Australian Journal of
Business and Social Science
An Open Access Journal

A Critical Review of Student Self-Learning from AI Platforms: Opportunities, Limitations, and Pedagogical Risks

W. C. Tang
BEng (Hons), MEd, PCEd, MSc, PhD, FIMA, FCMI, FIMC
Panel head of Mathematics, Lock Tao Secondary School, Hong Kong
twc@locktao.edu.hk

Abstract

The rise of AI-powered learning platforms, such as ChatGPT and Khanmigo, has changed the way students learn on their own. This critical review integrates evidence from educational technology research, cognitive science, and critical pedagogy to assess the assertions and realities of AI-assisted self-learning. The review looks at four areas: (1) the technical features of current AI learning platforms, (2) real-world evidence of learning outcomes, (3) the effects on cognitive and metacognitive processes for self-regulated learning, and (4) larger issues of pedagogy and fairness. The analysis uncovers a substantial disparity between promotional rhetoric and substantiated evidence. AI platforms can offer tailored practice, instant feedback, and access to extensive knowledge repositories; however, they may also promote passive learning, weaken metacognitive skills, reinforce biases, and encourage "answer-seeking" behaviours instead of true comprehension. The review finds that AI platforms are best seen as tools that work with, not instead of, structured teaching and human guidance. It gives important advice to students, teachers, and people who make platforms.

Keywords: AI in education, ChatGPT, critical pedagogy, educational technology, intelligent tutoring systems, metacognition, self-directed learning

1. Introduction

1.1 The Rise of AI in Self-Directed Learning

A student is sitting alone at a desk, having trouble with calculus. She doesn't want to wait for a teacher or look through a textbook, so she opens an AI tutoring platform. The system quickly creates a step-by-step explanation, adjusts to her speed, and gives her practice problems based on the mistakes she made. This situation is no longer in the future. Since ChatGPT was made available to the public in late 2022, the use of generative AI for self-learning has skyrocketed. A survey by the Higher Education Policy Institute (2025) found that 88% of students had used generative AI to help them with their homework. A large number of them also said they used AI as their main study aid outside of class.

There are many different AI platforms that are marketed for self-learning (Zhou, 2020). Some examples are general-purpose chatbots like ChatGPT, Claude, and Gemini; adaptive tutoring systems like Khan Academy's Khanmigo, Carnegie Learning, and Duolingo; content summarisation tools like Elicit and Scholarcy; and specialised subject tutors like Wolfram Alpha for math and Codecademy's AI assistant for programming. They all promise to make personalised, immediate, and infinitely patient instruction available to everyone. This is something that traditional classrooms, with their set pace and limited teacher attention, cannot do.

1.2 The Self-Learning Context

Self-directed learning, where the learner starts and manages their own learning without direct teacher control, is not a new idea. Students who are motivated have always used libraries, videos, and other people to help them learn in addition to what they learn in school. AI platforms, on the other hand, change things in a way that is different from other platforms. AI agents, on the other hand, interact, respond, and change in real time. They can act like Socratic dialogue (Akhmedov & Yuldasheva, 2023), make an endless number of practice problems, and give feedback right away. People often call this kind of interactivity revolutionary because it promises to change education from being focused on the teacher to being focused on the student.

But not all revolutions lead to progress. The same things that make AI useful (being responsive, fluent, and seeming like it knows what it's talking about) can also be dangerous. Students might think that being able to speak fluently means they know a lot. They might let the AI do the thinking for them, which cuts short the hard work that is necessary for deep learning. They might become dependent on others, which would hurt their independence instead of helping it. This review takes a close look at these tensions.

1.3 Scope and Methodology

This review combines peer-reviewed empirical studies, systematic reviews, theoretical critiques, and grey literature (such as industry reports and platform documentation) that were published between 2018 and 2026. The emphasis is on K-12 and higher education students participating in voluntary, self-directed learning through AI platforms, excluding teacher-directed AI utilisation in classrooms. There are five main parts to the review: Section 2 looks at the technical strengths and weaknesses of the platforms that are available right now. Part 3 looks at real-world proof of learning outcomes. Section 4 looks into the cognitive and metacognitive aspects, which are at the core of the self-learning critique. Section 5 talks about issues related to teaching and fairness. Section 6 gives important suggestions and a summary.

