



Journal of Emerging Trends in Social Sciences and Humanities

A Double-Blind, Peer-Reviewed, HEC recognized [Y-category](#) Research Journal

E-ISSN: [3006-7898](#) P-ISSN: [3006-788X](#)

Volume 3 Issue 2, 2026, 253-264



Artificial Intelligence and Student Learning in Higher Education: Opportunities, Challenges, and Future Directions

Shahbaz Sikandar¹ Aleeza Nasir² Shazia Afzal³

¹ MPhil Scholar, University of Education Lahore, Multan Campus, Punjab, Pakistan.

² MPhil Scholar, University of Education Lahore, Multan Campus, Punjab, Pakistan.

³ MPhil Education, Bahauddin Zakariya University, Multan, Punjab, Pakistan.

Corresponding Author: Shahbaz Sikandar, Email: shahbazsikandar481@gmail.com

Article Information

Received:

2026-04-20

Revised:

2026-05-25

Accepted:

2026-06-24

Keywords

Artificial Intelligence; Student Learning; Higher Education; Qualitative Research; Interviews; AI-supported Learning; Academic Integrity.

ABSTRACT

There is growing interest in how Artificial Intelligence (AI) is impacting on teaching and learning in higher education, presenting new opportunities and possibilities for personalized learning, academic support, feedback and engagement. Its increased adoption, however, has sparked debate on academic honesty, techno-dependency, privacy and security of data, access disparities and the evolving role of teachers and students. This qualitative study explored participants' perceptions and experiences of the use of AI in student learning within higher education. The data collecting method employed was semi-structured interview method among the participants of the higher education institutions. Thematic analysis was used to identify common themes linked to the educational advantages, obstacles, and prospects for the use of AI in learning. The results showed that AI tools can help to improve access to information, give instant feedback, enable personalized learning, increase understanding of challenging concepts and facilitate independent learning. The participants also raised worries about the lack of accurate or unbiased information, limited critical thinking, plagiarism, ethics, and low AI literacy, and overreliance on AI-generated text. The results further underscored the importance of institutional policies, teacher guidance, student training, ethical frameworks, and equitable access to AI technologies. However, the study concludes that AI shows great promise in enhancing student learning in higher education, provided that it is utilized responsibly as an educational aid, not a replacement for human reasoning, engagement with teachers, and active student participation. The study has implications for universities, teachers, and policymakers interested in pedagogically utilizing AI in higher education ethically, inclusively, and meaningfully.



© 2026 by the authors. Licensee Sikandar, Nasir & Afzal. This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license <https://creativecommons.org/licenses/by/4.0/>

Publisher: Institute for Educating Environmental Resilience and Governance

1. Introduction

A.I. in higher education is not just a curiosity anymore but a commonplace in the daily life of students, faculty, students' lives, and faculty's lives. Until recently, the use of AI in education was restricted to intelligent tutoring systems, adaptive courseware, automated assessment engines, and learning analytics platforms that are used at institutional level, and passengers, not drivers, of education. Large language models (LLMs) have turned the tables on that encounter. Students can now use generative tools, not needing to go through an institution, at low cost and no technical skills required. In seconds, without any teacher's knowledge, a student can get an explanation of a difficult concept, a critique of a draft, a translation, a summary of a dense reading or even an entire assignment. The focus of adoption has moved from an institutional to an individual one and higher educations are left with a responsibility that they should not have taken on and cannot easily manage (Batista et al., 2024).

This change has resulted in an ever-growing but strange polarization of literature. On the other side, AI is viewed as a tool of educational democracy, a patient tool which can provide the personalized attention that mass higher education has never been able to give, scaffold students who come underprepared, and help free teachers from menial work so they can focus on the relationship and thinking that only humans can offer. On the other hand, AI is cast as an enemy to the very essence of academia – to the objectivity of evaluation, the thinking process that leads to knowledge, the confidentiality of learners, and the expertise of teachers. Both narratives have a lot of truth to them and neither one is sufficient. The one thing that is still missing between them is a practical description of the use of AI by users and by learners who have to learn alongside AI (Kuleto et al., 2021).

This divide, in part, is due to the way the field has been examined. Empirical research has often been conducted through large scale surveys that measure attitudes to and acceptance of, and intention to use, AI or through experimental studies that measure the impact of a specified AI intervention on a defined learning outcome. Such designs set up distributions and effects with desirable precision, but are poorly suited to

capturing the reasoning underlying behaviour: why a student uses an AI tool at one time in an assignment and not another; what they think they receive and what they think they sacrifice; how they distinguish between legitimate use of AI and illegitimate substitution of it; and how those distinctions change when the assignment time limit is approaching. They are interpretive questions which require interpretive methods (Alotaibi, 2024).

