

Artificial Intelligence in Teaching and Learning

Dr Chethan H

Associate Professor

Sapthagiri NPS University, Bangalore

ORCID id - <https://orcid.org/0009-0008-4821-9518>



Date of Submission: 01-07-2026

Date of Acceptance: 15-07-2026

Date of Publication: 01-08-2026

Abstract— Artificial intelligence is increasingly reshaping how teaching and learning are designed, delivered, and experienced across educational institutions worldwide. This paper examines the role of artificial intelligence in teaching and learning, tracing its evolution from early intelligent tutoring systems to contemporary generative and large language model based tools such as ChatGPT. Drawing on published research from higher education and educational psychology, the paper synthesizes evidence on the pedagogical benefits of adaptive tutoring, automated feedback, and personalized learning pathways, while also examining the growing body of literature on the risks and limitations of these technologies, including concerns about academic integrity, over-reliance, and the weak connection between many artificial intelligence applications and established pedagogical theory. The paper further considers how artificial intelligence is being applied across different phases of the teaching and learning process, from content delivery and assessment to administrative support for educators. While the evidence reviewed here indicates that artificial intelligence based tools can produce measurable learning gains under appropriate conditions, the paper concludes that realizing these benefits depends heavily on thoughtful pedagogical integration, teacher preparedness, and continued attention to the ethical and equity implications of these rapidly evolving technologies.

Keywords— artificial intelligence, teaching and learning, intelligent tutoring systems, generative artificial intelligence, higher education, educational technology


INTRODUCTION

The integration of artificial intelligence into education has accelerated dramatically over the past decade, moving from a specialized area of educational technology research into a mainstream concern for teachers, students, and institutional leaders alike. Early applications of artificial intelligence in education were largely confined to intelligent tutoring systems designed to provide step by step guidance in well structured domains such as mathematics and physics. More recently, the emergence of generative artificial intelligence and large language models has expanded the scope of artificial intelligence in education dramatically, enabling tools that can draft essays, explain complex concepts in natural language, generate practice questions, and provide instant feedback across a much wider range of subjects and skill levels than earlier systems could support.

This rapid expansion has generated both considerable enthusiasm and considerable concern within the education community. On one hand, proponents argue that artificial intelligence offers a genuine opportunity to personalize instruction at scale, providing each learner with a pace and style of instruction tailored to their individual needs in a way that has historically only been possible through one on one human tutoring. On the other hand, critics caution that the rapid adoption of artificial intelligence tools in classrooms has often outpaced the development of sound pedagogical frameworks for their use, raising concerns about superficial engagement, academic integrity, and the potential erosion of foundational skills if students come to rely on artificial intelligence tools as



a substitute for, rather than a support for, their own thinking and learning processes.



As Popenici and Kerr (2017) observed in one of the earlier influential discussions of this topic, the real potential of technology in higher education lies in its capacity to extend human capabilities in teaching, learning, and research, rather than to serve as a wholesale replacement for sound pedagogy. This framing continues to inform much of the subsequent literature on artificial intelligence in education, which consistently emphasizes that the educational value of artificial intelligence tools depends less on the sophistication of the underlying technology and more on how thoughtfully that technology is integrated into existing pedagogical practice.

This paper reviews the existing literature on artificial intelligence in teaching and learning, drawing on research spanning intelligent tutoring systems, systematic reviews of artificial intelligence applications in higher education, and more recent scholarship addressing the implications of generative artificial intelligence and large language models. It then outlines the objectives and methodology of the study, before examining specific applications of artificial intelligence across the teaching and learning process. The paper subsequently considers the principal challenges and limitations associated with artificial intelligence in education, before concluding with recommendations for institutions and educators seeking to integrate these tools effectively.

LITERATURE REVIEW

The academic literature on artificial intelligence in education has developed substantially over several decades, beginning with foundational research on intelligent tutoring systems and expanding more recently to encompass generative artificial

intelligence applications. VanLehn (2011) conducted an influential review of experiments comparing the effectiveness of human tutoring, intelligent tutoring systems, and other forms of computer based tutoring, finding that the effect size of human tutoring was considerably lower than previously assumed, while intelligent tutoring systems were found to be nearly as effective as human tutors across the studies examined. This finding was significant because it challenged the long standing assumption that human tutoring was substantially more effective than computer based tutoring, suggesting instead that well designed intelligent tutoring systems could approach the effectiveness of individualized human instruction under many conditions.

