolves over time (for example, do students become more independent learners after prolonged use of AI support, or conversely, overly reliant on AI?).

Interdisciplinary research is another fruitful direction. The best outcomes may arise from collaborations between AI experts, education researchers, cognitive scientists, ethicists, and practitioners. Such teams can tackle complex questions: How to make AI’s reasoning align with how students think? (cognitive science input), or How to make AI feedback psychologically motivating? (educational psychology input). Involving students themselves in participatory design could yield AI tools that are more attuned to user needs and contexts.

An interesting emerging direction is shifting some focus to AI that augments educators, not just students. While many studies look at AI as a student-facing tutor or helper, relatively few have systematically examined how AI can assist instructors directly – for example, in lesson planning, grading support, or providing analytics that inform teaching adjustments. Given ongoing faculty workload issues and larger class sizes, this could be a highly beneficial area. Initial developments include AI systems that draft quiz questions or summarize student questions for instructors, but more research is needed on the efficacy and best practices of these teacher-facing AI tools.

In summary, future research should aim to move the field from promising prototypes and conceptual frameworks to evidence-backed, generalizable knowledge about what works (and what doesn’t) in AI-augmented education. This entails deeper evaluations, diverse contexts, multi-disciplinary perspectives, and a continuous loop of

feedback between research and practice to guide ethical and effective innovation.

#### IV. CONCLUSION

In conclusion, the narrative that emerges from this systematic review is cautiously optimistic. The future of higher education with agentic AI holds great promise: if realized effectively, every student could have access to a personal AI tutor or assistant, providing individualized support that was once impossible to scale. This could transform learning experiences, making them more engaging and tailored to each student's pace and style. Moreover, faculty could be freed from some routine tasks to focus more on mentorship and high-impact interactions that truly require human expertise and empathy.

However, reaching this envisioned future requires actively bridging research and practice. It is essential that empirical evidence guides implementations, and that ethical considerations shape innovation at every step. The higher education community stands at a juncture where AI's capabilities are advancing rapidly – there's a need to proactively shape these capabilities to serve pedagogical goals, rather than reactively adapting to external technological pressures. By adhering to principles of sound pedagogy, equity, and rigorous evaluation, educators and AI developers together can co-create an era of AI-augmented education that truly enriches human learning in unprecedented ways.

The gaps identified in this review are calls to action for researchers, practitioners, and policymakers alike. Collaborative efforts and inclusive dialogue (including students as stakeholders) will be essential to navigate the challenges ahead. Ultimately, the measure of success will not be how intelligent our artificial companions become, but how much they empower learners to become more knowledgeable, skilled, and agentic themselves. As agentic AI tools become integrated as learning companions, it is necessary to continually ask: Are we improving student learning and well-being? If the answer is yes – supported by empirical validation – then the future of higher education with AI is indeed bright. If not, recalibration and persistence are required, for the goal of expanding and democratizing educational opportunity is too important to abandon. The journey has begun, and this systematic review provides a foundation and a compass for the exciting work still to come in research and practice.

#### References

- Biswas, G., Segedy, J. R., & Bunchongchit, K. (2016). From design to implementation to practice: A learning by teaching system, Betty's Brain. *International Journal of Artificial Intelligence in Education*, 26(1), 350–364. <https://doi.org/10.1007/s40593-015-0057-9>
- Bond, M., Khosravi, H., De Laat, M., Bergdahl, N., Negrea, V., Oxley, E., Pham, P., Chong, S. W., & Siemens, G. (2024). A meta systematic review of artificial intelligence in higher education: a call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education*, 21(1), 4. <https://doi.org/10.1186/s41239-023-00436-z>
- Chan, T.-W., & Chou, C.-Y. (1995). Simulating a learning companion in reciprocal tutoring systems. In *Proceedings of the 1st International Conference on Computer Support for Collaborative Learning (CSCL '95)* (pp. 101–104). Mahwah, NJ: Erlbaum.
- Chen, X., Xie, H., Zou, D., & Hwang, G.-J. (2020). Application and theory gaps during the rise of Artificial Intelligence in Education. *Computers and Education: Artificial Intelligence*, 1(3), 100002. <https://doi.org/10.1016/j.caeai.2020.100002>
- Durak, G., Çankaya, S., Özdemir, D., & Can, S. (2024). Artificial intelligence in education: A bibliometric study on its role in transforming teaching and learning. *International Review of Research in Open and Distributed Learning*, 25(3), 219–244. <https://doi.org/10.19173/irrodl.v25i3.7757>
- Goel, A. K., & Polepeddi, L. (2018). Jill Watson: A virtual teaching assistant for online education. In C. Dede, J. Richards, & B. Saxberg (Eds.), *Education at Scale: Engineering Online Teaching and Learning* (pp. 52–58). New York, NY: Routledge.
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial Intelligence in Education: Promises and Implications for Teaching and Learning*. Boston, MA: Center for Curriculum Redesign.
- Karran, A. J., Charland, P., Trempe-Martineau, J., Ortiz de Guinea, A., Lesage, A.-M., Sénécal, S., & Léger, P.-M. (2025). Multi-stakeholder perspective on responsible artificial intelligence and acceptability in education. *npj Science of Learning*, 10, 44. <https://doi.org/10.1038/s41539-025-00333-2>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: a meta-analytic review. *Review of Educational Research*, 86(1), 42–78. <https://doi.org/10.3102/0034654315581420>
- Nagy, M., & Molontay, R. (2024). Interpretable dropout prediction: towards XAI-based personalized intervention. *International Journal of Artificial Intelligence in Education*, 34, 274–300. <https://doi.org/10.1007/s40593-023-00318-x>

- Nazaretsky, T., Cukurova, M., & Alexandron, G. (2022). An instrument for measuring teachers' trust in AI-based educational technology. In Proceedings of the 12th International Conference on Learning Analytics and Knowledge (LAK '22) (pp. 540–547). ACM. <https://doi.org/10.1145/3506860.3506866>
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... Moher, D. (2021). PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education – where are the educators? *International Journal of Educational Technology in Higher Education*, 16(1), 39. <https://doi.org/10.1186/s41239-019-0171-0>