There is an additional reason to pay attention to perception, in particular. According to Chan and Hu (2023), who are referring to Biggs's model of student learning, the students' cognitive experience of a task, a tool or an environment has a significant impact on their approach to the task and hence the quality of the learning outcome. One student who sees an AI tool as a tool for the end will use it one way, a student who sees it as a thinking partner or interlocutor will use it in another. The technology itself will not cause either deep engagement or its absence, and it is neither a characteristic of the technology nor of the software, but of understanding and framing. If so, then perception study is not a prerequisite for investigating the educational impacts of AI, rather it is one of the main components of studying the educational impacts of AI (AIBlooshi, 2026).

Zawacki-Richter et al., (2020), reviewing 146 studies of AI applications in higher education published between 2007 and 2018, reached a conclusion that has become one of the most frequently cited findings in the field: research on AI in higher education was being conducted overwhelmingly by computer scientists rather than educators, and displayed a striking absence of critical reflection on pedagogical and ethical implications. Only a small minority of studies were authored by researchers affiliated with faculties of education. The technical capability of the systems was being developed and evaluated in near-isolation from the educational questions that ought to have governed their design. That observation was made before generative AI reached students directly, and the imbalance it identified has become more consequential since (Xiao et al., 2025).

This study is an answer to that imbalance. It challenges higher education's approach to AI as a technology to be judged, but as a practice to be

comprehended, from the perspective of the people who practice it. It aims to provide insights into how AI is being integrated into student learning, the perceived benefits of AI, concerns about what AI is taking away, and the conditions that must be met for its responsible use, using semi-structured interviews with participants from across the higher education sector. The goal isn't to make a judgment or draw a conclusion on whether AI in education is good or bad, but to create a richer narrative about how AI and optimism are working side by side in the lives of educators who interact with the technology on a daily basis (AlBlooshi, 2026; Alotaibi, 2024; Kuleto et al., 2021).

The study is guided by three research questions:

1. How do participants perceive and describe the contribution of AI tools to student learning in higher education?
2. What challenges, risks, and concerns do participants associate with the use of AI in student learning?
3. What institutional, pedagogical, and ethical conditions do participants regard as necessary for the responsible integration of AI into higher education?

2. Literature Review

There is no one AI in education; it's a family of them, and members have very distinct pedagogical differences. Zawacki-Richter et al. (2020) identify four large categories of applications in higher education: Profiling and Prediction – using past information to profile students at risk or to inform admission and placement decisions; Assessment and Evaluation – automated grading, feedback generation and evaluation of teaching; Adaptive Systems and personalization – adapting content, sequence or difficulty to individual students; Intelligent Tutoring Systems – attempting to replicate components of one-to-one dialogic instruction. The distribution of both risk and benefit is unique in each domain and each situates the learner within a unique relation to the system (Aguado-García et al., 2025).

The introduction of generative AI is not clearly defined in this taxonomy, which is why it has caused such disruption. Large language models can produce natural-language text, code and analysis from natural-language queries, and they

can do this over an effectively infinite variety of tasks. In an early and influential study, Kasneci et al., (2023), identified the opportunities that such models provide for teaching and learning, and the difficulties that they pose for governance, since their breadth of applicability means that they are also more useful and more difficult to manage than the task-specific models that came before them. Unlike an adaptive courseware platform, a general-purpose language model has no pedagogical model embedded in it, no curriculum, and no institutional accountability. Its educational value is determined almost entirely by how it is used.

A further conceptual distinction concerns the locus of control. Institutional AI systems — learning analytics dashboards, plagiarism detection services, adaptive platforms licensed by a university — are adopted by institutions and applied to students. Generative tools are adopted by students and applied to their own work. This inversion has significant implications for policy: an institution can decide whether to deploy a learning analytics system, but it cannot in any practical sense decide whether its students use generative AI. It can only decide how to respond to the fact that they do. Wu et al., (2025), synthesizing ninety-nine studies of generative AI in higher education classrooms, document the resulting variability in institutional response, which ranges from prohibition through cautious permission to active curricular integration, often within the same national system and sometimes within the same university.

