Available Online: 30 October 2026 DOI: 10.54254/2753-7102/2026.38444

# AI Models empowering teaching in higher education: current applications, practical dilemmas and pathways for improvement

***Zhongqiang Sun*[1], *Shuyi He*[1,2,*]**

[1]Lingnan Normal University, Zhanjiang, China
[2]South China Normal University, Guangzhou, China

*Corresponding Author. Email: 19886096909@163.com

**Abstract.** As generative artificial intelligence continues to iterate, mainstream models such as ChatGPT, DeepSeek and Copilot are entering the teaching spaces of universities, and teaching processes, teacher–student relations and assessment practices are changing with them. Drawing on a literature review and case analysis, this paper maps the deployment landscape and application picture of mainstream Artificial Intelligence (AI) models in higher education teaching and assesses both the effects that have already emerged and the problems that have been exposed. The study argues that the evidence is relatively clear for gains in teaching efficiency, in the accessibility of personalized learning support, and in the extension of teaching across time and space, whereas risks to content reliability caused by hallucination, cognitive reliance and uneven literacy among teachers and students, and the absence of regulation together with ethical controversy constitute the main obstacles to wider diffusion. On this basis, the paper proposes a coordinated approach spanning technology, literacy and institutional design.

**Keywords:** generative artificial intelligence, large language models, higher education teaching, teaching empowerment, academic ethics

## 1. Introduction

After ChatGPT was launched in late 2022, general-purpose artificial intelligence moved from the laboratory into everyday use faster than most observers had anticipated. In the two years that followed, GPT-4o, Gemini, Claude, DeepSeek and other models were released in quick succession, with language understanding, reasoning and multimodal interaction improving steadily; their position in educational activity has also shifted, gradually, from that of an "assistive tool" toward that of a "participant in teaching". Universities have felt this wave most directly. They are more alert to new technologies because they carry the dual mandate of knowledge production and talent development; and the text generation, information synthesis and knowledge question-answering at which large models excel happen to cover a considerable share of routine teaching work: compiling lecture notes, writing exercises and giving first-pass comments on assignments can all be taken over, in part, by a model.

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

In China, universities have adopted AI models faster than expected. Liu et al.'s survey of 147 "Double First-Class" universities shows that, by August 2025, 125 of them had completed localized deployment of large models represented by DeepSeek, an overall deployment rate of 85.03%; by tier, the rate was 97.44% for Project 985 universities, 90.43% for Project 211 universities and 69.63% for Double First-Class universities outside Project 211, with the differences between tiers statistically significant ($p < 0.001$) [1]. Deployment followed an S-shaped diffusion curve, with Project 985 universities taking the lead from early January to mid-February 2025, Project 211 universities following from late February to mid-March, and the pace leveling off thereafter. Comprehensive deployment predominated, at 72%, covering teaching support, research assistance, teaching management and learning assessment. Application scenarios expanded alongside: in the autumn semester of 2025 Tsinghua University opened more than 220 pilot courses in AI-empowered teaching and provided a "24-hour AI learning companion", so that students could obtain academic support outside class hours [2].

Behind these figures lies a point that is easily overlooked: deployment and use are two different things. Bringing equipment online and opening accounts are actions at the infrastructure level; what actually determines teaching outcomes is design at the course level—at which point a teacher calls on the model, in what manner model output is brought into the teaching process, and whether what students gain is convenience or substitution. Surveys have repeatedly shown that the former advances quickly while the latter lags behind. The Report on Teachers' Digital Learning (2025) concludes that the application of AI in higher education has moved past the question of whether teachers can use it and into a stage defined by whether they can use it deeply, well and to real effect; in practice, however, teachers' use of GenAI remains concentrated in preparation tasks such as searching for materials, designing lesson plans and making slides, while intervention in personalized teaching, intelligent assessment and more complex teaching situations is quite limited. Alongside this are several tensions that are difficult to avoid: the erosion of content reliability by hallucination; the possibility that students' reliance on AI weakens their thinking; shallow application resulting from insufficient AI literacy among teachers; and the lag in building rules for academic integrity and data privacy [3]. These problems are unlikely to dissolve on their own through model iteration; they require a systematic response in educational philosophy, institutional arrangements and teaching practice.