Building on this line of inquiry, Kulik and Fletcher (2016) conducted a meta-analysis of fifty controlled evaluations of intelligent tutoring systems, finding that the median effect of intelligent tutoring was to raise test scores by approximately two thirds of a standard deviation over conventional instruction, equivalent to moving a typical student from the fiftieth to the seventy fifth percentile. Importantly, their analysis also found that the magnitude of this effect depended considerably on whether learning gains were measured using locally developed assessments or standardized tests, suggesting that the alignment between the content and objectives of an intelligent tutoring system and the assessment used to evaluate it is a critical determinant of measured effectiveness. This methodological caution remains an important consideration for subsequent research evaluating the impact of artificial intelligence tools on learning outcomes.

Beyond the specific domain of intelligent tutoring, several works have taken a broader view of artificial intelligence in education. Luckin, Holmes, Griffiths, and Forcier (2016) offered an influential argument for the potential of artificial intelligence in education, explaining the underlying goals and mechanisms of artificial intelligence in education to a non-specialist audience while making the case that these technologies could help address persistent challenges such as achievement gaps, provided that their design and deployment remain grounded in a clear understanding of how students actually learn. Their work emphasizes that artificial intelligence in education should be understood not as a single technology but as a diverse set of tools and approaches, each suited to different pedagogical purposes, ranging from adaptive practice systems to tools that support collaborative learning and metacognitive reflection.

Complementing this conceptual work, Holmes, Bialik, and Fadel (2019) provided a comprehensive treatment of the promises and implications of artificial intelligence in education,


examining how these technologies intersect with core questions of curriculum design, assessment, and the changing skill demands of a world increasingly shaped by automation and intelligent systems. Their analysis is notable for situating artificial intelligence in education within a broader discussion of what students need to learn in an artificial intelligence enabled world, in addition to how artificial intelligence can be used as a tool for teaching and learning, thereby connecting questions of educational technology to broader questions of curriculum reform.

A significant contribution to understanding the empirical landscape of this field comes from Zawacki-Richter, Marín, Bond, and Gouverneur (2019), who conducted a systematic review of research on artificial intelligence applications in higher education, screening 2,656 initially identified publications and including 146 articles in their final synthesis. Their review found that existing research on artificial intelligence in higher education could be broadly organized into four areas, including profiling and prediction, assessment and evaluation, adaptive systems and personalization, and intelligent tutoring systems. Notably, the review also found a striking absence of critical reflection on the challenges and risks associated with artificial intelligence in education within much of the reviewed literature, along with a weak connection between many studies and established theoretical perspectives in pedagogy, leading the authors to pointedly ask where the educators were within this rapidly growing body of research.

More recently, the emergence of large language models and generative artificial intelligence tools such as ChatGPT has prompted a new wave of scholarship examining both the opportunities and risks these tools present for education. Kasneci and colleagues (2023) provided an extensive analysis of the opportunities and challenges associated with large language models in education, highlighting potential applications including personalized learning support, the generation of practice materials, and assistance with code explanation and writing tasks, while also raising significant concerns regarding plagiarism, the accuracy and reliability of generated content, and the risk that easy access to fluent, on demand text generation could undermine the development of students' own writing and reasoning skills if not thoughtfully integrated into assessment design. Their analysis is particularly valuable for highlighting the need for educators to develop new competencies in evaluating and guiding the appropriate use of these tools, rather than treating their adoption as a purely technical decision.

Taken together, this literature traces an evolution from early intelligent tutoring systems, which were grounded in relatively

well defined models of student knowledge within narrow subject domains, toward more recent, broadly capable generative artificial intelligence tools whose educational applications and risks are still being actively mapped by the research community. Across this evolution, a consistent theme emerges: while the underlying artificial intelligence technology has become considerably more powerful and flexible over time, the central pedagogical challenge, ensuring that these tools are integrated in ways that genuinely support rather than substitute for student learning, remains as pressing today as it was in the earliest era of intelligent tutoring research.



OBJECTIVES OF THE STUDY

This study pursues three central objectives. First, it aims to trace the evolution of artificial intelligence applications in teaching and learning, from early intelligent tutoring systems through to contemporary generative artificial intelligence tools. Second, it seeks to synthesize existing empirical and conceptual research on the pedagogical benefits and limitations of these technologies across different phases of the teaching and learning process. Third, it aims to identify the principal challenges facing the effective and responsible integration of artificial intelligence into educational practice, and to propose directions for institutions and educators seeking to adopt these tools thoughtfully.