The educational case for AI rests on several distinct claims, which are worth separating because the evidence for them is not equally strong. The first and most frequently advanced is personalization. Mass higher education is characterized by a structural mismatch between the individualized attention that learning theory recommends and the resources institutions can supply. Adaptive systems and, more recently, conversational AI are proposed as a partial remedy: a system that never tires, never runs out of office hours, and adjusts its explanation to the specific point at which a particular learner is stuck. (Chan & Hu, 2023), surveying 399 undergraduate and postgraduate students in Hong Kong, found that students recognized personalized learning support as among the most valuable affordances

of generative AI, together with assistance in writing, brainstorming, research, and analysis.

The second is immediate feedback. Feedback is among the most powerful influences on achievement, and it is also among the most expensive to provide. Delay degrades its value: comments returned three weeks after submission arrive when the student has moved on and the cognitive context has dissipated. AI tools can provide response at the moment of composition. Crompton and Burke, (2023) note that this immediacy is one of the clearest educational advantages of AI systems, while also cautioning that feedback which is fast but opaque which tells a student that something is wrong without conveying why in terms they can act on may produce compliance rather than learning (Gasanova & Romanova, 2024).

The third is improved access to information and to explanation. For students working in a second language, for those with disabilities, for those studying at a distance, and for those in institutions with limited library or tutorial provision, an AI system that can explain, summaries, translate, and rephrase lowers a real barrier. Safiulina et al., (2024) emphasize the accessibility gains available to learners who have historically been poorly served by standard instructional formats. This is an argument about equity, and it stands in some tension with the equity concerns discussed below.

The fourth is support for independent learning. The availability of an always-present source of explanation may encourage students to pursue questions they would otherwise abandon, to explore beyond the syllabus, and to take greater responsibility for their own study. Whether this materializes depends heavily on how the tool is used: the same availability that supports exploration also supports avoidance (Bashir & Beydoun, 2026).

The empirical evidence on learning outcomes is accumulating but should be read carefully. Meta-analytic work synthesizing experimental studies of ChatGPT's effects reports improvements in academic performance, in affective and motivational states, and in higher-order thinking propensities, alongside reductions in perceived mental effort. That last finding is double-edged. Reduced mental effort is a benefit when the effort eliminated was extraneous to learning formatting,

searching, mechanical translation and a cost when the effort eliminated was the germane cognitive work through which understanding is constructed. Distinguishing between the two is precisely the pedagogical judgement that automated tools cannot make (Esakkiammal & Kasturi, 2024).

2.1 Challenges, Risks, and Concerns

Academic integrity has dominated institutional discussion, for understandable reasons. Doğan et al., (2025) argue that generative AI destabilises assessment practices that had assumed a reliable correspondence between a submitted artefact and the cognitive work of the student who submitted it. The difficulty is not simply that students may cheat they always could but that the boundary between legitimate assistance and illegitimate substitution has become genuinely unclear. Check grammar; create a full argument, this is not controversial; this vast middle ground that lies between these extremes is subject to norms that most institutions have not yet communicated and students must guess. There is no stable answer for detection: detection tools are not reliable, result in false accusations, and present a disadvantage to second-language writers with prose that may be identified and marked as machine-written (Baig, 2024).

There's a second, subtler issue — that of accuracy and bias. Generative systems create confident and fluent text regardless of the accuracy of the content, and can either cite or fabricate a citation, misattribute a finding or reproduce an error with absolute confidence. In the domain they are studying, they often are not, for expert users, they can be detected, but not for students who by definition are not experts with which to evaluate claims (Sposato, 2025). The models also generate biases that are found in training data, which comes disproportionately from specific languages, regions, and epistemic traditions. For students who are not within these traditions, their contexts may be misrepresented or underrepresented (Zawacki-Richter et al., 2020).

There has been increased focus on overdependence. (Kasneci et al., 2023), reported on the reliance on AI systems and the extent, as well as the costs, of overreliance in an experimental study of trust and reliance, reporting the extent to which users absorbed the information provided by the AI system, even when they had

the information to know that the information output provided by the AI system was incorrect. The worry in an educational context is that the cognitive effort that is being redirected by the AI is typically the same kind of effort that is used to acquire new knowledge. A student who gets a worked solution also has an answer; a student who works towards a solution has also, in the process, constructed the schema which enables him to solve the next problem. Overreliance, that is, shallow learning, in a significant number of the studies examined, has been reported in reviews of AI in education for programming (Wu et al., 2025).