Using literature review and comparative case analysis, this paper examines the current state of mainstream AI models in higher education teaching, summarizes the effects and problems that have emerged, and attempts to propose workable pathways for improvement, as a reference for university administrators and frontline teachers.

## 2. Data and methods

The materials come from three channels. The first is academic literature, retrieved from CNKI, Web of Science and Scopus using combinations of Chinese and English keywords such as "AI large models", "generative artificial intelligence", "higher education", "teaching empowerment" and "GenAI in higher education", covering the period from January 2023 to March 2026, of which 68 items were selected for in-depth analysis. The second is policy texts and practical cases, including Tsinghua University's *Guiding Principles for the Application of Artificial Intelligence in Education*, *the Guidelines for Regulating the Use of Artificial Intelligence Tools in Graduate Theses and Practical Outcomes* developed by the Chinese Association of Degree and Graduate Education, and deployment data and teaching cases published by universities [4]. The third is empirical survey data, mainly questionnaire and interview results obtained by Chinese and international teams working with university teachers and students, used to reflect actual patterns

of use and how teachers and students perceive them. Among these, Liu et al.'s deployment statistics for 147 Double First-Class universities, Fudan University's questionnaire survey and interviews with 118 teachers, and Hao et al.'s stratified random sampling survey of teachers and students in 31 universities nationwide, based on the UTAUT2 model, constitute the main quantitative basis of this paper [5].

The study combines literature review with case analysis. The literature component organizes existing findings around four dimensions—application scenarios, technology types, effectiveness, and problems and challenges—and identifies points of consensus and disagreement. The case component selects AI teaching practices at Tsinghua University, Fudan University, Wuhan University, Beijing Forestry University and elsewhere, and conducts a descriptive analysis in the light of publicly disclosed implementation details and outcome data. It should be noted that this paper does not aim at quantitative meta-analysis; it seeks to present an overall picture and the key issues through systematic literature synthesis and case comparison. Because most of the cases relied upon come from better-resourced Double First-Class universities, the scope within which the judgments below apply should be limited accordingly.

# 3. Results

## 3.1. The application landscape of mainstream AI models in higher education teaching

AI models have not entered university teaching along a uniform path. Deployment at the infrastructure level has advanced quickly, but penetration into specific teaching activities is uneven: the closer an activity lies to the core of the classroom, the later AI has arrived.

### *3.1.1. Deployment status: from early exploration to scaled access*

The year 2025 can be regarded as the "year of DeepSeek" for Chinese universities. With its open-source status, low deployment threshold and strong adaptability, DeepSeek lowered the technical and economic cost of localized deployment, and infrastructure building accelerated. The statistics cited above cover local on-campus deployment at 147 Double First-Class universities between 1 January and 30 August 2025; the results are shown in Table 1.

Regionally, coastal provinces such as Guangdong are in the lead. South China University of Technology completed localized deployment of the full DeepSeek-R1 model (671B) on its "Chengwu" scientific computing platform, with 19 servers and 76 A800 GPU cards; in addition to serving its own teachers and students, it provides backend sharing services to 42 universities in the province through the Guangdong Education and Research Network[1]. Shanghai Jiao Tong University's "Zhiyuan-1" platform has deployed eight mainstream large models, DeepSeek among them, serving 38,000 teachers and students on campus in 2025. A pattern in which leading universities build their own systems and radiate capacity regionally is becoming a practical way to ease the shortage of computing power at smaller institutions. Yet this sharing model, while alleviating the shortage, creates new problems: the model service is operated and maintained by a single university, and other institutions can hardly influence model behavior, data retention or the pace of updates, nor are teaching-specific customization needs easily met in time. For local institutions without independentoperational capacity, a considerable distance remains between accessing a shared service and forming their own AI teaching capability.

[1] More than 40 universities in Guangdong complete localized deployment of the DeepSeek education system [EB/OL]. Sohu News, 2025-03-19. https://news.sohu.com/a/873029488_121897422 (in Chinese)

**Table 1.** Local on-campus deployment of DeepSeek at Double First-Class universities

| University category | Deployment rate |
|---|---|
| Project 985 universities | 97.44% |
| Project 211 universities | 90.43% |
| Double First-Class universities outside Project 211 | 69.63% |

Note: The three categories together account for 125 deployed universities, an overall deployment rate of 85.03%, of which 38 are Project 985 universities. The original source does not disclose the number of institutions in each tier separately, so only deployment rates are reported here. Overall figures are cited from Liu et al. (2025) [1].

