

Pedagogical Equipose: Navigating The Affordances, Risks, and Institutional Readiness of Generative and Agentic Artificial Intelligence in Higher Education

Clinton Mokaya Nyakongo

University of Eldoret

DOI: <https://doi.org/10.47772/IJRISS.2026.100900135>

Received: 10 September 2026; Accepted: 15 September 2026; Published: 03 October 2026

ABSTRACT

Generative and agentic artificial intelligence (AI) tools are reshaping teaching, learning, assessment, and institutional practice in higher education, yet decisions about which tools to adopt and how to use them remain largely ad hoc. This paper introduces pedagogical equipose as an organizing concept for evaluating AI tools: a deliberate, evidence-informed and provisional balance between the instructional benefits a tool affords and the cognitive, ethical, and institutional risks it introduces, rather than a default assumption that a tool is either beneficial or harmful. Using a structured qualitative document analysis of 27 widely used generative and agentic AI tools, spanning conversational, visual/video, audio, and agentic/development categories, the study examines pedagogical affordances and constraints alongside institutional readiness. The analysis draws on Constructivist Learning Theory, Mayer's Cognitive Theory of Multimedia Learning, and Critical Digital Pedagogy. The paper is grounded in a thematic synthesis of peer-reviewed and authoritative grey literature published between 2020 and 2025, following Braun and Clarke's (2006) six-phase approach to thematic analysis. Four interconnected themes emerged: conversational AI as a pedagogical partner; multimodal AI as a resource for representation and expression; agentic AI for workflow automation; and an overarching tension between efficiency and cognitive engagement. The paper concludes by synthesizing these themes and their theoretical grounding into a decision framework intended to help institutions and educators operationalize pedagogical equipose in policy, professional development, and everyday practice.

Keywords: Artificial intelligence; generative AI; higher education; pedagogy; AI literacy; multimodal learning; educational technology; pedagogical equipose.

INTRODUCTION

The presence of generative and agentic AI in higher education is no longer a matter of speculation. These artificial intelligence tools now produce written content, images, video, and audio and summarize source material. In their agentic form, they chain several of these capabilities into multi-step workflows. Such developments create new instructional affordances. Nevertheless, they also confront institutions and educators with a practical and recurring choice: which tools to use, for which purposes, and under what conditions they support rather than undermine learning.

This paper develops the idea of pedagogical equipose as the frame of reference. Pedagogical equipose is neither a zero-sum balance, nor a simple decision rule, but an ongoing, discipline-specific, pedagogical judgement regarding the ability of AI to support productive struggle, teacher judgement, assessment validity, and student agency. It goes beyond responsible AI integration, human-centred AI, critical AI literacy, and technology-pedagogy alignment by bringing their questions, tensions, and considerations under one decision footing for educators and educational institutions. Accordingly, the research questions below are not merely about what AI can do, but what it should be allowed to do in a given learning situation.

The issue the present paper attempts to address is pedagogical, not technical. An AI tool may support learning when a student questions, tests, compares, and modifies an explanation derived from the AI but the same tool

may decrease learning if it eliminates reading, thinking, or revising from learning activity. These considerations align with critical approaches to the educational technologies (Selwyn, 2022) that frame tools as socially constructed technologies embedded within an institution's socio-technical milieu.

This paper was guided by the research questions outlined below:

1. What pedagogical affordances are posed by the new array of generative and agentic AIs and how should they be weighed through pedagogical equipoise?
2. What are the risks associated with these affordances cognitive engagement, equity, privacy, and academic integrity and how does equipoise inform their mitigation?
3. How will the roles of the educator and the readiness of institutions be reconfigured by an equipoise-oriented approach?

The paper accentuates the new and emerging array of AI tools under one pedagogical lens. The paper views the new tools as a collection of technological features. It answers questions such as; what educators is doing, what forms of learning a machine mediated-tool might produce, and what a favorable environment is for AI use. The present analysis centers on four broad themes and further guides its implications for teaching practice, institutional policies, and future research endeavors.

LITERATURE REVIEW AND THEORETICAL FRAMEWORK

Constructivist Learning Perspectives

Constructivist theories view learning as an activity in which individuals construct and modify knowledge through experience and interaction, conversation, and reflection. For a constructivist- perspective, the "utility of an intelligent system is in a scaffolding or cultural tool function, not as a replacement for the learner's efforts Collins,(1996). For example, conversational tutors can produce different explanations and questions for the students, however, students are to analyze evidence, recognize incorrect responses and build their own constructions of understanding. Additionally, social constructivism specifically illustrates, the importance of interaction and teamwork in learning. By assisting students in discussing potential solutions, coming up with project ideas, simulating expert feedback, or offering various viewpoints on an issue, AI tools can promote collaborative learning. However, since AI-generated information may contain mistakes, bias, or misleading explanations, meaningful interaction with peers and trainers is still crucial.

Generative AI reduces the efforts in response generation; it does not necessarily reduce the cognitive effort required in generating understanding of the response. The educational challenge will then be constructing activities for learners in which AI outputs become material for investigation and not an endpoint for the lesson.