Critical thinking is involved but in a different way. Habitually going to a machine with a new question instead of sitting with uncertainty, forming a tentative position, and testing it against evidence can atrophied as a result of disuse (Chan & Hu, 2023). There is indeed literature on both sides of the argument here. There are some reports that the usage of AI is related to an increase in higher-order skills, presumably due to students' critical engagement with the output of the AI; while other studies find that the excessive adoption of AI-generated content leads to a decrease in higher-order skills. The most likely solution is that both happen, in varying students and with different tasks, and that this is the difference between how AI use is pedagogically structured or not (Wu et al., 2025).

Data privacy raises questions that most institutions have not adequately addressed. Student prompts submitted to commercial AI services may include coursework, personal circumstances, health information disclosed in the course of seeking help, and identifiable details of third parties. These data leave institutional control entirely. UNESCO's guidance Holmes & Miao, (2023) identifies the protection of data privacy as a priority for regulation, observing that the absence of national regulatory frameworks in most countries leaves users unprotected and institutions unprepared to validate the tools their students are using.

Unequal access complicates the equity argument advanced in favour of AI. The most capable models are increasingly available only on subscription; effective use depends on stable connectivity, suitable devices, and critically the

linguistic and technical competence to construct effective prompts and evaluate the resulting output. A digital divide that once concerned access to hardware now concerns access to capability, and the second divide is harder to see and harder to remedy. There is a real risk that a technology defended on grounds of inclusion will in practice widen the gap between well-resourced and poorly-resourced learners (Xiao et al., 2025; Zadorina et al., 2024).

3. Research Methodology

The study adopted a qualitative research design located within an interpretive paradigm. This orientation was appropriate to the research questions, which concern meaning and experience rather than incidence or effect: the study sought to understand how participants make sense of AI in relation to student learning, not to measure how frequently they use it or to estimate its impact on attainment. Interpretive research treats participants as knowledgeable agents whose accounts of their own practice constitute legitimate evidence, and treats the researcher's interpretive work as an acknowledged element of the analysis rather than a source of contamination to be eliminated (Mustofa, 2023).

A qualitative approach was also warranted by the state of the field. Where the relevant constructs are still being defined and the practices under study are changing rapidly, imposing predetermined categories through a structured instrument risks measuring the researcher's assumptions rather than the participants' experience. In situations where the boundaries of a phenomenon are still being clarified, as they are in the current study, semi-structured interviewing allows for an exploration of issues that participants present that the researcher may not have had in mind initially (Coe et al., 2025).

The participants were selected using purposive sampling, where informants with first-hand experiences of AI use in higher education teaching and learning were chosen because they have the best understanding of the phenomenon.

Only semi-structured interviews were used to collect the data. The research questions were used to design an interview guide, which was divided into four sections: Participants' pattern of AI use in learning; perceived educational benefits;

perceived risks, concerns and difficulties; and institutional response, policy and future direction. Open questions with prompts were prepared and left in each section for the researcher to pursue lines of thought brought forward by the participants, whilst elucidating the key areas.

4. Data Analysis

The transcripts of the interviews were thematically analysed, following the six steps recommended by (Braun & Clarke, 2023). This first stage consisted of multiple readings of the entire set of transcripts to familiarize oneself with the data and to make initial analytic impressions. In the second, codes were given systematically to all data; coding was done inductively (codes were not imposed but were derived from the data from participants' own formulations). The third phase involved collating codes into candidate themes and the links between themes were explored. In the fourth phase, candidate themes were examined in relation to the coded extracts and also the whole data set, and any themes that did not have enough supporting data or lacked internal coherence were collapsed, subdivided, or eliminated. In the fifth phase, each theme was refined, given a name and the scope of the theme and its boundaries were outlined.

The aspects of trustworthiness were discussed on the four criteria suggested by (Enworo, 2023). Prior to interviews, all participants were given written materials describing the purpose of the study, the voluntary nature of participation, the purpose of data use, and that they were free to participate, or not, without any consequences. Reporting was anonymous: transcripts have been anonymized and participants are identified throughout using codes; institutions are not named. Recordings and transcripts were kept safely and only available to the research team. There was no monetary or non-monetary incentive to participate.

5. Findings

Thematic analysis resulted in three superordinate themes: several subthemes were part of each. The first is what the participants thought AI could offer in education; the second are the threats and concerns that they highlighted; and the third is what they believed were the conditions for responsible uses. Themes are given one-by-one

here to make them easier to follow, but they are often used in the same answer by the participants, and they come back to this complication in the discussion.