*3.1.2. Teaching preparation: efficiency gains and a tendency toward the superficial*

Teaching preparation is the earliest and most widespread point of AI involvement, concentrated on lesson design, resource generation and item writing. A survey of 2,003 teachers at the University of Iceland found that about 39% used AI to generate assignment ideas, 35% to develop course materials, 30% to write examination questions and 24% to produce assessment rubrics [6]. In China, teachers at Harbin Engineering University can obtain interdisciplinary case materials within seconds through a locally deployed DeepSeek, markedly shortening the time spent searching for and synthesizing information during preparation. Fudan University's questionnaire data, however, point to another side. Of 118 teachers, 64.79% used AI writing tools (such as ERNIE Bot and Kimi) and 61.97% used personalized learning platforms (such as XuetangX), while the use rates for AI teaching assistants and assignment grading tools were only 25.35% and 14.63% respectively[2]. Technology remains largely confined to resource generation and task management and has not yet penetrated classroom interaction or teaching feedback. This "hot outer ring, cold core" distribution suggests that AI empowerment currently remains at the level of an efficiency tool, still some distance from a redesign of teaching.

The reason for this distribution is not hard to grasp. Resource-generation tasks have clear goals and verifiable outputs; teachers can judge output quality with some confidence, and the cost of trial and error is low. Classroom interaction and teaching feedback, by contrast, involve real-time judgment about learning, the organization of discourse and emotional response, where a mistake carries more immediate costs—hence more hesitation.

*3.1.3. Classroom implementation: AI teaching assistants and AI learning companions*

The most active forms in classroom teaching are AI teaching assistants and AI learning companions. Drawing on its Fundamentals of Electrical Engineering course, Wuhan University built an AI teaching assistant system based on multiple large models: in class, the AI scores students' answers in real time and flags their weak points; outside class, it answers questions at any time. Its distinguishing design is that answers come from three different models, so students must compare and judge, which to some extent reduces dependence on the reliability of a single model's output.

Tsinghua University's exploration goes further. The Chemical Engineering Thermodynamics course in the Department of Chemical Engineering set up a "24-hour AI learning companion", allowing students to hold in-depth conversations with AI outside class hours; the University Physics course used a learning companion to

[2] Survey data on teachers' application of AI in teaching at Fudan University [R]. 2025 International Conference on the Scholarship of Teaching and Learning in China and the 7th National Conference on the Scholarship of Teaching and Learning in Chinese Universities, 2025. (in Chinese)

collect questions before class, and the system analyzed the most frequent questions across the class in real time, generated "response cards" and pushed them to the teacher's terminal, so that classroom instruction could address students' actual confusion [2]. The control engineering course team at Beijing Forestry University experimented separately with an AI-driven mode of thinking training, developing a "Socratic AI teaching assistant" that connects textual, audio and visual interaction channels and constructs a four-stage dialogue structure of "context activation – scaffolded guidance – probing questions – metacognitive summary". Such designs go beyond knowledge question-and-answer, attempting to guide critical thinking through structured dialogue. In practice, these forms differ in the depth of their involvement.

*3.1.4. After-class tutoring and academic assessment: broader support, assessment still auxiliary*

After-class tutoring and personalized learning are the areas where AI's fit is best demonstrated. In 2026, Dalian University of Technology introduced "Lianxiaoli", an all-scenario AI learning companion, for more than 6,000 new students, covering teaching, learning, administration, assessment, research and innovation and responding in parallel to needs such as coursework questions, thesis writing, innovation competitions and career planning. A quasi-experimental study of 192 students found that the experimental group receiving GenAI personalized learning support significantly outperformed the traditional blended-teaching group on post-test scores and on self-regulated learning indicators.

