DOI: 10.54254/2753-7048/2026.LD36020

Educational Purpose under AI's Technical Logic: Goal Substitution and Stage-Sensitive Governance in Basic and Higher Education

Yang Han[1*], Teng Wang[2]

School of Education, Beijing Foreign Studies University, Beijing, China
**Corresponding Author. Email: han13@bfsu.edu.cn*

Abstract. Artificial intelligence is becoming part of everyday teaching, assessment, and knowledge production in schools and universities. Debate has concentrated on the functions of these systems, but the educational consequences always depend on the logic through which they operate. AI applications are commonly organized around optimization, automation, prediction, datafication, and scale. Education follows a different set of commitments, including developmental appropriateness, sustained cognitive effort, professional judgement, human interaction, and responsibility for knowledge. Using critical conceptual analysis and directed content analysis of international policy documents and academic literature, this paper examines how these logics interact in basic and higher education. The analysis identifies goal displacement as the central problem. Indicators that are easy to measure and optimize can gradually redefine what institutions treat as learning. The consequences differ across educational stages. In basic education, premature cognitive outsourcing may interrupt the formation of foundational capabilities and reduce meaningful interaction. In higher education, AI-assisted knowledge production raises concerns about verification, authorship, disciplinary judgement, and academic responsibility. A Pedagogical Primacy Framework is proposed to guide educational decisions through four connected considerations: educational purpose, cognitive necessity, human agency, and accountability. The framework supports teacher-governed use in basic education and transparent, reviewable use in higher education.

Keywords: artificial intelligence, educational rationality, basic education, higher education, pedagogical primacy

1. Introduction

Artificial intelligence has moved from a specialist educational technology to an everyday part of teaching, assessment, administration, and academic work. Adaptive platforms recommend exercises, analytics flag patterns of participation, and generative systems produce explanations, essays, code, images, and research summaries. These tools can widen access and reduce repetitive work, yet their value rests on instructional design. International guidance places educational use under human

DOI: 10.54254/2753-7048/2026.LD36020

purposes [1], while recent evidence shows that unrestricted task outsourcing can improve a product without producing comparable learning [2].

The key question is how AI's operating logic fits educational purposes. Most systems work toward predefined outputs, draw predictions from data, and convert activity into forms that can be processed at scale. These features favor outcomes that are easy to record and compare. A score, completion rate, engagement pattern, or polished product can therefore gain more institutional weight than the slower development it is meant to represent. Learning also depends on interpretation, dialogue, uncertainty, and gradual change, all of which resist simple technical measures.

Generative AI makes this tension harder to ignore. Earlier research on AI in higher education was already dominated by technical and quantitative perspectives, with educators playing a limited role [3]. Current systems can now perform substantial parts of reading, writing, coding, and analysis. Their fluent output may conceal the difference between developed understanding and accepted assistance. At the institutional level, automation also promises faster assessment and student support, encouraging adoption before its educational value has been clearly established.

Educational decisions therefore require a clear account of purpose. Biesta shows that measurement always carries a judgement about what is worth pursuing [4]. Once an indicator guides funding, intervention, or evaluation, it begins to shape practice. Research on automation and datafication similarly locates AI within policy, procurement, and commercial arrangements [5]. This paper examines how those arrangements can displace educational aims and redistribute cognitive work, judgement, and responsibility. It then compares basic and higher education and develops a Pedagogical Primacy Framework for stage-sensitive decisions.

The comparison is stage-sensitive because the same assistance can carry different educational meanings. School students are still acquiring the capabilities that make later independence possible, whereas university students are expected to participate in disciplinary inquiry and accept responsibility for their claims. The relevant boundary therefore changes with learner maturity, the purpose of the task, and the consequences of error. This approach avoids a universal rule for every tool while retaining a common standard: technology should strengthen the human capacities on which education depends.

2. Technical logic and educational rationality

Technical logic refers to recurring tendencies produced by the design and use of AI. Optimization privileges specified outcomes; automation transfers parts of a task to a system; prediction projects earlier patterns forward; datafication turns activity into machine-readable traces; and scale rewards standardization. An adaptive platform may combine all five by recording behaviour, predicting the next task, and optimizing progress against a predefined model. The difficulty begins when categories built for the system become the main categories through which an institution understands learning.