4.1 Theme 1: AI as an enabler of student learning

AI tools were consistently described as saving effort and time in finding and pulling information together. What was highlighted was not only speed, but the breaking down of a barrier at the first point of contact with an unknown topic; a place where the difficulty of getting started had previously deterred inquiry. Multiple participants compared AI to traditional search, saying AI gives a summary instead of a list of sources to assess and reconcile — a step that participants appreciated and, as will be explored later, distrusted.

One of the most valuable features of the AI tools was found to be the ability to respond immediately, according to the folks. The folks identified the ability to respond immediately as one of the most valuable features of the AI tools. Participants found the timing of the feedback and feedback and the work to be important: feedback received when the work is in its draft stage can be taken on board, but comments given after the work has been submitted are likely to be much less effective. For some, AI was used as a way to get a "first draft" and then have teachers read it, rather than to replace teachers.

Participants felt that AI adapts to their individual challenge in a way that generalized instruction does not. One theme that emerged repeatedly in these stories was that participants felt they could get an AI system to explain something, then ask it to explain something they thought they should already know, without feeling shame for being in a room with a teacher or in a classroom where they can't do it themselves. This affective dimension of the affordance was raised repeatedly and appears to be an important part of its perceived value.

Related but distinct was the use of AI to obtain alternative explanations of material that had not been understood at first presentation. Participants described requesting analogies, simplified accounts, worked examples, or explanations pitched at a different level, and reported that the ability to iterate — to say that an explanation had not helped and request another — was central to

the benefit.

Participants described AI as extending their capacity to pursue questions without waiting for scheduled contact with teachers, and some framed this in terms of increased autonomy and confidence. This subtheme was, however, the one most frequently qualified by participants themselves, several of whom noted in the same breath that the line between working independently and having the work done for them was not always easy to locate.

4.2 Theme 2: Risks, concerns, and unintended consequences

Participants expressed substantial concern about the reliability of AI-generated content, describing instances of confident error. What made this concerning to participants was specifically the mismatch between the confidence of the presentation and the reliability of the content: errors delivered fluently and without hedging are difficult to detect, particularly for a learner who lacks the domain expertise to evaluate the claim. The small scope of AI outputs and lack of representation of their context were also pointed out by some participants.

Participants expressed a worry that relying on AI continuously will take away the cognitive efforts that lead to understanding. The concern was not so much a shortfall as a gradual erosion that would result from a consistent stream of use over time, over the course of a degree program. Multiple participants reported that they observed this pattern in themselves and made an effort to avoid it, indicating that this concern is not just an outside-one's-bubble thing but an inside-one's-bubble thing.

Participants described considerable uncertainty about where legitimate assistance ends. Notably, the concern was frequently expressed not as a worry about deliberate cheating but as anxiety about inadvertent transgression: participants reported not knowing whether particular uses of AI were permitted, and reported that the absence of clear institutional guidance left them to make such judgements alone. Participants also raised questions about the fairness of assessment when students differ in their willingness to use AI and in their access to more capable tools.

Beyond assessment, participants raised concerns

about the handling of personal and institutional data submitted to AI services, about the use of AI to generate content presented as original scholarship, and about the opacity of the systems themselves — the impossibility of knowing on what basis an output was produced or what happens to the material submitted.

Participants described widespread gaps in the capability required to use AI well: in constructing effective prompts, in recognizing the limits of the systems, in verifying outputs, and in understanding what the technology is and is not doing. Some participants pointed out that the assumption that there was a generational fluency with digital tools was incorrect, and that people may be competent with digital devices but not with AI.

Doing things because of AI, not in spite of it, was the most commonly stated worry of being dependent on AI and relying on it as the first choice for all academic tasks. Participants described this as a gradual and largely unnoticed shift, and some connected it explicitly to a perceived decline in their own tolerance for difficulty and willingness to persist with a problem before seeking help.

4.3 Theme 3: Conditions for responsible integration

Participants were near-unanimous in identifying the absence of clear institutional policy as a problem, and equally clear that prohibition was neither desirable nor enforceable. What was sought was specificity: guidance identifying which uses are permitted in which contexts, how AI use should be disclosed, and how assessment will account for it. Participants noted that inconsistency between teachers within the same institution — one prohibiting what another required — was a particular source of confusion.