The main point is that AI platforms can help people learn on their own but using them without thinking about it

could hurt the skills (metacognition, persistence, and critical evaluation) that self-directed learning tries to build. A critical perspective is not opposed to technology but rather supportive of pedagogy: the inquiry is not whether AI can facilitate learning, but under what circumstances, with what direction, and for which learners.

2. Affordances and Technical Limitations of AI Learning Platforms

2.1 What Current Platforms Offer

Chatbots that use generative AI (ChatGPT, Claude, Gemini) -- These large language models (LLMs) can explain ideas in different levels of difficulty, make analogies, answer follow-up questions, and come up with practice questions. Their conversational interface makes it easier for students to get help because they can ask simple questions without worrying about being judged. They are great at covering a wide range of topics and at being flexible with language, changing how they explain things to fit different reading levels and learning styles.

Khanmigo and Carnegie Learning are examples of adaptive tutoring systems -- These platforms are made for teaching, unlike regular chatbots. They use ideas from learning science, like scaffolding, worked examples, mastery learning, and spaced repetition. For example, Khanmigo uses Socratic questioning to help students find answers instead of giving them directly. But they don't cover as many topics, and they aren't as good at having conversations as general LLMs.

Tools that are specific to a field (Wolfram Alpha, Codecademy AI, Duolingo Max) -- These combine AI interfaces with reasoning that is specific to a domain, like programming execution, symbolic mathematics, and speech recognition. Wolfram Alpha can solve problems and show how to do them step by step. With Duolingo Max, you can have conversations with AI characters in role play. These tools work better in their specific areas, but they are less flexible.

2.2 The Hallucination Problem

The most well-known problem with generative AI is hallucination, which means making information that seems real but is not. A 2025 study that looked at how accurate ChatGPT was in different academic subjects found that the error rates were between 3% (for simple mathematics) and more than 30% (for medical diagnosis and historical dates). Hallucinations are especially dangerous for students (Danyaro et al., 2025) who are learning on their own because they do not have the background knowledge to spot mistakes. A student studying the French Revolution might believe a made-up date is true, and a student learning to code might copy a code snippet that does not work.

Platforms have put in place ways to reduce the risk of misinformation, such as retrieval augmented generation (RAG), which bases answers on verified sources, confidence scoring, and clear disclaimers. But no current system gets rid of hallucinations. The self-learning student inadvertently assumes the role of a fact checker, for which they are inherently unprepared.

2.3 The Reasoning Ceiling

A more fundamental limitation is that LLMs do not reason in the way that humans do. They use statistical patterns in the training data to guess what the next token will be. When a physics problem is common, the AI finds the right answer by filling in the gaps in the pattern. Performance drops when the problem is new or needs real inference. A study in 2024 showed that changing the numbers in a mathematics problem (for example, from "train X travels at 60 km/h" to "63 km/h") made some models fail. This means that they had memorised problem templates instead of understanding the rules.

This makes it hard to learn on your own. Students who use AI to solve problems might get the right answers to common questions, but not to new ones. If the AI cannot reliably move knowledge from one problem to another, it can't teach skills that can be used in many situations. The student learns how to solve that specific problem, not problems like that.

2.4 Personalisation: Real or Simulated?

AI platforms say they can make learning more personal by changing based on how well each person does. Adaptive tutoring systems do this in a reliable way: they keep track of mastery, find gaps, and change the difficulty of problems. However, general chatbots use conversation memory to make it seem like they are personalising things, but they do not have strong student models. If a student says, "I don't understand", they may get a different explanation, but the system does not keep track of their knowledge, misunderstandings, or learning path over time.