Educational rationality concerns what teaching is for and what learners must experience to reach that purpose. It takes account of age, prior understanding, sustained participation, professional judgement, and the relationships through which students become more independent. Efficiency matters only in relation to these aims. A quick route to a correct answer may help when students are applying an established procedure. The same shortcut can weaken learning when they need to formulate a problem, weigh evidence, or revise an argument. Automated feedback, prediction, and personalization must therefore be judged by the educational work they enable or remove.

The point of tension is institutional judgement. AI products often present their categories as neutral: a learner is classified as engaged, at risk, or on track, and an automated score appears

DOI: 10.54254/2753-7048/2026.LD36020

objective. Selwyn treats adoption as a social and political choice because every system carries assumptions about desirable outcomes, legitimate knowledge, and control [6]. Technical outputs can inform professional judgement, but time pressure and performance targets may turn them into default decisions. Teachers and students then remain accountable for choices they could not fully examine.

Recent guidance has begun to address this problem. UNESCO's student and teacher frameworks connect technical competence with critical judgement, ethical responsibility, pedagogy, and human agency [7, 8]. European guidance adds legal and data-protection concerns [9]. These principles move the debate beyond access, but they do not settle what should happen in a reading lesson, an automated placement decision, a university essay, or a research project. Such decisions depend on the capability being developed and the consequences attached to the task.

Educational stage is crucial because support, competence, and responsibility change over time. School students are still developing foundational literacies, attention, self-regulation, and social capacities. A review of K-12 uses reports gains in personalization and accessibility alongside dependency, privacy, reliability, and developmental concerns [10]. University students have greater independence and must learn to use tools common in professional practice. Their standard is the ability to understand the task, evaluate the output, explain the reasoning, and accept responsibility. Over-reliance can erode these abilities when fluent answers are accepted without verification [11].

3. Methodology

The study combines critical conceptual analysis with directed qualitative content analysis. Conceptual analysis clarifies how terms such as optimization, personalization, efficiency, agency, and responsibility change meaning when applied to education. Directed analysis connects this argument to a purposive corpus of policy documents and scholarly research. The study focuses on mechanisms through which AI can alter educational purposes and practices; theme frequency and full systematic coverage fall outside its scope.

The corpus comprises thirteen texts published between 2009 and 2026. Policy sources cover generative AI, student and teacher competence, ethical use, and implementation. Scholarly sources address educational purpose, automation, datafication, school integration, cognitive over-reliance, academic integrity, and university policy. Sources were included when they directly connected technological functions with educational aims or helped distinguish basic from higher education. Policy recommendations supplied normative guidance; classroom implementation required separate evidence.

The analysis first coded how each source described AI's intended outcomes and performance, using optimization, automation, prediction, datafication, and scale as organizing concepts. It then traced effects on capability formation, cognitive effort, professional judgement, interaction, authorship, and responsibility. Comparison across sources revealed a common mechanism: a measurable representation of education can acquire more authority than the purpose it was meant to serve. This mechanism is termed goal displacement. The interpretation was conducted by one researcher, so the framework remains an analytical proposal for testing through classroom observation, interviews, and comparative cases.

The corpus is deliberately narrow and conceptually focused. It captures a selected line of argument in a rapidly expanding field and cannot show how widely the identified patterns occur in practice. Intended functions were also distinguished from institutional effects. A system designed as support may still narrow judgement when it is linked to performance targets or used under severe

DOI: 10.54254/2753-7048/2026.LD36020

time pressure. This distinction keeps the argument attentive to the conditions of adoption as well as to system design.

4. How AI reshapes educational purpose

4.1. Goal displacement and the redefinition of learning

AI systems need a representation of the outcome they are expected to improve. In education, that representation is usually a proxy: correctness, completion, time on task, predicted attainment, participation, or similarity to a model answer. Each can provide useful evidence. Goal displacement occurs when the proxy begins to define success. A platform may raise the number of correct answers while reducing open inquiry; a writing assistant may improve coherence while hiding whether the student can form an argument; and automated scoring may exclude qualities that require interpretation.

This shift is reinforced by institutional pressure. Schools and universities need visible evidence of performance, administrators want comparable data, and providers need to demonstrate value. Dashboards and usage rates answer all three demands. Once embedded in intervention, procurement, or reporting, their categories become self-reinforcing: teachers respond to what the dashboard records, and activity within a purchased platform is read as successful implementation. Digital modernization is then judged by the presence and use of systems, with less attention to learning transfer, independence, dialogue, or professional judgement.

