International Journal of Education, Vocational and Social Science
https://e-journal.citakonsulfindo.or.id/index.php/IJEVSS
E-ISSN : 2963 4954
IJEVSS

Reconstructing the Aims of Higher Education in the Era of Generative AI: A Philosophical Account of Student Agency and Justified Delegation

Deni Yuliana
Universitas Sultan Ageng Tirtayasa
Aceng Hasani
Universitas Sultan Ageng Tirtayasa
Erwin Salpa Riansi
Universitas Sultan Ageng Tirtayasa

Corresponding Author: *1,3 deni.universal627@gmail.com, salpariansierwin@untirta.ac.id

Abstract

The arrival of generative artificial intelligence (AI) has made it necessary to revisit an old philosophical question: how is the production of academic work related to the formation of the person who produces it? This article explores how the aims of higher education might be rebuilt around student agency, drawing on epistemological, ethical, and humanistic lines of inquiry. The study takes the form of a conceptual literature review based on purposively selected publications from 2020 to 2025. Twelve full-text sources, consisting of journal articles and a UNESCO guidance document, shaped both the methodological design and the substantive analysis. Empirical findings, conceptual arguments, and policy recommendations were kept apart during the analysis before being brought together in a normative synthesis.

Three interrelated educational commitments emerge from this synthesis: understanding that students are able to explain and critically revise, responsibility for decisions taken with the help of AI, and participation in learning that honours personal growth and fair access. On this basis, the article offers a deliberative framework for judging delegation to AI, organised around five considerations: learning aims, the activities handed over, verification, accountability, and student participation. The framework allows for substantial assistance where it widens access and supports learning, yet it asks for evidence that students remain answerable for the intellectual work under assessment. What the article contributes is a philosophical account of justified

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

delegation, not an empirically tested assessment tool. Its practical implications point toward process-sensitive assessment, clear rules for AI assistance, and institutional responsibility for creating the conditions in which students can exercise meaningful agency.

Keywords: generative artificial intelligence, philosophy of education, higher education, student agency, educational aims

INTRODUCTION

Generative AI can now draft explanations, summarise readings, and propose answers within seconds. This capacity confronts higher education with a pressing question about what learning is ultimately for. Once an academic product can be generated with the help of a system, the quality of that product no longer tells us much about what a student has actually understood. The issue therefore touches the relationship among knowledge, action, and the formation of the person. Questions about how AI should be used must be accompanied by a more fundamental one: which capacities ought to develop in students when part of their intellectual work can be delegated?

Scholarship on AI in education offers a useful backdrop for this question. Holmes and Tuomi (2022) show that discussions of educational AI carry implicit assumptions about learning, assessment, and human control. Kasneci et al. (2023) describe how language models could support educational activities, while also stressing the continued need for fact-checking, critical thinking, and human oversight. Both works help situate technology within pedagogical reasoning. Even so, what a system is capable of doing does not, by itself, settle which aims higher education ought to pursue.

Students' own experiences also deserve attention, provided that we do not ask more of the evidence than it can give. A survey by Chan and Hu (2023) involving 399 students from six universities in Hong Kong revealed a wide range of views on the benefits and challenges of generative AI. These findings give good reason to involve students in shaping how such tools are governed. Because the data are self-reported, however, they show neither that AI improves understanding nor that it weakens students' thinking. Keeping perceptions, pedagogical possibilities, and measured effects distinct is a precondition for responsible argument.

Broader reviews similarly point to the need for attention to ethics and to the quality of knowledge about educational AI. Bond et al. (2024), in a meta-review of 66 reviews, highlight the

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

need for methodological rigour, collaboration, and ethical care. Mouta et al. (2024) identify issues in educational ethics that remain underrepresented in the body of work they examined. These two sources support a critical look at the grounds on which decisions are made, but they are not used here to suggest that every AI practice suffers from the same weaknesses. It should also be noted that the scope of educational AI in these reviews is wider than generative AI alone.

Within the philosophy of education, Herman and Lara-Steidel (2025) offer a discussion that speaks directly to the purposes of the university. They treat understanding as an educational achievement distinct from the mere availability of information. The present article takes up this line of thought by asking under what conditions delegation to AI is acceptable: when does assistance enlarge students' opportunities to understand, and when does it take over the very activities that learning is meant to cultivate? Neither setting a quota on AI use nor banning every form of assistance answers that question.

In this article, student agency is understood as a capacity supported by genuine opportunities to decide on learning steps, weigh reasons, revise decisions, and answer for one's actions. This is a working definition for the purposes of analysis; it does not presume that all philosophical traditions share a single concept of agency. Students act through their relationships with lecturers, peers, scholarly sources, and technologies. Agency, therefore, does not require isolated self-sufficiency. On the other hand, student involvement loses its meaning when important decisions simply follow system output, without understanding and without any real possibility of rejecting it.

Three perspectives guide the analysis. The epistemological perspective asks why an answer deserves to be trusted and how understanding can be demonstrated. The ethical perspective concerns accountability, fairness, and the limits on how others may be treated. The humanistic perspective serves as a normative orientation that values personal development, student voice, and educational relationships. This last sense should not be equated with every school of humanism, nor with the technical label of human-centred technology design. Holding these distinctions in place prevents technical vocabulary from standing in for arguments about educational value.