To truly personalise learning, you need to do a diagnostic assessment to find out what the student doesn't understand, not just that they are confused. Current AI systems have a hard time making precise diagnoses. They might be able to tell that a student got the answer wrong, but they might not be able to tell if the mistake was due to a lack of understanding, a mistake in the process, or a misunderstanding of the question. Personalisation is only skin deep without a correct diagnosis.

2.5 Feedback Quality

A well-known benefit of AI is that it gives you feedback right away. But the quality of the feedback can be very different. AI feedback can be accurate and helpful for tasks that are well-structured, like multiple choice questions or short answers with only one right answer. For poorly structured tasks like essays, open-ended projects, and creative work, AI feedback usually only looks at surface features like grammar, structure, and keyword presence. It does not look at substance, the quality of the argument, or originality.

Also, AI feedback is often too positive. Models that are trained on data about what people like learn to avoid criticism because users like positive responses more. A student may be praised for a subpar essay, receiving no encouragement to enhance it. One critical pedagogue said, "The AI tutor that never says, 'this is wrong' is not a tutor but an enabler."

3. Empirical Evidence on Learning Outcomes

3.1 The State of the Evidence

Even though a lot of people use AI to help them learn on their own, there is not much strong evidence that it works. A systematic review conducted in 2025 of 47 experimental and quasi-experimental studies on LLM-based tutoring (Ramos & Casas, 2025) revealed that only 12 employed randomised controlled designs. Most studies evaluated short-term gains (hours or days) on limited, AI-generated assessments, resulting in a circular validation issue. Long-term retention (weeks or months) and the transfer to novel problems were infrequently investigated.

The review indicated a median effect size of $d = 0.43$ for AI tutoring relative to no tutoring, yet only $d = 0.11$ when contrasted with active controls (e.g., textbook study, video lectures). This means that AI is better than doing nothing, but not much better than other self-learning tools.

3.2 Mixed Findings Across Domains

The best results come from math and programming. A study involving 500 high school students utilising an AI tutor for algebra (Henkel et al., 2024) revealed that participants who engaged with the tutor for 30 minutes each day exceeded control group performance by 0.3 standard deviations on standardised assessments. But the effect went away when students were given problems that required more than one step of reasoning that was not in the tutor's practice set. AI code assistants like GitHub Copilot help beginners write working code faster, but research shows that they also make it harder to understand and fix code. Students depended on AI-generated solutions that they were unable to elucidate or rectify.

Learning a new language is a mixed bag. Duolingo's AI features (like explanations and role play) make people more interested and boost their self-reported confidence, but objective tests of proficiency do not show any improvement over the non-AI version. AI conversation partners can help you practise your fluency, but if the AI doesn't always correct your mistakes, it might make them worse. AI feedback tools like Grammarly and ProWritingAid help with surface correctness but not with making big changes to the writing. Stadnik et al. (2025) found that students who used AI feedback made more changes to their writing but fewer big changes to its structure.

The evidence for the humanities and social sciences is the weakest. Students who used AI to summarise readings or make essay outlines did worse on comprehension tests than students who read and outlined by hand. The AI group showed "cognitive offloading" by putting less effort into encoding information because they thought the AI would fill in the gaps. When students were asked open-ended questions that required them to put together information from different sources, AI helped them write more fluently, but their answers were less accurate and less original.

3.3 The Expertise Reversal Effect

A significant finding from cognitive load theory is the expertise reversal effect (Sweller, 2024), which posits that instructional methods beneficial for novices may impede experts, and conversely. This means that the same AI features (like step-by-step solutions) help students who are having trouble but hurt students who are already good by taking away their ability to work hard. Most AI platforms provide uniform support, neglecting to tailor assistance to the learner's proficiency. Instead of letting a student figure out the answer on their own, they are shown the answer, which takes away their chance to solidify their understanding.