Application in assessment is more cautious, focused mainly on automatic scoring, the generation of formative feedback and the analysis of learning data. Xidian University's "Xidian Smart Evaluation" system analyses data on student learning behavior to give teachers visual support for process-oriented assessment; Huazhong University of Science and Technology used agents to build a simulated project-based learning environment in its Automotive Suspension System Analysis course; Sun Yat-sen University experimented with intelligent marking in a University Physics Laboratory course. This caution is well founded: scoring is not merely a judgment of whether an answer is right or wrong but also a composite judgment about reasoning, process and quality of expression—dimensions that are precisely the weak points of current models. Most universities keep people in the judgment role, with models responsible only for initial screening, marking and data organization.

## 3.2. Effects that have emerged

Judging from existing research and frontline feedback, the teaching effects of AI models fall broadly into three directions. The most direct is improved teaching efficiency. AI can quickly retrieve, organize and generate teaching resources, freeing teachers from part of their repetitive work. An exploratory survey at Dresden University of Applied Sciences showed that GenAI provides clear support for routine tasks in content development, teaching delivery and formative assessment; the Report on Teachers' Digital Learning (2025) likewise shows that nearly 70% of university teachers have made large models part of their daily learning toolkit and that more than half use them to support teaching [3]. The second is making personalized learning support accessible. In traditional higher education, personalized tutoring is constrained by faculty numbers and time costs and is difficult to supply at scale. Through adaptive path design, immediate feedback and precise delivery, AI moves such support from a "scarce good" toward an "accessible good". A practical study of a university Java Programming course noted that an AI-based personalized learning support system can effectively improve learning outcomes and course satisfaction, while the learning companion in Tsinghua's Chemical Engineering Thermodynamics course extends teaching support from inside the classroom to beyond it. The third is the extension of teaching interaction across time and space. Students can obtain immediate help outside teaching hours, which matters especially for groups whose learning rhythms differ widely and which, to some extent, mitigates the divergence produced by a single uniform pace.

Of these three, the evidence for the first two is relatively solid, while the third rests more on teaching observation. It should also be noted that these effects have mostly arisen at better-resourced universities and in courses with higher levels of cooperation, so the sample itself carries selection bias. Once AI support enters settings with strained faculty, large classes or wide variation in student preparation, whether its stability and effects can be maintained lacks reliable empirical evidence.

## 3.3. Main problems

### *3.3.1. Technical level: hallucination and content reliability*

The "hallucination" of large language models—generating content that appears plausible but is factually wrong or fabricated—poses a direct reliability risk in teaching. In the University of Iceland survey, 76% of responding teachers named AI hallucination as a major obstacle to adopting AI tools. In disciplines with high demands on accuracy, such as medicine, law and engineering, errors in AI-generated content can directly mislead teaching. Universities generally lack mechanisms for fact-checking AI-generated teaching content; teachers mostly rely on personal experience when using AI for preparation or for generating materials, without a stable quality-assurance process. More troublesome is that hallucination often gives no obvious warning. Models deliver correct and incorrect content in equally fluent, equally assured tones, and users who lack relevant knowledge in the field can hardly detect anomalies while reading. This is especially disadvantageous for junior students, who do not yet have the background knowledge against which to compare and are prone to accepting the model's statements as authoritative conclusions.

### *3.3.2. Actor level: cognitive reliance and uneven literacy*

Problems on the student side deserve particular attention. A quantitative study of university students found an intriguing paradox: GenAI users reported higher self-rated confidence, yet scored significantly lower than low-frequency users on problem-solving without assistance [7]. The researchers called this the "hollow skills effect": reliance on AI creates an illusion of competence and erodes the cognitive processes on which in-depth analysis and strategic decision-making depend [7]. Another longitudinal study likewise showed that while ChatGPT can improve learning outcomes in the short term, students with weak background knowledge who rely heavily on AI show a declining trend over time in independent thinking and problem-solving. In a 2025 survey of more than 1,000 university teachers by the American Association of Colleges and Universities, 90% of respondents believed that generative AI would weaken students' critical thinking.

This conclusion, too, calls for restraint. Existing studies mostly use quasi-experimental or cross-sectional designs and cannot rule out self-selection: students who were already more reliant on AI and less motivated may also show lower problem-solving scores. Whether AI use causes a decline in ability, or whether those with weaker abilities are more inclined to use AI, the current evidence is insufficient for a definite answer. Regardless of the causal direction, however, the combination of "high confidence, low ability" is itself a signal that teaching needs to address. Moreover, when students hand over to AI the thinking they should do themselves, the core educational aim of higher education may be quietly conceded—and this concern does not disappear because the causal evidence remains unresolved.