Personalization shows how a useful function can narrow educational vision. A system can adjust pace, sequence, or format from recorded behaviour, which may improve access and practice. Personal education also requires knowledge of a learner's interests, relationships, uncertainty, and changing aspirations. Data-driven models place the learner within a pattern, whereas teachers sometimes need to challenge that pattern. Past struggle may call for a demanding opportunity even when the model recommends another simplified task. Apparent disengagement may reflect anxiety, language transition, family circumstances, or mistrust that behavioural traces cannot explain.

Prediction sharpens the problem because education aims to change what is currently probable. Early-warning systems may direct support to a student, but a provisional risk score can also lower expectations or restrict opportunity. Group-level accuracy alone cannot decide how a classification should guide an individual case. Educational use must therefore focus on what follows the prediction. A useful signal opens professional inquiry and additional support; it should remain explainable, revisable, and open to challenge.

4.2. Cognitive outsourcing, agency, and responsibility

Automation redistributes cognitive work: for example, calculators, search engines, translation tools, and statistical software already allow learners to pursue problems that would otherwise exceed their available time or current skill level. Generative AI extends this support to explanation, synthesis, writing, coding, and analysis. Its educational value depends on the operation delegated. Removing clerical work can create space for deeper thought. Delegating the activity through which competence is formed can weaken the learning task.

Some effort is productive because it exposes and develops understanding. Retrieval strengthens memory; drafting reveals gaps; comparison develops judgement; and revision connects feedback with intention. Generative systems can shorten each process by supplying an answer before retrieval, a polished paragraph before an argument, or a summary before the source has been read.

DOI: 10.54254/2753-7048/2026.LD36020

The result may meet the task specification while offering little evidence of learning. Confusing instructions and repetitive work still obstruct learning. The practical distinction lies between removing barriers to participation and removing the cognitive activity that gives a task its educational value.

Over-reliance becomes more likely when users lack the knowledge needed to judge fluent output. Students with weaker subject knowledge may be least able to detect error and most tempted to accept a plausible answer. Zhai and colleagues associate repeated dependence on AI dialogue systems with weaker critical thinking, analytical reasoning, and decision-making [11]. Accuracy alone is insufficient: even a correct answer can interrupt learning when it arrives before the learner can evaluate how it was produced.

Redistributed work also affects agency and interaction. Teachers may use AI to prepare materials, interpret data, or draft feedback, but their authority weakens when recommendations are difficult to question or institutionally expected. Meaningful oversight requires time, information, competence, and the power to reject an output. The same principle applies to students placed on adaptive pathways. Dialogue matters because it allows uncertainty, misunderstanding, and intellectual risk to become visible. A human presence offers little protection when the system still determines the sequence, categories, and available choices.

Responsibility becomes decisive in assessment and knowledge production. Data selection, category design, thresholds, and institutional action remain human choices even when an automated result appears objective. A student must be able to challenge a consequential assessment, and the institution must explain how it was reached and who can revise it. In higher education, responsibility also covers evidence, interpretation, and claims. Academic integrity research favours assessment that keeps reasoning, source use, and authorship visible [12]. University policies increasingly permit AI under course-level conditions and varied forms of assessment [13]. The central test is whether students can account for their work and whether instructors can still judge learning validly.

Human interaction deserves separate attention because feedback and dialogue do more than transmit information. They help teachers interpret misunderstanding, communicate expectations, and build the trust required for intellectual risk. Immediate machine feedback may be useful, but continuous mediation can reduce opportunities to ask uncertain questions, observe how others reason, and negotiate meaning. These losses rarely appear in performance data. Their importance lies in the quality of participation and in students' willingness to remain engaged when learning becomes difficult.

5. Stage-sensitive governance of artificial intelligence

5.1. Basic education and the formation of capability

Basic education should give priority to formation before augmentation. Children and adolescents are developing the practices that make later independence possible: sustained reading, clear writing, mathematical reasoning, attentive listening, self-regulation, and revision. Early substitution can interrupt this development. Generated summaries may let a student complete a reading task without identifying its structure or evidence. Automatic composition may produce acceptable prose before the student has gained control of sentences and arguments. Immediate performance can therefore conceal a developmental loss.