3.4 Publication Bias and the Halo Effect

There is a publication bias in the literature on AI in education: positive results are more common than null or negative results. The AI halo effect, which is the tendency to think that AI-generated outputs are more credible than they really are, affects both researchers and students. Research comparing "AI tutor" to "control" frequently neglects to account for novelty effects, engagement effects, or placebo effects. A student using an AI platform may put in more effort just because it's new and exciting, not because AI is better at teaching. Not many studies have long-term follow-ups, and even fewer consider the fact that novelty wears off over time.

4. Cognitive and Metacognitive Dimensions

4.1 Metacognition and Self-Regulation

Self-directed learning (Bhandari, Chopra & Singh, 2020) is heavily based on metacognition, which is the ability to keep track of your own understanding, plan how to learn, check your progress, and make changes as needed. Metacognition is what let’s a student say, "I don't really get this", even though the AI gives a good answer.

AI platforms represent a significant threat to metacognitive advancement. When an AI answers questions right away, the student doesn't have to think about how sure they are. The student doesn't have to choose the right challenges when an AI automatically changes the difficulty level. The student doesn't have to figure out what they don't know when an AI makes practice problems for them. Each of these features, which are sold as "convenience", systematically takes away chances to practise metacognition.

The results can be measured. A longitudinal study conducted in 2024 involving 300 university students revealed that participants utilising AI tutors (Pillai et al., 2024) for one semester exhibited considerable reductions in self-reported metacognitive awareness and in performance on metacognitive monitoring tasks, such as predicting their own test scores prior to examination. The decrease was contingent upon dosage: increased AI utilisation forecast enhanced metacognitive deterioration. This discovery is concerning as metacognition serves as a more significant predictor of academic success than mere intelligence or prior knowledge.

4.2 The Generation Effect and Productive Struggle

Cognitive psychology has consistently evidenced the generation effect: information that learners actively produce (as opposed to merely receiving) is retained more effectively and transferred more efficiently. Struggling with a problem that is just out of reach of one's current skills is an important part of deep learning. AI platforms are made to make things easier. They come up with answers,

explanations, and fixes. The student stops making their own understanding and starts using what others have made.

Think about two students who are learning how to solve quadratic equations. Student A tries for twenty minutes, fails, looks over the textbook, and then finally figures it out. ChatGPT gives Student B a step-by-step answer, which they then copy. Tomorrow, both can answer a similar question, but only Student A has learned how to be flexible and resilient when it comes to solving problems. Student B has learned how to do what they are told. When the problem changes a little, Student B fails, but Student A changes to fit it. The AI has not taught mathematics; it is taught people to depend on it.

4.3 Illusions of Competence

AI platforms make it look like they are very good at what they do. When a student watches an AI explain a concept, they think they understand it. This is called the fluency illusion. The explanation is clear, the logic seems to work, and the student feels like they know what they're doing. But just because you understand an explanation does not mean you can make one yourself. The student is unable to elucidate the concept without the AI, apply it to novel contexts, or identify flaws in the AI's reasoning.

The AI's interactivity makes this illusion even stronger. Students think that when they ask a follow-up question and get an answer, they are having a conversation and understanding each other. But the AI does not get it; it just patterns. The student thinks that the AI's ability to speak well means that they understand it. Studies on human-AI interaction refer to this phenomenon as over trust: the consistent inclination to ascribe greater intelligence and dependability to AI systems than their actual capabilities justify.

4.4 Attention and Deep Processing

To learn on your own, you need to pay attention for a long time and think deeply about what you learn. This means rehearsing what you have learned, connecting new information to what you already know, and thinking critically about it. AI platforms might make people process information quickly and in a shallow way because they respond quickly and have a conversational style. When students talk to AI, it's like when they talk to people on social media: they have quick conversations, get what they want right away, and switch tasks all the time. When a student reads a chapter from a textbook for thirty minutes, they are doing deep, uninterrupted processing. When a student asks the AI a series of short questions, they are doing shallow, broken processing.