On the teacher side, the problem is an uneven distribution of literacy. Fudan University's survey of 118 teachers showed that 69.5% had received no structured AI skills training and relied mainly on their own exploration in teaching; 55.93% found it difficult to link AI tools effectively with teaching objectives and the content system; and 22.03% considered AI platform configuration difficult as shown in Table 2. Differences in usage frequency across disciplines are equally clear: 46.3% in medicine and 21.9% in engineering, against 16.1% in the humanities and 5.09% in the social sciences as shown in Table 2. Regression analysis further shows that the need for systematic training has the most significant effect on teachers' perception of difficulty

($\beta$ = 0.27, $p$ < 0.001), followed by the difficulty of curriculum integration ($\beta$ = 0.19, $p$ = 0.006), while technical complexity and student adaptability did not reach statistical significance. The policy implication is fairly clear: the primary bottleneck restricting deeper teacher use of AI lies not in technical operation but in the absence of training support and of a methodology for integrating AI into teaching.

**Table 2.** Teachers' use of AI [3, 8]

| Indicator | Data |
|---|---|
| Already use AI tools in teaching | about 60% |
| Use frequently with good results | only 9.3% |
| No structured AI training | 69.5% |
| Difficult to align AI with teaching objectives | 55.93% |
| Consider AI platforms complex to operate | 22.03% |
| AI as a daily learning tool | nearly 70% |
| Regard AI hallucination as an adoption barrier | 76% |

*3.3.3. Institutional level: absence of regulation and ethical risk*

Academic integrity governance is caught in a cat-and-mouse standoff. Since 2025, AIGC detection has entered university academic evaluation at scale, with more than 60% of undergraduate institutions and 80% of graduate training units incorporating AIGC detection into thesis review [8]. The upper limits set for the proportion of AI-generated content differ across institutions: Sichuan University requires no more than 20% for the humanities and 15% for science, engineering and medicine; Tianjin University of Science and Technology and Guangdong Medical University set 40% as an early-warning threshold; Fuzhou University links the "AI rate" directly to thesis scores and stipulates that the detected value must not exceed 20% [9]. Beyond inconsistent standards, the limitations of detection technology itself create new problems: a scholar's original paper, based on three years of fieldwork, was once misjudged as AI-written.

Detection tools also carry a structural limitation. AIGC detection relies on statistical features of text, and these features are closely related to writing style and language habits: writers who favor concise sentences and orderly prose are more easily misjudged, and international students whose first language is not Chinese face a higher risk as well. Linking detection results directly to grades or to eligibility for thesis defense amounts to transferring the uncertainty of technology onto students. If governance relies only on technical detection, without clearly defining what constitutes reasonable use and without accompanying pedagogical guidance, it is likely to fall into a cycle in which detection and evasion continually outbid each other.

On institution building, some universities have begun forward-looking exploration. Tsinghua University's Guiding Principles for the Application of Artificial Intelligence in Education, released in 2025, provided the first comprehensive, tiered guidance and regulation for AI application on campus, establishing a stance of "active yet prudent" and setting out five principles: primary responsibility, compliance and integrity, data security, prudent and critical thinking, and fairness and inclusiveness. It explicitly prohibits the use of AI to replace academic training that should be undertaken by the individual, prohibits ghostwriting, plagiarism and fabrication, and also prohibits defense committee members and peer reviewers from using AI in review [4]. This approach—drawing red lines while also lighting green lights—is more constructive than a purely prohibitive style of management. Overall, however, teaching-related AI regulation remains absent at most universities, and key issues such as data privacy protection, disclosure mechanisms for AI use, and the boundaries of teachers' responsibility have yet to be adequately addressed.

# 4. Discussion

## 4.1. From "tool empowerment" to "ecosystem reconstruction"

Most universities still use AI models at the level of "tool empowerment", treating AI as an auxiliary means of raising the efficiency of existing processes. But AI's impact on teaching does not stop at efficiency. When models can generate teaching content, provide personalized feedback and even deliver conversational explanations, the underlying logic of teaching loosens: knowledge becomes less scarce and the teacher's position as "knowledge authority" is challenged; immediate feedback replaces delayed feedback and the rhythm of internalizing knowledge is rewritten; and the focus of assessment needs to shift from "correct results" to "credible process".