The boundary should follow the learning objective and the student's current competence. Text-to-speech can widen access for a learner with a visual impairment or reading difficulty, translation can clarify instructions, and automated practice can provide timely support. These uses enable

DOI: 10.54254/2753-7048/2026.LD36020

participation. Generating the core response that a student is expected to learn has a different effect. Teachers should decide which part of a task may be supported, for how long, and when the support should be withdrawn.

The same judgement applies to devices and platforms. Equipment can extend access to information, simulation, and communication; it can also fragment attention, intensify surveillance, and place a commercial system between teacher and learner. Usage rates and device coverage are therefore weak measures of progress. Evaluation should examine transfer beyond the platform, student independence, quality of interaction, teacher workload, and differences in access and benefit. Technology is valuable when it creates opportunities for observation, explanation, practice, and discussion that the lesson could not otherwise provide.

Prediction requires particular restraint in schools because younger students have limited power to understand or challenge classifications. Risk scores should prompt inquiry, not fix pathways or lower expectations. Teachers need to know the basis of a recommendation and retain authority to reject it; students and families need a route to appeal. Institutions should also test whether the underlying data reproduce earlier inequalities. These safeguards keep uncertainty visible and prevent a probable outcome from becoming a fixed expectation.

Teacher authority provides the practical basis for these safeguards. Educators should govern access, timing, prompts, assessment, and withdrawal, and they need enough information and training to exercise that control. AI is useful when it enlarges professional understanding and leaves contextual judgement intact. A school that purchases systems without investing in teacher capacity may increase technical dependence while weakening the people responsible for education.

5.2. Higher education and epistemic responsibility

Higher education requires a wider scope for use because students are entering disciplines and professions in which AI is already common. A general ban would separate university work from contemporary practice, while unrestricted generation would make a submitted product a poor measure of understanding. The boundary should follow the capability being assessed and the responsibility attached to the final claim.

Students may use AI to generate alternatives, debug code, reorganize notes, improve language, explore examples, or support data analysis when those operations are outside the assessed competence. The process should remain transparent and open to review. Students need to verify sources, explain important decisions, reproduce key reasoning, and identify the system's contribution. Acceptable assistance will differ across fields and tasks. Language editing may be appropriate in one assignment and undermine another that assesses original composition; code generation may support a design project but invalidate a test of foundational programming.

Course-level guidance is more useful than a single rule for every task. Universities need common expectations for transparency, privacy, academic responsibility, and misuse, while instructors specify permitted use in relation to learning outcomes. Assessment can combine final products with annotated drafts, version histories, source audits, oral explanation, in-class application, or reflective commentary. These forms provide enough process evidence for a valid judgement without surrounding every assignment with surveillance. Requiring students to explain and defend their work keeps the focus on inquiry and understanding.

Universities must also teach students how generative systems work, where their limitations arise, and how their use affects authorship, labour, privacy, and professional standards. Epistemic responsibility includes knowing when a tool should be set aside. Plausible syntheses can hide missing evidence, generated citations may be false, and dominant training patterns may narrow

DOI: 10.54254/2753-7048/2026.LD36020

disciplinary perspectives. Responsible use therefore requires verification, critical distance, and an awareness that efficiency can change the character of academic work. Faster drafting should create room for stronger questions and deeper interpretation while the expected volume of output remains proportionate.

Academic integrity follows from this wider responsibility. Text detection is too unreliable to serve as a complete solution and can draw attention away from assessment design. A stronger approach makes important stages of student work visible and requires meaningful disclosure. A useful disclosure shows how AI affected the argument, evidence, data, code, organization, or language. This allows readers and instructors to identify the student's contribution and judge the work on appropriate grounds.

5.3. The Pedagogical Primacy Framework

The Pedagogical Primacy Framework translates this analysis into four questions for classroom and institutional decisions. First, what capability, understanding, relationship, or form of participation is the activity meant to develop? Second, which parts of the work must the learner perform for that purpose to be achieved? Third, can teachers and students understand, question, modify, and reject the system's output? Fourth, who must explain the decision or claim, and who can correct it? Educational purpose comes before tool selection, and the four questions are considered together. Strong technical performance has little educational value when learning is bypassed or responsibility remains unclear.