RESEARCH METHODOLOGY

This study adopts a descriptive and analytical research design based on secondary sources. Evidence was drawn from peer reviewed journal articles, systematic reviews, and influential edited reports addressing the application of artificial intelligence technologies to teaching and learning, spanning both foundational research on intelligent tutoring systems and more recent scholarship on generative artificial intelligence and large language models in education. The literature was selected to represent both the historical development of the field and its current state of practice, allowing the analysis to capture how understanding of artificial intelligence's educational potential

and limitations has evolved over time. The analytical approach applied in this paper is qualitative and interpretive, synthesizing findings across studies to identify recurring themes, methodological considerations, and unresolved challenges, rather than testing a specific quantitative hypothesis. This method is well suited to the objective of the paper, which is to provide an integrative overview of a rapidly evolving and methodologically diverse field.

Applications of Artificial Intelligence in Teaching and Learning

Artificial intelligence applications in teaching and learning can be organized around several major functions that together span the instructional process from initial content delivery through to ongoing assessment and support.



The first and most extensively studied function is intelligent tutoring, in which artificial intelligence systems provide step by step or substep by substep guidance to students working through problems in domains such as mathematics, physics, and computer programming. As demonstrated by both VanLehn (2011) and Kulik and Fletcher (2016), these systems have been shown to produce meaningful learning gains relative to conventional instruction, particularly when the granularity of feedback provided by the system closely matches the level of detail at which students need support. Modern intelligent tutoring systems increasingly incorporate models of student knowledge that update dynamically as a student progresses, allowing the system to adjust the difficulty and type of problems presented in real time.

A second major function is personalized and adaptive learning, in which artificial intelligence systems adjust the sequence, pace, and content of instructional material based on an individual student's demonstrated performance and learning patterns. Luckin et al. (2016) emphasize that this capability represents one of the most promising applications of artificial intelligence in education, since it offers the possibility of tailoring instruction to individual learners at a scale that would

be impractical for a single human teacher managing a large classroom. Adaptive learning platforms are increasingly used to support mastery based progression, in which students advance to new material only once they have demonstrated sufficient proficiency with prerequisite content.

A third function concerns automated assessment and feedback. Artificial intelligence systems are increasingly capable of scoring student responses, including relatively open ended written responses, and providing immediate, formative feedback that would be impractical for a human instructor to deliver at the same speed and scale, particularly in large enrollment courses. Zawacki-Richter et al. (2019) identify assessment and evaluation as one of the major established areas of artificial intelligence application in higher education, noting that these tools are used both to support summative grading and to provide low stakes formative feedback that helps students identify and correct misunderstandings before high stakes assessments.

A fourth function, which has expanded dramatically with the emergence of generative artificial intelligence, involves content generation and explanation. Large language model based tools can now generate practice questions, explain concepts in multiple ways tailored to a learner's apparent level of understanding, summarize complex material, and assist students with drafting and revising written work. Kasneci et al. (2023) note that these capabilities offer considerable potential to support learning, particularly for students who benefit from additional explanation or practice beyond what is available in a standard classroom setting, while also cautioning that the same capabilities create new risks around academic integrity and the potential for students to bypass rather than engage in the cognitive work that supports genuine learning.

Finally, artificial intelligence is increasingly applied to support educators themselves, through tools that assist with lesson planning, the generation of assessment items, and the analysis of learning analytics data to identify students who may be at risk of falling behind. Holmes et al. (2019) note that these administrative and analytical applications of artificial intelligence, while less visible to students than tutoring or content generation tools, may have an equally significant impact on teaching and learning by freeing educators from routine administrative tasks and allowing them to focus more of their attention on direct engagement with students.

CHALLENGES AND LIMITATIONS

Despite the substantial potential benefits documented in the literature, several challenges continue to constrain the effective

integration of artificial intelligence into teaching and learning. The first and most frequently discussed challenge concerns academic integrity. As generative artificial intelligence tools have become more capable of producing fluent, plausible written work, concerns have grown regarding the potential for students to submit artificial intelligence generated work as their own, undermining the validity of traditional assessment approaches. Kasneci et al. (2023) argue that addressing this challenge requires rethinking assessment design itself, rather than relying solely on detection tools, since assessments that emphasize process, reasoning, and in-person demonstration of understanding are inherently more resistant to inappropriate artificial intelligence use than assessments that rely primarily on the production of a finished written product.

A second challenge concerns the risk of over-reliance on artificial intelligence tools in ways that may undermine the development of foundational skills. If students routinely turn to artificial intelligence tools to generate answers or explanations rather than working through problems independently, there is a risk that they may fail to develop the deeper conceptual understanding and problem solving skills that sustained engagement with challenging material is intended to build. This concern echoes long standing debates in educational psychology regarding the appropriate balance between providing support and allowing productive struggle, a balance that becomes considerably more difficult to manage as artificial intelligence tools become more capable and more readily accessible.