An eye-tracking study in 2025 looked at how students learned from an AI chatbot and a textbook (Becerra, 2025). People who used AI spent more time looking at the answer area and less time looking back at the question. They looked at key terms less often and didn't show as much evidence of rereading or cross-referencing. Their gaze patterns showed that they were passively taking in information instead of actively making sense of it. When tested, they did worse on questions that required them to connect ideas (inferential questions) but did the same on questions that required them to remember facts. The AI helped with memorisation but not with understanding.

4.5 Curiosity and the Path of Least Resistance

Curiosity (the inherent drive to investigate and clarify ambiguity) fuels self-directed learning. AI platforms can make people curious. The good thing is that they let students go off on tangents and ask "what if" questions without feeling anxious around other people. A student who is too shy to raise their hand in class can ask the AI as many questions as they want.

On the downside, AI answers make it so cheap to ask questions that they take away the motivational power of not knowing. When there is a gap between what someone knows and what they want to know, curiosity grows. This is an information gap that makes them look forward to learning more. The excitement goes away when the AI closes that gap quickly and easily. The student never feels good about finding an answer on their own, and the reward pathway for self-directed exploration gets weaker. Students may become less curious, less willing to deal with uncertainty, and more reliant on people outside of the classroom to give them answers over time.

5. Pedagogical and Equity Concerns

5.1 The De-Skilling of Teachers and Learners

One less talked about effect of AI self-learning is that it could make both students and teachers less skilled. As previously stated, metacognitive and problem-solving skills may deteriorate for students. For educators, the extensive accessibility of AI tutors might diminish the motivation to cultivate robust pedagogical content knowledge. Why pay for expert teaching if students can learn calculus from AI? If this logic were used, it would strip professional teaching of its value, leaving only technical support roles. There would be a two-tier system: rich schools with human experts who teach how to learn, and poor schools with AI terminals that give information but not wisdom.

5.2 The Digital Divide Articulated

Not all AI self-learning platforms are easy to get to. To use high-quality platforms, you need a reliable internet connection, modern devices, and sometimes a paid subscription. Khanmigo costs $4 to $8 a month, and advanced ChatGPT features cost $20 a month. Students who don't have these things can't participate. But the digital divide is now at a second level: being ready for digital. Students who have learned how to learn on their own through previous schooling, parental support, or their own natural tendencies use AI more effectively. Students who don't have these skills use AI in a shallow or harmful way. AI makes achievement gaps bigger instead of closing them. The wealthy become more knowledgeable, while the less fortunate have fewer chances to learn new things.

5.3 Bias and Cultural Homogenisation

AI training data primarily embodies Western, educated, industrialised, rich, and democratic (WEIRD) viewpoints (Baek & Doleck, 2024). For students who learn on their own, this means they are only exposed to a small number of cultural works. When a student uses AI to learn about history, they get a story that is centred on Europe. A student studying literature is primarily advised to read white, male, English-speaking authors. A student studying

ethics encounters utilitarian or deontological frameworks while Indigenous or Eastern philosophies are overlooked.

AI platforms may also actively discourage views that are not standard. When a student asks a question that goes against the mainstream story, the AI often gives a "balanced" answer that defaults to the mainstream view, which pushes other views to the side. Self-learning is meant to help people think critically and learn about different points of view. AI platforms are made to give safe, average, and consensus outputs, which is the opposite of intellectual diversity.

5.4 Privacy and Data Exploitation

Self-learning produces detailed information about a student's difficulties, confusions, learning speed, mistakes, and emotional responses in typed questions. Most AI platforms gather, keep, and analyse this data to make their models better, often without clear or full consent. Students and their parents do not often understand the data lifecycle. Can a platform take advantage of a student's confusion about fractions to train a model that is then sold to an ed tech company? Yes, in many cases, because terms of service give people a lot of rights.

The risks over the long term are very high. Data from learning could be used to make profiles, predict how well someone will do in school, or even screen job applicants. If the data from an AI tutor about a student's early problems with mathematics is shared or sold, it could affect their future chances. Self-learning turns into self-surveillance (Alleblas & Dorrestijn, 2025), with the learner being both the subject and the product.