Table 1. Stage-sensitive application of the Pedagogical Primacy Framework

Dimension	Basic education	Higher education
Educational purpose	Foundational development, inclusion, and meaningful interaction	Disciplinary learning, inquiry, and preparation for responsible practice
Cognitive necessity	Core literacy, numeracy, reasoning, and self-regulation remain with the learner until competence is established	Support is permitted outside the assessed competence; verification and explanation remain necessary
Human agency	Teachers govern access, timing, prompts, and withdrawal; classifications remain open to challenge	Students and teachers choose tools within clear course and task rules
Accountability	Teachers and schools retain responsibility for feedback, placement, and consequential decisions	Students are responsible for submitted claims; instructors remain responsible for assessment judgements
Appropriate role	A temporary, teacher-governed learning scaffold	A transparent and reviewable tool for knowledge production

The framework changes procurement and evaluation. Adoption should begin with a documented educational problem and comparison with non-technological options. Pilots should examine transfer beyond the system, student independence, teacher control, and the distribution of benefits and risks. Satisfaction and short-term task performance provide only partial evidence. Systems involved in consequential decisions also need understandable grounds for review, a human route for appeal, and a clear allocation of authority.

Equipment creates educational value through the tasks, relationships, and judgements it supports. More technology may widen access, yet weak design can also increase distraction and dependence.

DOI: 10.54254/2753-7048/2026.LD36020

AI use is justified when its educational contribution outweighs the cognitive, relational, ethical, and equity costs. That judgement may support extensive use in one setting and limitation in another.

6. Conclusion

AI can alter what institutions value, which activities students perform, how teachers exercise judgement, and who is responsible for decisions. The main risk is goal displacement: measurable outputs gradually stand in for broader educational aims. Cognitive outsourcing, predictive classification, and automated decisions matter because they reshape capability, agency, and accountability.

Governance should reflect educational stage. Schools need teacher-controlled uses that protect foundational development and meaningful interaction. Universities can allow broader use when reasoning, verification, authorship, and responsibility remain visible. The Pedagogical Primacy Framework links both settings through educational purpose, cognitive necessity, human agency, and accountability. AI integration should be judged by the capacities it strengthens.

References

[1] Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. *UNESCO*. https://unesdoc.unesco.org/ark:/48223/pf0000386693

[2] OECD. (2026). OECD Digital Education Outlook 2026: Exploring effective uses of generative AI in education. *OECD Publishing*. https://doi.org/10.1787/062a7394-en

[3] Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education, 16*, 39. https://doi.org/10.1186/s41239-019-0171-0

[4] Biesta, G. (2009). Good education in an age of measurement: On the need to reconnect with the question of purpose in education. *Educational Assessment, Evaluation and Accountability, 21*(1), 33-46. https://doi.org/10.1007/s11092-008-9064-9

[5] Williamson, B., Macgilchrist, F., & Potter, J. (2023). Re-examining AI, automation and datafication in education. *Learning, Media and Technology, 48*(1), 1-5. https://doi.org/10.1080/17439884.2023.2167830

[6] Selwyn, N. (2019). Should robots replace teachers? AI and the future of education. *Polity Press*.

[7] Miao, F., Shiohira, K., & Lao, N. (2024). AI competency framework for students. *UNESCO*.

[8] Miao, F., & Cukurova, M. (2024). AI competency framework for teachers. *UNESCO*. https://doi.org/10.54675/ZJTE2084

[9] European Commission. (2022). Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators. *Publications Office of the European Union*. https://doi.org/10.2766/153756

[10] Marzano, D. (2025). Generative artificial intelligence in teaching and learning processes at the K-12 level: A systematic review. *Technology, Knowledge and Learning*. https://doi.org/10.1007/s10758-025-09853-7

[11] Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. *Smart Learning Environments, 11*, 28. https://doi.org/10.1186/s40561-024-00316-7

[12] Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International, 61*(2), 228-239. https://doi.org/10.1080/14703297.2023.2190148

[13] Wang, H., Dang, A., Wu, Z., & Mac, S. (2024). Generative AI in higher education: Seeing ChatGPT through universities' policies, resources, and guidelines. *Computers and Education: Artificial Intelligence, 7*, 100326. https://doi.org/10.1016/j.caeai.2024.100326