Cognitive Theory of Multimedia Learning

Mayer's (2021) Cognitive Theory of Multimedia learning is valuable when considering the generation of images, videos, and audio. The theory focuses on how our students possess limited cognitive channels through which learning information is filtered and also the importance of not placing an excessive cognitive load upon the learner. Generating multiple modes of media through AI may effectively promote understanding of abstraction, alternative depictions or learners who require additional audio or visual input to assist learning; however, more media is not always better. Excessive or poorly created multimedia can compete for limited cognitive resources with other presented media. Therefore, it is with the objective of enhancing or understanding of the subject rather than visual novelty, which the worth of the generated multi-media should be determined by, not its aesthetic appeal.

Critical Digital Pedagogy and AI Literacy

Selwyn (2022) presents the argument for critical use of educational technologies in that students should think critically about power, value, institutions and interests surrounding the technologies they use. As far as AI is concerned, the researcher notes that this means critical and inquisitive thinking by teachers and learners: Who

built this? What kind of data and assumption have fed it, and what kind of system do they have; the needs of whom does it serve, and who is ultimately to blame if things go wrong?" AI literacy consequently does not just mean knowing how to write a prompt but also, how to check a system's response, check the reliability of AI-generated sources, detect bias in AI's output and reasoning, be aware of limitations, secure your data and privacy, and know when not to use AI.

Theoretical Integration

The three frameworks reinforce one another, the first addresses active construction of knowledge, the second addresses processes and the third addresses power values and the use of technology responsibly. They therefore provide criteria for evaluating AI devices based upon the learning processes not their technological novelty. Each of the below observations directly relates back to theory.

AI availability does not by itself produce better learning. Effective AI usage occurs when AI is pedagogically integrated to support active knowledge construction, cognitively appropriate information processing, and critical, ethical and responsible engagement.

Pedagogical Equipose: Definition and Differentiation

Pedagogical equipose is a contextually sensitive professional judgement where teachers compare the affordances of generative and agentic AI to the cognitive, relational, ethical, equity, and institutional risks of using it. It is neither a universal law, nor a neutral copasetic ratio, nor a technical cost/benefit analysis. It is an ongoing stance to assessing whether the use of AI sustains productive struggles, judgements, assessments, and agency for students. The stance of pedagogical equipose contrasts with approaches to responsible AI which prioritise ethics, fairness, transparency, and accountability. It describes the impact of AI on various aspects of education; with a human-centred AI approach which prioritises human agency and wellbeing, but where the pedagogical choice is left implicit. Through critical AI literacy it offers questions about power, data, and biases, and which illuminates where to look for inequity and bias; and with technology-pedagogy alignment which help teachers ensure that tools align with learning aims and how they are used, but which do not ask whether to use a particular tool, or how. Pedagogical equipose sits within and combines these approaches to arrive at a practical stance for teachers and institutions. Pedagogical equipose is deliberately narrower than several adjacent concepts with which it is sometimes conflated, Responsible AI integration is primarily a governance and compliance frame, concerned with bias mitigation, safety, and institutional policy compliance. Pedagogical equipose is instead an epistemic stance held at the point of instructional decision-making, about how confident one can be that a specific use of a tool will help or harm a specific learning task.

Human-centred AI (Shneiderman, 2020) is a design-level framework concerned with how AI systems and interfaces should be built to preserve human agency and control. Pedagogical equipose is not a design principle for building tools; it is an evaluative posture educators and institutions adopt toward tools that already exist. Critical AI literacy describes a set of competencies that students and teachers should hold, such as interrogating data provenance and power relations (Selwyn, 2022). Pedagogical equipose operates one level up, as a decision-orientation concept: critical AI literacy supplies the skills that make a well-informed equipose judgment possible, but the two are not synonymous. Technology-pedagogy(-content) alignment, most fully developed as TPACK (Mishra & Koehler, 2006), addresses the fit between a technology, a pedagogical approach, and content knowledge. It is largely a static integration heuristic. Pedagogical equipose adds an explicitly temporal and uncertainty-holding dimension: it is a stance of withheld, revisable judgment pending evidence, rather than a model of a good technology-pedagogy-content fit at a single point in time.

METHODOLOGY

Research Design

The paper utilizes qualitative document analysis, specifically thematic analysis. The unit of analysis is defined as the pedagogical potential of an AI tool depicted in the literature and manifested in the recorded tool

affordances. Interpretation is the main paradigm, not experimentation: this paper did not directly test the 27 tools in a student environment and cannot assert that their pedagogical potential has a direct causal link to student learning outcomes.

Search Strategy and Evidence Selection

The review was carried out using a documented search and evidence-selection method; searches were conducted in Scopus, Web of Science, ERIC, IEEE Xplore, ACM Digital Library, Google Scholar, and selected institutional and think-tank repositories with the string ("generative AI" OR "large language model" OR "agentic AI" OR ChatGPT) AND ("higher education" OR university OR college) AND (pedagog* OR assessment OR integrity OR "AI literacy" OR accessibility). The search was constrained to sources from 2020-2025 in English, but earlier theoretical work prior to 2020 was considered for background only. Screening proceeded through title/abstract review, full-text assessment, and backward/forward citation checking. Because this study employed a purposive, interpretive document analysis rather than a systematic review, a formal PRISMA flow diagram was not produced; instead, the inclusion criteria and coding matrix (Section 3.5) document the evidence-selection process transparently.