5.5 The Absence of Social Learning

Self-learning from AI is inherently solitary. A large body of research, on the other hand, shows how important social learning is: explaining ideas to peers, arguing about different interpretations, getting emotional support, and building knowledge together. AI cannot offer authentic social engagement (Ezeanya et al., 2024). It can act like it cares and encourages you, but these are just simulations. Students who use AI to learn miss out on chances to work together, learn from different points of view, and feel the emotional parts of learning, like frustration, breakthroughs, and shared joy.

Some platforms try to add social features like discussion forums and peer comparison, but these features are often not used enough or are just for show. The basic structure of AI self-learning is based on individualism, which fits with neoliberal ideas about education as a way to learn skills for yourself rather than for the common well or as a group effort.

6. Recommendations and Conclusion

6.1 For Students: Becoming Critical AI Users

If students use AI platforms in the right way, they can be very useful. First, do not use AI as an oracle for answers; use it as a Socratic partner. Tell the AI to ask you questions instead of answering them. "Explain photosynthesis" is not as helpful as "Ask me five questions to check my understanding of photosynthesis". Second, make the AI show you how it got to its answer. For mathematics and science, ask for step-by-step reasoning and then explain it in your own words. Third, make sure important facts are correct by checking them against reliable sources like textbooks and peer-reviewed articles. Fourth, let AI study sessions happen where you solve problems, write summaries, or come up with explanations without using any digital tools. Fifth, think about how you used AI after each session: Did you learn something that you couldn't have learned any other way? Did you think you should have done it yourself when you hired someone else to do it?

6.2 For Educators: Reclaiming Pedagogical Agency

Teachers should not ignore AI self-learning; instead, they should include it in a bigger plan for how they teach. Explicitly teach metacognition: show students when they are relying too much on AI, how to keep track of their own understanding, and how to use AI as a tool instead of a crutch. Make tests that AI cannot pass, like oral examinations, performances in class, process portfolios, and new problem situations. Make AI activities that students can work on together, like having them compare AI outputs, criticise AI explanations, or use AI to come up with arguments against their own views. Demonstrate critical AI use by talking out loud while using AI yourself and showing how to check your work and ask yourself questions.

6.3 For Platform Designers: Building for Learning, Not Engagement

The business model for current platforms is to get people to spend more time on the site and ask more questions. A design that focuses on learning would work best for different results. Design for metacognitive prompting: add forced pauses where the AI asks, "Do you really get this?" "Explain it to me before I go on." Support diagnostic assessment by keeping track of not only right and wrong answers, but also patterns of mistakes, and giving feedback that addresses specific misunderstandings. Use fading support: as the student shows that they have mastered the material, slowly take away the scaffolding, which forces them to work hard. Be open: show students how sure they are of each answer, where the facts come from, and what the AI doesn't know. Do not let engagement hacking get to you: stay away from infinite scrolling, notification badges, and other design patterns that make you want to use something more than learn.

6.4 For Policymakers: Regulation and Investment

Policymakers need to think about the structural conditions that make AI self-learning possible. Require AI platforms used in education to be open and easy to check they should have to tell you where they get their training data, how accurate it is in each area, and what biases they know about. Set rules for protecting data in schools AI: Do not let businesses use student learning data for commercial purposes without clear, detailed, and revocable consent. Put money into digital literacy programs so that all students learn how to judge AI outputs, handle cognitive offloading, and keep their metacognitive skills sharp. Pay for thorough, unbiased reviews of AI learning platforms, with planned trials and follow-ups over a long period of time. Don't replace human teachers with AI. Instead of buying AI

licenses to save money, use the money to make classes smaller and give teachers more training.

6.5 Conclusion

This review has contended that AI platforms provide authentic yet constrained opportunities for student autonomous learning. They can help you practise, give you feedback right away, and let you explore without any risk. But they also put metacognition, deep processing, critical thinking, and fairness at a lot of risk. The evidence base is less robust than promotional rhetoric indicates, and the pedagogical risks are undervalued.