Understanding empowerment as tool substitution also tends to obscure another change: once technology enters the teaching system, what changes is often not only the efficiency of a particular step but the positional relations among the actors. Teachers move from lecturers to designers and gatekeepers, students from recipients to initiators of dialogue, and assessors from final arbiters to recorders of process. Once such changes occur, existing management rules, timetabling and assessment methods tend to look ill-fitting.

Kohout-Diaz's interview study of Chinese university teachers provides supporting evidence [10]. The teachers interviewed generally recognized the value of AI in teaching assistance and personalized feedback, while identifying three key deficiencies: the destabilization of teaching authority, the absence of emotional support, and students' tendency toward over-reliance. This suggests that AI's empowerment of teaching should not be reduced to a question of technical efficiency but examined within a coordinated framework of "technology – education system – teaching actors". When technological capability iterates rapidly while institutional arrangements and the capacity preparation of actors fail to keep pace, empowerment can turn into imbalance.

## 4.2. Disciplinary differences and differentiated strategies

AI models perform markedly differently across disciplines. The interview study cited above notes that STEM teachers reported more effective experiences of use, whereas humanities teachers saw clear limitations in AI's explanatory depth and historical contextualization, which constitute a ceiling for application. Fudan University's data further confirm this: the AI usage rate among medical teachers (46.3%) is more than nine times that in the social sciences (5.09%).

The source of the difference lies in the nature of knowledge production in different disciplines. Knowledge in STEM fields is highly structured and codifiable and fits better with the way models represent knowledge; the core value of the humanities lies in interpretive depth, sensitivity to context and the individuality of judgment—precisely what current models cannot easily replace. It follows that AI teaching empowerment should not pursue a one-size-fits-all standardized solution but requires differentiated design according to disciplinary characteristics: in STEM fields, AI-assisted problem solving, experimental simulation and immediate feedback can be advanced more actively; in the humanities, AI is better treated as a "thinking partner", using critical dialogue to sharpen argumentation rather than to generate conclusions directly.

A further difference relates to the types of task within a discipline. Even in the humanities, literature searching, format proofreading and corpus organization are equally suited to being delegated to models, whereas the construction of arguments and the judgment of viewpoints should not be outsourced. A more accurate formulation, then, may not be that "discipline A suits AI and discipline B does not", but that within the same discipline the degree of fit varies across teaching tasks. This suggests that differentiated design needs to operate at the level of the task rather than stopping at the level of the discipline.

### 4.3. Literacy and institutions: two key variables

Teacher literacy and institutional regulation are the two key variables constraining deeper empowerment. Regression results show that the need for systematic training explains more of teachers' perception of difficulty than technical complexity or student adaptability. In other words, the main contradiction in advancing AI teaching at universities today is not "whether the technology works well" but "whether teachers can use it and use it well".

Existing training focuses mostly on tool operation and lacks systematic training in methodologies for integrating AI into teaching. Teachers need not only to be able to use AI but also to understand the limits of a model's capability, to master ways of critically appraising AI-generated content, and to embed AI appropriately in instructional design. This means that training should be upgraded from "tool training" to "pedagogy training", incorporating AI literacy into the framework of Technological Pedagogical Content Knowledge (TPACK). At the institutional level, the problem is a preference for constraint over guidance. The "active yet prudent" stance of Tsinghua University's Guiding Principles offers a model worth drawing on: regulation should not stop at prohibition but should give teachers and students clear guidance on how to use AI reasonably [4]. A more constructive approach is to embed usage rules in course design — specifying at the course level which activities permit the use of AI, which require independent completion, and how AI use should be declared and assessed—so that regulation becomes an organic part of the teaching process rather than an external constraint.

From the perspective of implementation sequence, institution building is not equally difficult in all respects. Disclosure of use and the delineation of permitted scope at the course level are low-cost starting points that can proceed once a department reaches consensus; teacher training and the adjustment of assessment standards, by contrast, involve resource investment and organizational coordination and take longer. A pragmatic path is to clarify the boundaries first and then address capacity and incentives step by step.

## 5. Conclusions

AI models are substantively changing the practice of teaching in higher education. The judgments of this study can be summarized in three points.