Data Source and Tool Selection

The tool selection was based on three factors: First, direct applicability and pedagogical resonance for higher education teaching, learning, assessment, content production, accessibility, and workflow. Secondly, distinctness of the AI's function in representation of specific emergent AI functionality, and finally availability of enough publicly available documentation and discussion to formulate and defend its pedagogical interpretation and use. Tools were excluded when there was insufficient documented pedagogical value. The paper focused on its teaching, learning, assessment, content creation, and accessibility uses; AI platforms change quickly and the sample inventory should only be viewed as a temporal snapshot. Table 1 documents each tool, its classification, and its pedagogical purpose.

The study drew its analytical data from a working inventory of 27 representative artificial intelligence (AI) tools selected to reflect the range of AI applications currently relevant to higher learning. The tools were purposively grouped into four broad categories according to their primary functionality: (i) conversational AI tools, (ii) visual and video generation AI tools, (iii) audio and voice AI tools, and (iv) agentic and developmental AI tools.

The conversational AI category comprised ChatGPT, Claude, Gemini, Perplexity, DeepSeek, Copilot, Grok, Meta AI, and Kimi. These tools primarily support text-based interaction, information retrieval, content generation, explanation, brainstorming, and other conversational learning activities.

The visual and video generation category comprised Leonardo AI, NanoBanana, Canva AI, Veo 3, Sora 2, Kling, Runway, Pika, and Pictory. These tools support the generation, editing, and transformation of visual and video-based learning content.

The audio and voice AI category comprised ElevenLabs, HeyGen, Murf AI, LOVO AI, Synthesia, Descript, and AssemblyAI. These tools provide capabilities related to speech synthesis, voice generation, audio and video production, transcription, and related multimedia applications. The agentic and developmental AI category comprised Lovable and Manus, which support AI-assisted software development, task automation, workflow execution, and agentic interaction. Accordingly, the complete working inventory consisted of the following 27 AI tools: ChatGPT, Claude, Gemini, Perplexity, DeepSeek, Copilot, Grok, Meta AI, Kimi, Leonardo AI, NanoBanana, Canva AI, Veo 3, Sora 2, Kling, Runway, Pika, ElevenLabs, HeyGen, Pictory, Murf AI, LOVO AI, Synthesia, Descript, AssemblyAI, Lovable, and Manus.

The categorization was intended to provide a representative rather than exhaustive coverage of AI applications relevant to higher learning. The classification was based primarily on each tool's dominant functionality, recognizing that some AI platforms have capabilities spanning more than one category.

Table 1: AI Tool Inventory, Classification, and Pedagogical Purpose

Tool	Classification	Purpose/ Use in Education
ChatGPT	Conversational (LLM)	AI All-purpose tutor, lesson planner, quiz generator, and brainstorming partner.
Claude	Conversational (LLM)	AI Analyzes long documents (entire syllabi, research papers, or novels) for quick summarization.
Gemini	Conversational (LLM)	AI Multimodal research—analyzes text, images, and YouTube videos together for rich lessons.
Perplexity	Conversational (LLM)	AI Academic research with real-time web search and cited sources—teaches proper fact-checking.
Deepseek	Conversational (LLM)	AI Handles extremely long contexts—can support summarizing entire textbooks or lecture transcripts.
Copilot	Conversational (LLM)	AI Integrates with Microsoft 365 to draft emails, rubrics, lesson slides, and parent communications.
Grok	Conversational (LLM)	AI Engages students in Socratic debate and critical thinking by challenging their arguments.
Meta AI	Conversational (LLM)	AI Great for social studies and media literacy; integrates with WhatsApp/Instagram for casual Q&A.
Kimi	Conversational (LLM)	AI Processes massive PDFs and research papers quickly—ideal for graduate-level literature reviews.
Leonardo AI	Image Generation	Creates detailed historical scenes, scientific diagrams, or fantasy settings for visual learning.
NanoBanana	Image Generation	Quick, simple image generation—great for younger students or fast visual aids during lectures.
Canva AI	Image & Design	The go-to for making infographics, posters, interactive slides, and classroom visuals with zero design skills.
Veo 3	Video Generation	Produces high-quality, cinematic videos—can support student documentaries and film projects.
Sora 2	Video Generation	Turns text prompts into narrative videos—excellent for storytelling and creative writing assignments.
Kling	Video Generation	Creates realistic video clips—useful for science simulations (e.g., biology processes, chemical reactions).
Runway	Video Generation & Editing	Professional-grade AI video editing—ideal for media arts classes and advanced student projects.
Pika	Video Generation	Quick, fun animation generation—great for younger students or rapid prototyping of ideas.
ElevenLabs	Audio / Voice Synthesis	Produces the highest-quality, natural-sounding voiceovers for presentations and podcasts.
HeyGen	Audio / Avatar Video	Creates an AI avatar teacher that delivers pre-recorded lessons in multiple languages.
Pictory	Audio / Video Editing	Automatically extracts short, shareable highlight clips from long lecture recordings.
Murf AI	Audio / Voice Synthesis	Multi-language voiceovers—excellent for ESL, foreign language classes, and dubbing.
Lovo AI	Audio / Voice Synthesis	Voice synthesis with emotional tone—brings literary character dialogues to life.