Participants argued that teachers require both preparation and authority in this area: preparation because many teachers were described as uncertain about the technology themselves, and authority because participants looked to teachers rather than to policy documents for the practical judgements that govern day-to-day use. Some participants observed that teachers who used AI openly and modelled critical engagement with its output were more useful than those who either

prohibited it or ignored it.

Participants advocated explicit and structured instruction in AI use rather than reliance on informal learning. The training anticipated was to include aspects of the evaluative capacities described above, in addition to operational skill, that is, verifying outputs, recognizing bias, understanding limitations, and making principled judgements about appropriate use.

Participants urged the need for mechanisms on how to deal with data privacy, transparency in disclosure, accountability for work done with AI tools, and institutions' accountability in choosing and endorsing tools as a framework, and several participants suggested that this process should be done together with students rather than imposed on them.

Participants noted a gap between access to more capable systems and access to them for students who cannot afford subscription access, and also a division between access to connectivity, devices and linguistic confidence to use more capable systems for students who do not have access to those systems. Others said, "It's no longer a matter of convenience, it's a question of equity if institutions are providing AI tools.

5. Discussion

The results form a picture which is neither the positive one that the literature has largely proceeded from, nor the negative one, but a third, that makes both positive and negative aspects structurally linked (Batista et al., 2024). What is most interesting about the accounts of the participants is that the affordances and the risks are not two distinct phenomena that are "competing" against each other, but are often the same phenomenon discussed from different perspectives. Immediate feedback is good because it eliminates delay and is bad because it eliminates the time a student could work through the problem on her own. The ability to access synthesized information is useful because it makes it easier to get started, but harmful because it is a way to avoid the exercise of judgment that requires reconciling sources. Inspecting the situation, support for independent learning and excessive dependence are two affordances in different contexts of use. Any policy that makes separation between benefits and harms, allowing benefits and

banning harm does not describe what it regulates (Kuleto et al., 2021).

This helps to understand why the empirical research on AI and higher-order thinking is conflicted, with some finding benefits and others harm. If the affordance is indeed bivalent, then it won't matter how the technology is used; it will depend upon the conditions of use: task design, framing provided by the teacher, prior competence of the student, pressures of use. This aligns with the pedagogical leverage suggested by Chan and Hu (2023) that student perception influences student learning approaches and hence learning outcomes. Students do not need to be asked if they should use AI, but when they should use AI in such a way that the positive aspects are more salient than the negative (Alotaibi, 2024).

The accuracy and bias results are in line with the concerns raised by Kasneci et al. (2023); however, the focus of the participants brings a value to the results. What disturbed participants was not the error, but the confident way in which the error was presented, that is, its absence of the hedging and uncertainty markers that normally highlight the need to verify a claim in a reader. This is a type of epistemic problem that this type of tool systematically eliminates the cues through which learners fine-tune their level of trust. It is obvious that AI literacy education should not be delivered through the prism of 'AI gets it wrong sometimes', as students are aware of this. It has to develop the substantive skill of verification from other sources, which is much more difficult to teach (AlBlooshi, 2026).

The academic integrity findings stand out in an informative manner from the typical manner in which the issue is discussed. The scholarly and institutional discourse about integrity has focused on detection and deterrence, how to spot AI-generated work and how to dissuade its creation. In this research it was much more a question of "what can I do? or "what is allowed?": the participants were uncertain about what was allowed and this uncertainty was an area of anxiety rather than opportunity. This is closer to the position taken by Sullivan, Kelly and McLaughlan (2023), who argue that the challenge is one of assessment redesign and norm articulation rather than surveillance. If students' dominant experience is confusion rather than

temptation, then clarity is a more effective intervention than detection, and considerably cheaper (Aguado-García et al., 2025).

The concern about dependence corresponds to what Klingbeil, Grützner and Schreck (2024) demonstrate experimentally: that users defer to AI outputs even when equipped to detect their errors. The participants' contribution is phenomenological. They described dependence as gradual, unremarked, and recognised only in retrospect a drift rather than a decision. This has a direct implication for intervention design. If dependence were a choice, information about its costs might be sufficient to counter it. If it is a drift, then it must be countered structurally, through task and assessment designs that require the student to demonstrate reasoning rather than to produce artefacts, and that make the process of work visible alongside its product (Zadorina et al., 2024).

The prominence of AI literacy in participants' accounts of what is needed supports the case made by Kong, Cheung and Zhang (2021) and formalised in UNESCO's competency framework (Miao & Shiohira, 2024). Participants' explicit rejection of the assumption that younger students are natively competent with AI is worth noting: fluency with digital interfaces is not the same as understanding what a language model is doing, and treating the two as equivalent leads institutions to omit instruction that is genuinely needed (Xiao et al., 2025).