The most profound critique is philosophical: learning transcends the mere acquisition of information or skills. It is the learner's change—the growth of their intellectual character, curiosity, resilience, and judgement. These traits are developed through adversity, social engagement, and the gradual process of grappling with concepts until they are internalised. AI cannot go through this change, no matter how fluent it is. It can only pretend to be its products.

We must not give in to the temptation of frictionless education if we want to keep the human part of learning. We need to make AI tools that help learners instead of replacing their own work. We need to teach students not only how to use AI, but also what it can and can't do, as well as what it doesn't teach them. The goal is not to stop AI from learning on its own, but to use it in a way that is critical, open, and always in the best interest of people. The student who is sitting alone at a desk with an AI on the screen is still a person on a journey. Technology should help with that journey, not take its place.

References

Alleblas, J., & Dorrestijn, S. (2025). 'Care of the Self' and Discipline in Smart Cities: Sensors in Singapore. In *Contemporary Practices of Citizenship in Asia and the West* (pp. 173-193). Routledge.

Akhmedov, B. A., & Yuldasheva, T. A. (2023). Socratic dialogue as a response to the challenge of the Epoch. *Sciental Journal of Education Humanities and Social Sciences*, *1*(1), 1-9.

Baek, C., & Doleck, T. (2024). Learning Analytics: A Comparison of Western, Educated, Industrialized, Rich, and Democratic (WEIRD) and Non-WEIRD Research. *Knowledge Management & E-Learning*, *16*(2), 217-236.

Becerra, R. C. A. (2025). Integrating Eye-Tracking and Artificial Intelligence for Human-Centered Visual Attention Analytics in Online Learning. *IE Comunicaciones: Revista Iberoamericana de Informática Educativa*, (42), 21-32.

Bhandari, B., Chopra, D., & Singh, K. (2020). Self-directed learning: assessment of students' abilities and their perspective. *Advances in physiology education*, *44*(3), 383-386.

Danyaro, K. U., Abdullahi, S., Abdallah, A. S., & Chiroma, H. (2025). Hallucinations in large language models for education: Challenges and mitigation. *International Journal of Teaching, Learning and Education*, *4*(6), 639993.

Ezeanya, C. U., Ukaigwe, J. A., Ogbaga, I. N., & Kwanashie, A. (2024). Enhancing Social Engagement among Online Learners' Using AI-Driven Tools: National Open University of Nigeria Leaners' Perspective. *ABUAD Journal of Engineering Research and Development*, *7*(2), 78-85.

Henkel, O., Horne-Robinson, H., Kozhakhmetova, N., & Lee, A. (2024). Effective and scalable math support: Evidence on the impact of an AI-tutor on math achievement in Ghana. *arXiv preprint arXiv:2402.09809*.

Higher Education Policy Institute. (2025). *Student Generative AI Survey 2025.* Oxford: Higher Education Policy Institute.

Pillai, R., Sivathanu, B., Metri, B., & Kaushik, N. (2024). Students' adoption of AI-based teacher-bots (T-bots) for learning in higher education. *Information Technology & People*, *37*(1), 328-355.

Ramos, L. T., & Casas, E. (2025). Applications, trends, and perspectives of large language models in education: A literature review. *Authorea Preprints*, 1-18.

Stadnik, M., Lokas Ćošković, A., Banko, T., & Ružak, D. (2025). AI and Writing Skills: Students' Attitudes towards Using AI to Enhance their Writing Based on the Example of Algebra University Students. *Interdisciplinary Description of Complex Systems: INDECS*, *23*(1), 46-71.

Sweller, J. (2024). Cognitive load theory and individual differences. *Learning and Individual Differences*, *110*, 102423.

Zhou, J. (2020). Design of AI-based self-learning platform for college English listening. In *2020 2nd International Conference on Machine Learning, Big Data and Business Intelligence (MLBDBI)* (pp. 544-547). IEEE.