Synthesia	Audio / Avatar Video	Professional AI avatar videos—can support flipped classrooms, online courses, and substitute plans.
Descript	Audio / Video Editing	Edits audio/video by simply editing a text transcript—ideal for teacher podcasts or video feedback.
Assembly AI	Audio / Speech-to-Text	Provides real-time captions, transcribes lectures, and analyzes student pronunciation for language learning.
Lovable	Development Builders	Builds custom educational apps (quiz games, flashcard tools, classroom polls) without coding.
Manus	Development Automation	General-purpose AI agent that automates multi-step tasks like grading, file organization, or data entry.

Evidence Inclusion Criteria

The initial stage involved both the searching for and selecting of literature. Preference was given to sources that had undergone the peer review process: these included research in scholarly journals and conferences, books and authoritative publications from institutions and research bodies published in 2020 to 2025. Earlier conceptual and theoretical works were included when their concepts were used in the analysis such as Braun and Clarke's (2006) thematic analysis and Popenici and Kerr's (2017) analysis of AI within higher education. Relevant articles covered educational implementation, processes of learning, assessment, academic integrity, accessibility, ethical implications and the implications for an institution.

Coding and Thematic Analysis

Data were analyzed using an organized thematic approach in line with Braun & Clarke (2006): Familiarisation of literature & list of tools, Initial coding, Organization of codes into collated areas of relevance, Developing and refining candidate themes, defining and naming themes, Derivation of analysis/findings from themes. The majority of the coding approach was deductive as the coding was influenced by the questions of the research and the existing theoretical context, however, inductive codes remained when repeated issues came up which were not identified by the initial framework. A coding matrix was employed which documented information on each tool, including its primary pedagogical role, probable learner benefits, risks or limitations of the tool, underlying conceptual thinking, and supporting literature. Themes were also cross-examined in relation to the coded data so as to reduce duplication and also to ensure themes represent relevant issues reflected in more than one of the tools. Since this was a single-reviewer qualitative synthesis statistics for inter-rater reliability could not be computed. Claims are limited to those reflected in the selected literature and are clearly articulated as analytical claims rather than empirical findings, in order to preserve credibility.

Analytical Categories

The ultimate categories were divided based on pedagogical role instead of technological type as instructional practices tend to take into account how technologies will support instructional tasks. Category division separates the type of function of an instrument into; those used in dialogue and explanation; those used for building multi-modal learning artifacts; and those used for facilitating or automating tasks and processes. Such division based on the form function also helps inter category comparisons based on typical concerns: engagement, privacy, accessibility and teacher roles.

Findings: Thematic Analysis of AI Affordances

Each theme is examined through the lens of pedagogical equipoise, asking whether the identified affordance preserves or displaces productive learning effort.

Theme 1: Conversational AI as a Pedagogical Partner

Conversational AI can aid the higher education curriculum as explanation tools, to demonstrate concepts via

examples and to ask clarifying questions and provide feedback and brainstorm and discourse with the user. There are several interactive AI systems available, including ChatGPT, Gemini, Claude and Perplexity, and students have the capability to continuously ask follow-up questions and ask for other kinds of explanations when they interact with these programs. The best possibly application of technology like that will be in an interactive iterative reply system.

However, this exact affordance can transform into a learning hindrance if AI is used to outsource cognition instead of enhancing it. Research has addressed the usage of generative AI and its implications on academic integrity, its reliability, and potential to outsourcing work on AI agents (Dwivedi et al., 2023; Vargas-Murillo et al., 2023). Rather than banning it entirely, rather, pedagogy could employ tasks which require students to interact with the AI, to ask questions about and possibly challenge or correct its responses and supporting evidence, to query for how claims in AI answers were supported, to provide documented citation to its interactions with AI agents.

Perhaps one constructive place that this technology has is in an interactive, reiterative response framework. The student is able to utilize what the AI gives him as input into an ongoing discussion-however, it will be incumbent upon the student to refute and revise the input. Adopting a critical AI literacy framework, students can use the agent and its response as input that needs investigation of authority of source and evidence of support; we should adopt an stance where AI output is just another participant in the learning dialogue for which ultimately the student is ultimately responsible for informed judgment. In equipoise terms, conversational AI is pedagogically warranted only when students are required to question, verify, and revise AI output rather than accept it as an endpoint.

Theme 2: Multimodal AI as a Resource for Representation and Expression

The ability to generate image, video, audio and avatars broaden how academic ideas are presented. Image/diagram generation from Leonardo AI and Canva AI will facilitate producing charts, posters, artwork or presentation media. Such videos creation tools are widely applied for the purpose of scenario building and demonstrations, or stories and fantastics productions (For example Veo 3, Sora 2, Kling, Runway, and Pika). In the meantime, such tools like text/voice generations (ElevenLabs, HeyGen, Synthesia, Murf AI, Descript) might be used by instructor to provide the voiceover narration, practicing language, transcript and/or for representation of the content, if the instructor does not wish to implement other ways of presentation.

All these tools have an effect on teaching which depends on the purpose of teaching. Visualizations and Other multimodal representations help student to have a grasp over something abstract or difficult concept, but there is no connection between the quality of the presentation and the quality of learning (using high fidelity presentation doesn't necessarily means having learning of high quality). Their framework of multimedia learning (Mayer, 2021) suggests that by providing with too many images, animations or narrative forms, the presentation will cause cognitive overdose, as human capability of consuming information is of a limited supply.