Finally, the equity findings complicate the accessibility argument that has featured prominently in the case for AI in education. The claim that AI democratises access to explanation is not false — participants' accounts of asking questions without social cost provide real support for it — but it is conditional on access to capable tools and on the competence to use them. Where the most capable systems sit behind subscription and effective use depends on linguistic and technical confidence, a technology defended on inclusion grounds may in practice stratify. This is the concern raised in UNESCO's guidance (Miao & Holmes, 2023) regarding institutional unpreparedness, and it deserves more attention than it has received in a literature substantially produced in well-resourced settings (Rangavittal, 2024).

Taken together, these findings support a position that might be summarised as follows: AI is educationally valuable in proportion to the deliberateness with which its use is structured, and educationally harmful in proportion to the absence of such structure. This places the burden squarely on institutions and teachers rather than on students, and it identifies the central failure mode not as misuse by students but as abdication by institutions that have left students to work out the norms of a consequential new practice on their own (Crompton & Burke, 2023; Gasanova & Romanova, 2024; Jafari & Keykha, 2024)

6. Implications

Teachers require professional development in AI both as users and as designers of learning. The results indicate three priorities for designing assessment: Making reasoning visible, not just products, as tasks that require students to reason are more vulnerable to substitution; modelling critical engagement with AI output in class, including deliberate examination of AI errors, which aids in teaching verification better than warnings; and stating expectations for AI use explicitly at the level of the individual assignment, where students actually look for guidance.

7. Limitations

The results are limited by a few restrictions. The study is qualitative and context specific, more than statistical generalization; the intent was analytic, and the results are meant to enlighten a set of experiences, not provide anything like a representative sample of a larger population. The study relied entirely on data obtained from interviews, thus reporting what participants said about their practice, not the practice itself, and self-report in this area may be influenced by social desirability, as with the area of academic integrity, for which participants may underestimate uses of the practice that they contest. This was a purposive sample drawn on the basis of AI experience and thus does not include the viewpoints of non-users who would contribute independent information to the reasons they do not use AI. The study is cross-sectional and reflects perceptions at one time, in a rapidly changing technological environment; the capabilities of the tools that are discussed will have changed by the time the article is published. Lastly, the process of thematic analysis is an interpretive process, and although the

trustworthiness procedures outlined above were adhered to, another theme might have structured the data differently.

8. Conclusion

This study aimed to explore the perceptions and experiences of higher education students with AI in their learning. The results show that AI tools are perceived to enhance information accessibility, offer instant feedback, facilitate personalized learning, assist in understanding complex content and facilitate self-guided learning. They also reflect a significant worry concerning the accuracy and bias of the generated output, the loss of critical thinking, the doubts related to academic integrity, the ethical use of AI, the lack of AI knowledge and skills literacy, and the increasingly growing dependence on AI-generated content. Participants stated that the conditions that may

lead to responsible integration include institutional policy, teacher guidance, student training, ethical frameworks, and equitable access.

The main thrust is that these benefits and risks are not separate magnitudes to be traded off against each other but are different uses of the same affordances. It is not the technology itself, but the pedagogical and institutional framework around it that determines which is realized. In some cases, AI can contribute significantly to the learning process when it acts as an aid to human logic, human interaction with the teacher, and student engagement, but when it replaces them, it can undermine learning. The distinction between the outcomes lies with the institutional design and pedagogical judgement, and falls squarely on the shoulders of universities and teachers and not the students who are the ones currently figuring it out.

Conflict of Interest

The authors showed no conflict of interest.

Funding

The authors did not mention any funding for this research.