First, the effectiveness of AI models in empowering higher education teaching has been preliminarily verified, with notable performance in raising efficiency, expanding access to personalized learning support and extending teaching across time and space; but the realization of these effects depends heavily on contextual fit and on teachers' instructional design capacity, and technology does not empower automatically. About 60% of teachers already use AI in teaching, yet only 9.3% "use it frequently with good results", with the lack of training and the difficulty of curriculum integration the main constraints.

Second, the application of AI models faces three challenges—technical reliability, actors' cognitive reliance and the absence of institutional regulation—of which cognitive reliance is the most fundamental. The illusion of competence revealed by the "hollow skills effect" indicates that the use of AI may imperceptibly erode the educational aims of higher education. This problem is difficult to resolve through technical solutions; it requires returning to instructional design itself and maintaining the depth of students' cognitive engagement through carefully designed human–machine collaborative tasks.

Third, the key to addressing these difficulties lies in building a coordinated pathway of "technical optimization – literacy development – institution building": on the technical side, strengthening fact-checking and disciplinary fit; on the actor side, moving teachers' AI literacy from tool operation toward pedagogical integration; and on the institutional side, establishing a framework of norms that combines guidance and

constraint. The "active yet prudent" stance of Tsinghua University's Guiding Principles is a case worth observing in this direction.

The analysis also suggests that discussion of the relationship between AI and teaching needs to shift from "whether to introduce it" to "how to organize it". The technology is already present, and if debate remains confined to prohibition versus permission, real teaching improvement is unlikely to occur. More valuable questions are how to arrange the division of labor between people and machines in specific courses, how to identify students' actual state of understanding, and how to make assessment reflect the process of thinking rather than the final text.

The acceptance of AI by higher education should not be merely a passive response to technology but should become a self-aware educational transformation. Technology is a means, not an end, and the educational essence of teaching should not be diluted by its introduction. How to harness AI's capabilities while safeguarding the core values that make us human—critical thinking, creative expression and responsible judgment—will be a question that higher education teaching reform must continue to address for a long time to come.

## Fundings

This study was supported by the Guangdong Provincial Higher Education Teaching Reform Project-"Reform and Practice of AI-Empowered Physical Geography Teaching Model" (PX-38251098), and 2024 University-Level Education and Teaching Reform Project of Lingnan Normal University-"Research on the Reform and Practice of Physical Geography Teaching Model Based on AI Technology" (PX-92463)

## References

[1] Liu, J., Xue, M. J., & Su, F. G. (2025). The rhythm and direction of technology deployment in higher education digitalization: A case study of DeepSeek deployment in 147 universities. *China Higher Education, 2025*(21), 47–55.

[2] Jiao, Y. X., & Yan, Z. X. (2026, May 28). After the "human–machine duel": An observation of Tsinghua University's exploration of deep integration of AI with education and teaching. *China Education Daily*.

[3] Teacher Work Department of the Ministry of Education & China Distance Education Magazine. (2026). Report on teachers' digital learning (2025).

[4] Chen, B. (2025, December 3). Not only drawing "red lines" but also lighting "green lights": Tsinghua University issues its first guiding principles for the application of artificial intelligence in education. *China Science Daily*.

[5] Hao, Z. X., Ma, L. T., Wen, W., & Zhang, Y. (2026). University teachers' and students' acceptance of GenAI and its influencing factors: Based on a survey of teachers and students in 31 universities nationwide. *Modern Educational Technology, 36*(4), 43–52.

[6] Schram, Á. B., Jóhannesdóttir, S., Einarsson, H., Bjarnadóttir, M. K., & Rogers, B. L. (2025). *Teachers' attitudes and perceptions of the usefulness of AI in academia: How should universities respond to a changing educational landscape?* University of Iceland.

[7] Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Article 1121, pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778

[8] He, X. Y., & Li, Q. (2025, July 2). The first graduation season after the new AIGC detection rules. *China Youth Daily*.

[9] Jing, X. Q. (2025, December 11). Multiple universities light "red and green lights" for artificial intelligence. *Science and Technology Daily*.

[10] Kohout-Diaz, M. (2026). Making sense of AI in teacher education: A qualitative study of perceptions, practices and pedagogical tensions. *Teaching and Teacher Education, 171*, Article 105342.