With this, AI-generated materials can benefit learning by removing certain impediments or by facilitating learners with additional helpful illustrations. Nevertheless, the excessive amount of varying visual form should not be provided due to possible overwhelming or overload to the cognition. Apart from the above, authenticity of AI-generated material can also be concern as a fake human avatar or voice might promote inclusivity, but this can also create problems. Students may be confused between what is real and what is simulated. The teaching guidelines for college and universities on ownership and copyright for generated material, permission and authenticity of the AI-synthesized elements are important. How, not if, the institution will integrate such content forms into the classroom to benefit students is something that needs serious consideration as its effect on students' learning process and self-assurance would vary for different context. In equipoise terms, multimodal AI is justified when it reduces extraneous cognitive load and supports the learning objective, not when it adds visual novelty or unnecessary media.

Theme 3: Agentic AI, Workflow Automation, and Development

Third, AI systems can support employees with tasks ranging from content creation and coding to workflow assistance, prototyping and step-by-step tasks. These factors indicate Microsoft Copilot, Lovable, and Manus signals a shift away from 1-step or request-response type interactions and more toward a sequence of step-based interactions and support. Actually, support may be present within a given use case to compose letters, collate data, design teaching material, generate simple prototypes of very small learning applications, to perform routine, very high frequency, administrative tasks. Such automation could free up worker tasks and allow them additional time to student interaction and feedback.

However, this delegation also raises questions of accountability. The grader and sum of evidence/recommendations provided by an AI (when creating an assessment, when grading) must, in all instances, be authenticated by the educator for the validity, fairness, and educational standards of the information. The greater the independence of the AI, the greater the oversight needed. Another issue apparent in this theme is the dichotomy of aid versus replacement; an educator may perhaps be responsibly assigned to create or reformat correspondence or format a paper, but a step-in student progress, where some level of judgment is inherent, would warrant increased human involvement and validation. It would be appropriate for an institution to judge applications of AI on a risk scale rather than as simply an "either/or" entity. In equipoise terms, agentic AI may be delegated routine tasks, but decisions involving judgement, assessment, or student progress require retained human oversight.

Theme 4: Efficiency Versus Cognitive Engagement

A trade-off is constant throughout these three functions; the balance between efficiency and cognitive engagement. AI is able to decrease the time needed to search, summarize, draft, translate, media create, or prototype. These are efficiencies which might be a considerable assist especially in context such as the number of tasks the educator may accomplish or where financial constraints can prohibit access by educators or students to such capabilities.

Conversely, some of these things we enable with artificial intelligence-reading closely, comparing sources, drafting, revising, computing, drawing up or fixing - are the issues students learn by doing and have heretofore depended on human interaction. This efficiency has no merit and, in fact, there can even be good efficiency since it has an altogether different quality to substitute intelligibility. In an assignment, whereby college students were set the task of writing a final essay, a key part of developmental work would be skipped, where the focus for a significant chunk of the essay writing would have to rely on the initiative of the AI; however tasks could be reimagined, whereby students were to critiquing an essay written by the AI; identifying the verification of its assertions, referencing where support to claims were sourced, rebutting or rethinking arguments and engaging through the Lens of construction not the loss of critical thought. Everything eventually comes back to constructivist pedagogy and critical AI literacy theories.

DISCUSSION

Equity, Accessibility, and the Digital Divide

The equity question is not peripheral to pedagogical equipoise; it is constitutive of the judgment itself. Equipoise presupposes a learner whose capacity to interrogate, verify, and revise AI output is unevenly distributed. The digital divide is best understood as stratified across three levels: material access to devices and connectivity (van Dijk, 2020); the skills needed to use digital resources productively (Hargittai, 2002); and the outcomes that accrue from use (van Deursen & Helsper, 2015). Generative and agentic AI intensifies all three simultaneously, raising material thresholds through subscription tiers, demanding a critical AI literacy not yet systematically taught (Long & Magerko, 2020), and rewarding those already positioned to convert fluent AI use into disciplinary advantage (Selwyn, 2022). Availability of a platform is therefore not equivalent to enablement of a learner. The capability approach clarifies why. Sen (1999) distinguishes between resources and capabilities—between what is nominally provided and what an individual is actually able to do. A student with nominal access

to a conversational AI but without the disciplinary background, language proficiency, or critical literacy to interrogate its output has not been equally enabled. Equity, on this reading, is a condition of realisable pedagogical capability, not of formal availability. UNESCO (2023) makes a parallel argument: without deliberate investment in teacher capacity and equitable infrastructure, generative AI risks amplifying rather than attenuating educational stratification. Equity thus becomes an internal condition of any defensible pedagogical judgment, not an external constraint upon it. For institutional readiness, this reframes AI adoption away from procurement. Institutions that distribute licences and publish tool inventories risk enacting a performative equity—measurable in deployment statistics but not in learner capability. A defensible approach integrates three interdependent commitments: infrastructure equity, ensuring connectivity and devices do not determine participation; capability equity, embedding critical AI literacy into curricula and professional development (Long & Magerko, 2020); and epistemic equity, calibrating AI use to learners' developing capacity rather than assuming it uniformly (Dwivedi et al., 2023; Vargas-Murillo et al., 2023). An equipoise judgment that weighs affordances against risks without accounting for positioned capability is not neutral; it is biased toward the already advantaged. Institutions should therefore evaluate AI adoption against capability outcomes rather than deployment metrics.