References

- Aguado-García, J.-M., Alonso-Munoz, S., & De-Pablos-Heredero, C. (2025). Using artificial intelligence for higher education: An overview and future research avenues. *Sage Open*, 15(2), 21582440251340350.
- AlBlooshi, S. (2026). Artificial intelligence in higher education, opportunities, and challenges: A review. *Frontiers in Education*, 10, 1683968.
- Alotaibi, N. S. (2024). The impact of AI and LMS integration on the future of higher education: Opportunities, challenges, and strategies for transformation. *Sustainability*, 16(23), 10357.
- Baig, M. A. A. (2024). *The Future of AI in Higher Education: Opportunities & Challenges Ahead: AI and Higher Education*.
- Bashir, S., & Beydoun, A. R. (2026). Transforming Higher Education Using Artificial Intelligence: Innovations, Future Directions, and Risks. In *Higher Education in the Arab World: Artificial Intelligence* (pp. 143–168). Springer.
- Batista, J., Mesquita, A., & Carnaz, G. (2024). Generative AI and higher education: Trends, challenges, and future directions from a systematic literature review. *Information*, 15(11), 676.
- Braun, V., & Clarke, V. (2023). Toward good practice in thematic analysis: Avoiding common problems and being a knowing researcher. *International Journal of Transgender Health*, 24(1), 1–6.
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(1), 43.
- Coe, R., Waring, M., Hedges, L. V., & Ashley, L. D. (2025). *Research methods and methodologies in education*. SAGE Publications Limited.
- Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education*, 20(1), 22.
- Doğan, O., Demir, M. C., & Yaşar, H. (2025). Artificial intelligence in higher education management: Opportunities, challenges, and future perspectives. *Selçuk Üniversitesi Akşehir Meslek Yüksekokulu Sosyal Bilimler Dergisi*, 20, 227–243.
- Enworo, O. C. (2023). Application of Guba and Lincoln's parallel criteria to assess trustworthiness of qualitative research on indigenous social protection systems. *Qualitative Research Journal*, 23(4), 372–384.
- Esakkiammal, S., & Kasturi, K. (2024). Advancing Educational Outcomes with Artificial Intelligence: Challenges, Opportunities, And Future Directions. *International Journal of Computational and Experimental Science and Engineering*, 10(4), 1749–1756.
- Gasanova, R. R., & Romanova, E. A. (2024). Artificial intelligence in higher education: problems, opportunities, risks. *RUDN Journal of Informatization in Education*, 21(4), 501–515.
- Holmes, W., & Miao, F. (2023). *Guidance for generative AI in education and research*. Unesco Publishing.
- Jafari, F., & Keykha, A. (2024). Identifying the opportunities and challenges of artificial intelligence in higher education: a qualitative study. *Journal of Applied Research in Higher Education*, 16(4), 1228–1245.
- Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., & Hüllermeier, E. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274.

- Kuleto, V., Ilić, M., Dumangiu, M., Ranković, M., Martins, O. M. D., Păun, D., & Mihoreanu, L. (2021). Exploring opportunities and challenges of artificial intelligence and machine learning in higher education institutions. *Sustainability*, 13(18), 10424.
- Mustofa, M. (2023). Epistemological paradigms in positivism, interpretivism, and action research in educational research: A literature review. *Journal of Office Administration: Education and Practice*, 3(3), 214–224.
- Rangavittal, P. B. (2024). Transforming higher education with artificial intelligence—Benefits, challenges, and future directions. *International Journal of Science and Research (IJSR)*, 13(5), 1635–1642.
- Safiulina, E., Labanova, O., Uukkivi, A., Petjärv, B., & Vilms, M. (2024). Integrating artificial intelligence in higher education: A literature review of current trends, challenges, and future directions. *ICERI2024 Proceedings*, 1666–1675.
- Sposato, M. (2025). Artificial intelligence in educational leadership: a comprehensive taxonomy and future directions. *International Journal of Educational Technology in Higher Education*, 22(1), 20.
- Wu, F., He, L., Liu, C., Wang, W., Li, R., & Dang, S. (2025). *Changing trend of serum tumor necrosis factor-alpha level during pegylated interferon-alpha treatment in inactive HBsAg carriers and its association with HBsAg clearance*.
- Xiao, N., Pei, Y., Yuan, C., Bu, Y., & Cai, Z. (2025). *Transforming education with artificial intelligence: A comprehensive review of applications, challenges, and future directions*.
- Zadorina, O., Hurskaya, V., Sobolyeva, S., Grekova, L., & Vasylyuk-Zaitseva, S. (2024). The role of artificial intelligence in creation of future education: Possibilities and challenges. *Futurity Education*, 4(2), 163–185.
- Zawacki-Richter, O., Marin, V., Bond, M., & Gouverneur, F. (2020). Einsatzmöglichkeiten Künstlicher Intelligenz in der Hochschulbildung—Ausgewählte Ergebnisse eines Systematic Review. *Digitale Bildung Und Künstliche Intelligenz in Deutschland: Nachhaltige Wettbewerbsfähigkeit Und Zukunftsagenda*, 501–517.