A third challenge, highlighted by Zawacki-Richter et al. (2019), concerns the weak connection between much of the existing research on artificial intelligence in education and established theoretical perspectives in pedagogy. Many studies of artificial intelligence applications in education have focused primarily on technical performance metrics, such as prediction accuracy or system responsiveness, without adequately grounding their design or evaluation in theories of how students learn, potentially limiting the pedagogical value of otherwise technically sophisticated systems.

A fourth challenge relates to equity and access. The benefits of artificial intelligence in education are not necessarily distributed evenly across institutions or student populations, since access to high quality artificial intelligence tools often depends on institutional resources, reliable internet connectivity, and appropriate devices, all of which remain unevenly distributed both within and across countries. Without deliberate attention to these disparities, the introduction of artificial intelligence tools risks widening rather than narrowing existing educational inequalities.

A fifth challenge concerns teacher preparedness. Effective integration of artificial intelligence into teaching and learning requires educators to develop new competencies, including the ability to critically evaluate the outputs of artificial intelligence tools, to design assessments that are resistant to inappropriate use of these tools, and to guide students in using artificial intelligence as a support for, rather than a substitute for, their own thinking. Many educators currently lack access to the professional development needed to build these competencies, creating a gap between the pace of technological adoption and the pace of pedagogical adaptation.

RECOMMENDATIONS AND THE WAY FORWARD

Addressing these challenges requires coordinated action across several dimensions. Educational institutions should invest in professional development programs that help educators build the competencies needed to critically evaluate and thoughtfully integrate artificial intelligence tools into their teaching practice, rather than leaving individual instructors to navigate these decisions without support. These programs should address not only the technical functioning of artificial intelligence tools but also the pedagogical principles that should guide their use.

Assessment design should be reconsidered in light of the capabilities of generative artificial intelligence tools, with greater emphasis placed on assessment approaches that evaluate process, reasoning, and the application of knowledge in novel contexts, in addition to the production of finished written work. This shift can help preserve the validity of assessment while also modeling for students the kinds of higher order thinking that artificial intelligence tools are best used to support rather than replace.

Institutions should also give explicit attention to equity considerations in the adoption of artificial intelligence tools, ensuring that access to these technologies does not become an additional source of educational advantage for already well resourced students and institutions. This may involve prioritizing investment in artificial intelligence tools and supporting infrastructure for under-resourced institutions and student populations.

Finally, continued research investment is needed to strengthen the connection between artificial intelligence applications in education and established pedagogical theory, addressing the gap identified by Zawacki-Richter et al. (2019). Closer collaboration between artificial intelligence researchers, learning scientists, and practicing educators throughout the design and evaluation of these tools is likely to produce

applications that are both technically sophisticated and genuinely grounded in an understanding of how students learn.

with which they are integrated into the broader practice of teaching and learning.

CONCLUSION

Artificial intelligence has become an increasingly significant force shaping teaching and learning across educational contexts, offering genuine potential to personalize instruction, provide timely feedback, and extend the reach of high quality educational support beyond what has traditionally been possible through human instruction alone. The literature reviewed in this paper, spanning several decades of research from early intelligent tutoring systems to contemporary generative artificial intelligence tools, demonstrates that these technologies can produce meaningful learning gains under appropriate conditions, while also revealing persistent challenges related to academic integrity, over-reliance, weak pedagogical grounding, equity, and teacher preparedness. Realizing the genuine potential of artificial intelligence in teaching and learning will require sustained attention to these challenges, alongside continued investment in professional development, assessment redesign, and research that keeps pace with the rapid evolution of the underlying technology. As artificial intelligence tools continue to grow more capable and more widely accessible, the central lesson of this literature remains clear: the educational value of these tools depends far less on their technical sophistication than on the thoughtfulness

REFERENCES

- [1] Holmes, W., Bialik, M., & Fadel, C. (2019). Artificial Intelligence in Education: Promises and Implications for Teaching and Learning. Center for Curriculum Redesign.
- [2] Kasneci, E., Sessler, K., Kuchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Gunnemann, S., Hullermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., Stadler, M., Weller, J., Kuhn, J., & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274.
- [3] Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), 42-78.
- [4] Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). *Intelligence Unleashed: An Argument for AI in Education*. Pearson.
- [5] Popenici, S. A. D., & Kerr, S. (2017). Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning*, 12, Article 22.
- [6] VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197-221.
- [7] Zawacki-Richter, O., Marin, 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, Article 39.