Privacy, Intellectual Property, and Data Governance

Privacy, intellectual property, and data governance are not separate compliance tasks to be handled one after another; together they form a single governance complex that conditions pedagogical equipoise. The first difficulty is classification. Institutions must distinguish between student records, assessment artefacts, sensitive personal information, and institutional research outputs, and decide which of these may legitimately be processed by AI tools. Existing regulations such as FERPA and GDPR were not written for generative systems, and their application becomes unclear when a student's assignment is uploaded for feedback or when an agentic tool chains several such operations across platforms (Malomo et al., 2025). Practical classification frameworks i.e public, internal, confidential, and regulated data offer a starting point, but they must be extended to account for AI-specific behaviours such as persistent retention, training on submitted content, and opaque vendor processing (Dwivedi et al., 2023). A parallel problem arises with intellectual property, because generative AI destabilizes authorship at the point of creation. When AI drafts teaching materials or student work, ownership does not fit neatly into existing policy or copyright categories. This matters beyond individual disputes: academic content is increasingly treated as an asset, with platforms and publishers licensing scholarly outputs and student work for model training, often without clear consent or benefit-sharing (Komljenovic, 2026). Institutions therefore carry a dual obligation—protecting the intellectual property of staff and students, and establishing defensible conventions for attribution, disclosure, and credit when AI contributes to knowledge production.

Data governance, in turn, is less a technical matter than a question of institutional readiness. Institutions that treat governance purely as restriction will struggle to support the context-sensitive judgments that equipoise demands. A more workable approach rests on three commitments: clear classification of data types and permissible AI uses; cross-functional oversight, so that academic, IT, and student services staff jointly review AI use cases as tools change; and transparency, so that students and educators understand what data is collected, how it is processed, and what rights they retain (Entsua-Mensah et al., 2026). Read through pedagogical equipoise, governance is not an external brake on AI adoption but a condition of its defensibility. An institution that cannot say where student data flows, who owns AI-generated artefacts, or how retention and access are governed has not made an equipoise judgment—it has postponed one. Governance, in this sense, is what allows pedagogical innovation to proceed without displacing institutional and learner accountability.

Environmental and Political-Economic Considerations

This growth also has implications beyond the immediate classroom space. Running and training these larger computer systems necessitates significant digital infrastructure, along with the corresponding energy use. Furthermore, many tools developed for education are now owned by, operated by, and supported by corporations, which makes them susceptible to platform constraints, corporate priorities, and product visions, the extent to which they effect practice will be conditioned by institutional dependencies established by them. Through these concerns, a critique could question more than if the tool works. It should also raise questions about what kind of institutional dependencies are developed by having teaching and learning be inextricably linked to a private

product. This provides critical digital pedagogy, from Selwyn (2022) to tool adoption based on institutional dependencies within institutional dependencies.

The Changing Role of the Educator: From Content Delivery to Curation

The shift from content delivery to curation is not simply a change in the teacher's task inventory; it reconstitutes the epistemic and ethical character of the educator's role. Where transmission pedagogy positioned the teacher as the authoritative source of knowledge, AI-mediated environments reposition the teacher as the site of judgment where competing knowledge claims—human and synthetic—are evaluated, reconciled, or rejected. This is a normative shift, not merely a functional one: the teacher is now responsible less for supplying information than for sustaining the conditions under which knowledge claims can be examined at all (Biesta, 2017; Popenici & Kerr, 2017). The curator role therefore encompasses several distinct judgments. The educator must assess whether an AI-generated explanation is factually sound, whether a multimodal artefact supports or distracts from the learning objective, whether an assessment task validly represents the construct it claims to measure, and whether a particular use of AI is ethically defensible in context. These are not technical decisions that can be delegated to policy or vendor guidance; they are pedagogical judgments requiring disciplinary expertise, knowledge of learners, and the capacity to hold uncertainty (Shulman, 1986; Mollick & Mollick, 2023). To exercise this role, a specific competence set is required: disciplinary expertise, pedagogical and design knowledge, AI and information literacy, assessment literacy, awareness of data and privacy issues, ethical judgement, and the ability to teach students to scrutinise AI output rather than accept it. These competences map onto the TPACK framework (Mishra & Koehler, 2006; Koehler & Mishra, 2009), which insists that technology be subordinated to pedagogy and content rather than driving either. Read through pedagogical equipoise, however, the educator's role involves more than integration. The educator is the practitioner who exercises equipoise—weighing affordances against risks in relation to specific learners, tasks, and institutional conditions. The competences above are the preconditions of that judgment, but the judgment itself—provisional, context-sensitive, and revisable—is the educator's distinctive professional contribution in an AI-mediated environment (Selwyn, 2022). Technology may be integrated; pedagogy remains the educator's responsibility.

Theory-to-Finding Integration

Table 2: Theory-to-Finding Integration Across Four AI Themes

Theme	Constructivist perspective	Multimedia learning	Critical AI literacy
Conversational AI	AI can scaffold dialogue, questioning and revision; learning weakens when AI replaces knowledge construction.	Less central; applies when text, audio or visual explanations are combined.	Students verify claims, question authority and recognize hallucinations or bias.
Multimodal AI	Students can construct and express knowledge through multiple representations.	Central: media should support the learning objective without unnecessary cognitive load.	Students consider authenticity, disclosure, intellectual property and synthetic-media risks.
Agentic AI	Automation should preserve meaningful human decision-making and reflection.	Applies to generated instructional materials and interfaces.	Users assess accountability, data governance, bias and platform dependence.
Efficiency–cognition	Tasks should retain active reasoning, revision and reflection.	Avoid unnecessary media and information overload.	Students learn when and why AI should or should not be used.

Taken together, these five considerations indicate that the pedagogical promise of generative and agentic AI cannot be assessed independently of the equity, governance, political-economic, and professional contexts in which it is deployed. Equitable access, sound data governance, awareness of environmental and corporate dependencies, and a redefined educator role are not peripheral add-ons to pedagogical decision-making; they are

constitutive of it. The theory-to-finding integration presented in Table 2 makes explicit that constructivist, multimedia, and critical AI literacy perspectives converge on a single conclusion: the value of any AI tool is contingent, provisional, and context-dependent, precisely the stance that pedagogical equipoise asks educators and institutions to adopt, rather than a fixed judgement that a tool is inherently beneficial or harmful.

Limitations and Future Directions

This study has several limitations. First, the AI landscape is evolving rapidly. Many of the tools analysed here are subject to frequent updates, changes in functionality, or discontinuation, meaning the findings represent a temporal snapshot rather than a stable or exhaustive account. New tools and capabilities continue to emerge, and the pedagogical interpretations offered in this paper may require revision as the technology matures. Second, this study employs qualitative document analysis rather than empirical testing. It identifies affordances and potential pedagogical risks based on published literature, theoretical frameworks, and available documentation, but it cannot establish a causal relationship between tool use and student learning outcomes. The 27 tools were not tested directly with teachers or learners, and no experimental or quasi-experimental design was employed. Consequently, the findings should be read as analytical claims and pedagogical interpretations, not as empirically demonstrated effects.

Third, the evidence base is not exclusively empirical. While preference was given to peer-reviewed sources, the analysis also draws on theoretical writing, review articles, and authoritative grey literature. This was necessary given the emerging nature of the field, but it means that some claims rest on conceptual argument rather than measured outcomes. Fourth, the thematic analysis was conducted by a single researcher. Although a coding matrix and explicit inclusion criteria were used to support transparency, inter-coder reliability could not be calculated. This limits the confirmability of the thematic structure. Future research would benefit from multiple coders, expert-coded systems, or independent audit trails to strengthen the credibility of the analysis. Fifth, the study focused primarily on the pedagogical and functional dimensions of AI tools. The institutional and commercial dimensions such as procurement processes, platform dependencies, vendor lock-in, and long-term sustainability were less fully explored. These factors are likely to shape how AI is adopted and sustained in higher education, and they merit dedicated investigation.

Finally, the concept of pedagogical equipoise is offered as a conceptual and practical framework. It has not been empirically validated in this study, and its usefulness for educators and institutions remains to be tested. Future research should examine how pedagogical equipoise is understood and applied in real institutional contexts.

Building on these limitations, several directions for future research are proposed. Longitudinal studies should examine the impact of AI tools on teaching and learning across an academic year or longer. Comparative studies should investigate different pedagogical uses of AI within and across disciplines. Mixed-methods designs should combine outcome measures with qualitative accounts of teaching and learning experiences. Institutional case studies should explore how governance, professional development, privacy, procurement, and equity considerations shape AI adoption. Further research should also examine the effect of AI on the construct validity of assessment and on the definition of literacy in higher education. Finally, the concept of pedagogical equipoise should be tested empirically to determine whether it offers a useful decision framework for educators and institutions navigating the complexities of generative and agentic AI.

These limitations establish important boundaries for interpreting the findings and also indicate the need for a progressive programme of empirical research. The present analysis should therefore be understood as a theoretically informed and temporally situated examination of the pedagogical possibilities and risks associated with emerging AI tools, rather than as evidence of their demonstrated effects on learning outcomes. Future empirical studies can build on this foundation by examining how different forms of AI use operate across disciplines, learner groups, institutional contexts, and levels of AI readiness. In particular, research that combines learning outcomes with qualitative evidence of learner and educator experiences could provide stronger evidence concerning when AI supports learning and when it displaces essential cognitive and pedagogical activity. Such work would also provide an empirical basis for refining and validating pedagogical equipoise as a practical framework for AI-related decision-making in higher education

CONCLUSION

AI is starting to be implemented within higher education but its pedagogic relevance cannot be ascertained by complex of technologies. Analysis of 27 AI tools uncovered potential benefits from the implementation of dialogic scaffolding, multiple modalities, accessibility, content creation, automation of workflow and application building; however, a host of interconnected risks accompany these benefits, stemming from displacement of human cognition, student cheating, disparity in access and effects, privacy and intellectual property. The manipulation of synthetically derived media, the ecological impact of generative systems, and reliance on closed-loop proprietary platforms all factor into the use and governance of synthetic media. AI is fundamentally a pedagogical as well as institutional concern within higher education. Teachers ought to design tasks where AI complements (rather than substitute for) human reasoning, reflection, dialogue, and engagement with content. Institutions should pair access to tools with AI literacy, continuing professional education, strong data governance, ethical clear guidelines, and appropriate assessment methods. The role of the educator should move from an end-user of the technology to a curator and critic able to judge where AI offers instructional affordances and where human thought must take center stage. The contribution of this paper is the concept of pedagogical equipoise, which requires educators and institutions to judge not only what AI can do, but what it should be permitted to do in a given learning context. Pedagogical equipoise thus moves the debate beyond tool adoption toward critical, context-sensitive, and ethically grounded pedagogical judgement.

Ultimately, the central implication of this analysis is that the question facing higher education is no longer simply whether AI should be used, but how its use can be situated within sound pedagogical, ethical, and institutional judgement. The analysis of the 27 tools illustrates that generative and agentic AI can extend opportunities for dialogue, representation, accessibility, content development, and workflow support, while simultaneously creating risks relating to cognitive engagement, equity, privacy, academic integrity, accountability, and institutional dependence. Pedagogical equipoise provides a way of holding these possibilities and risks in productive tension by requiring decisions to remain evidence-informed, context-sensitive, and open to revision. AI should therefore be integrated where it demonstrably supports the purposes of learning while preserving the human capacities that higher education seeks to develop. The enduring principle is not technological adoption for its own sake, but the deliberate alignment of AI use with human reasoning, learner agency, disciplinary knowledge, ethical responsibility, and meaningful educational purposes.

REFERENCES

1. Biesta, G. J. J. (2017). *The rediscovery of teaching*. Routledge.
2. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
3. Collins, A. (1996). Design issues for learning environments. In S. Vosniadou, E. De Corte, R. Glaser, & H. Mandl (Eds.), *International perspectives on the design of technology-supported learning environments* (pp. 347–361). Lawrence Erlbaum Associates
4. Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023). "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
5. Entsua-Mensah, F. O., Kwanya, T., & Hoskins, R. (2026). Balancing innovation and privacy: A framework for personal data protection in AI-enhanced higher education in Africa. *Education for Information*. <https://doi.org/10.3233/EFI-240012>
6. Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. *First Monday*, 7(4). <https://doi.org/10.5210/fm.v7i4.942>
7. Imran, M., & Almusharraf, N. (2024). Google Gemini as a next generation AI educational tool: A review of emerging educational technology. *Smart Learning Environments*, 11(1), Article 22. <https://doi.org/10.1186/s40561-024-00310-z>.

8. Koehler, M. J., & Mishra, P. (2009). What is technological pedagogical content knowledge? *Contemporary Issues in Technology and Teacher Education*, 9(1), 60–70
9. Komljenovic, J. (2026). *Code acts in education: AI and the amplification of academic content assetization*. Education International.
10. Lee, D., Arnold, M., Srivastava, A., Plastow, K., Strelan, P., Ploeckl, F., Lekkas, D., & Palmer, E. (2024). The impact of generative AI on higher education learning and teaching: A study of educators' perspectives. *Computers and Education: Artificial Intelligence*, 6, Article 100221. <https://doi.org/10.1016/j.caeai.2024.100221>
11. Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). ACM. <https://doi.org/10.1145/3313831.3376727>
12. Malomo, O. O., Adekoya, A. A., Donald, A. M., Eyob, E., Omojokun, E., Mummalaneni, V., & Garuba, M. (2025). AI as asset and liability: A dual-use dilemma in higher education and the SPARKE Framework for institutional AI governance. *Online Journal of Applied Knowledge Management*, 13(2), 57–76.
13. Mayer, R. E. (2021). *Multimedia learning* (3rd ed.). Cambridge University Press.
14. Mishra, P., & Koehler, M. J. (2006). Technological pedagogical content knowledge: A framework for teacher knowledge. *Teachers College Record*, 108(6), 1017–1054. <https://doi.org/10.1111/j.1467-9620.2006.00684.x>
15. Mollick, E. R., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts [Working paper]. Wharton School Research Paper. <https://doi.org/10.48550/arXiv.2306.10052>
16. 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. <https://doi.org/10.1186/s41039-017-0062-8>
17. Selwyn, N. (2022). *Education and technology: Key issues and debates* (3rd ed.). Bloomsbury Academic.
18. Sen, A. (1999). *Development as freedom*. Oxford University Press.
19. Shneiderman, B. (2020). *Human-centered AI*. Oxford University Press.
20. Shulman, L. S. (1986). Those who understand: Knowledge growth in teaching. *Educational Researcher*, 15(2), 4–14. <https://doi.org/10.3102/0013189X015002004>
21. UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://doi.org/10.54675/EWZM9535>
22. Van Deursen, A. J., & Helsper, E. J. (2015). The third-level digital divide: Who benefits most from being online? In *Communication and information technologies annual* (Vol. 10, pp. 29–52). Emerald. <https://doi.org/10.1108/S2050-206020150000010002>
23. Van Dijk, J. A. G. M. (2020). *The digital divide*. Polity Press.
24. Vargas-Murillo, A. R., de la Asuncion Pari-Bedoya, I. N. M., & de Jesús Guevara-Soto, F. (2023). Challenges and opportunities of AI-assisted learning: A systematic literature review on the impact of ChatGPT usage in higher education. *International Journal of Learning, Teaching and Educational Research*, 22(7). <https://doi.org/10.26803/ijlter.22.7.7>