This study aims to reformulate the aims of higher education by examining how understanding, responsibility, and personal development relate to one another in the use of generative AI. The research question is: how can delegation to AI be philosophically justified while still preserving student agency? The article contributes a normative framework that links learning

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

aims with the activities being delegated, verification, accountability, and participation. The framework is meant to support reasoning about learning design; it is not intended to yield a universal metric or to make untested claims about effectiveness.

RESEARCH METHODOLOGY

This research is a conceptual literature study carried out through critical philosophical analysis. The argument was developed in line with the principle of traceability, that is, keeping visible the links among the selection of literature, the use of concepts, and the conceptual contribution, as discussed by Jaakkola (2020). This design serves as a cross-disciplinary methodological reference rather than as a substantive theory of education. No respondents, experiments, or impact measurements were involved. Since sources were chosen purposively and were not meant to cover the entire body of publications, the study is not presented as a systematic review.

The publication window was set at 2020–2025, covering six calendar years inclusively. Searches and access checks were conducted on 26 October 2026 through web searches, publisher websites, academic repositories, and official UNESCO sources. Search strings combined terms such as generative artificial intelligence, higher education, educational aims, understanding, student agency, ethics, and human-centred AI, supplemented by title or DOI searches. The year of searching is distinct from the publication cut-off, so sources published in 2026 were not used to expand the corpus.

Sources were selected according to four considerations: relevance to educational aims or agency, clear publication details, fit with the time window, and the availability of a full text that could be examined. The final corpus comprised 12 sources: ten substantive journal articles, one methodological article, and one UNESCO guidance document. Two candidate sources whose full texts could not be obtained were excluded from the analysis. Author versions were used when they were accessible through repositories or open archives, while citations still refer to the published journal versions. Sources predating the widespread use of generative AI were drawn on for conceptual grounding in educational AI, not as evidence about experiences with the technology that emerged later.

The analysis proceeded through five interrelated steps. First, sources were mapped by type so that survey results, review syntheses, philosophical arguments, and policy recommendations would not be treated as evidence of equal standing. Second, relevant passages

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

were examined to identify concepts, supporting reasons, and the limits of each claim. Third, the relationships among ideas were compared through epistemological, ethical, and humanistic categories. Fourth, the argument was tested against conceptual objections concerning delegation, accessibility, and inequality. Fifth, the results of this comparison were synthesised into a framework for judging justified delegation.

Analytic notes recorded each source's details, the sections examined, how the source was used in the argument, and its limitations. This procedure was intended to leave a trail of reasoning, without claiming independent coding or inter-rater reliability. The outcomes of the study take the form of normative propositions and pedagogical implications. Throughout this article, the word results refers to the products of synthesis and argument rather than to new field findings. The synthesis is presented together with the discussion and is followed by its philosophical interpretation and evaluation.

DISCUSSION

Synthesis of the Review

Comparing the literature yielded three conceptual relationships that underpin the discussion. First, access to answers should be distinguished from the ability to grasp the reasons that support those answers. Second, the use of assistance does not remove the need for accountability; rather, responsibility has to be shared in proportion to the control each party holds. Third, student control depends on educational conditions that make participation possible, not merely on a formal right to accept or refuse a technology. These relationships represent the article's interpretive synthesis rather than the outcome of any joint measurement across all sources.

The first relationship draws chiefly on discussions of the epistemic aims of education and the limits of AI output. The second rests on work concerning educational ethics and the governance of technology use. The third gains its context from studies of human control and stakeholder participation. These sources differ in purpose, method, and technological scope. For this reason, the synthesis does not add their results together as statistical evidence; instead, it examines how each line of reasoning helps address the same problem.

The central proposition is that delegation can be justified when AI assistance aligns with learning aims and still leaves students adequate opportunity to show understanding, make decisions, and explain their responsibility. Such alignment is context-dependent. An activity that

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

may reasonably be delegated in one course might be the core skill students must practise in another. The aims of higher education therefore need to be translated into reasons that can be scrutinised at the level of tasks, assessment, and institutional policy.

Epistemological Perspective on Understanding and Delegation

The gap between product and learning process becomes visible when an essay meets every linguistic criterion, yet its author cannot explain how the thesis connects to the evidence. The problem is not simply who composed the sentences, but what can rightly be attributed to the student as a learning achievement. Following the question of understanding raised by Herman and Lara-Steidel (2025), this article places the ability to explain, connect, and revise reasons at the centre of education's epistemic aims. Factual knowledge remains necessary, because evaluating reasons requires a command of the subject matter.

Put simply, learning aims determine which activities students need to go through. If the purpose of a task is to build an argument, handing the choice of premises and the weighing of evidence entirely to AI may strip away the chance to practise that very skill. If the purpose is to analyse a social problem, help with spelling may remove an obstacle without taking over the analysis. No list of activities can be declared always permitted or always forbidden without reference to aims. Each decision must make clear how the assistance relates to the capacity being developed.

Kasneci et al. (2023) give grounds for attending to verification and critical thinking when language models are used. The implication drawn here is that verification should itself become a guided learning activity. It is not enough to ask students to make sure an answer is correct without giving them access to comparison sources and relevant knowledge. In the early stages, lecturers might provide a handful of trustworthy sources along with examples of flawed reasoning. At more advanced stages, students can be asked to set standards of evidence and to explain why a particular source is fit for use. Demands for verification should grow with students' competence and with the stakes of the task.

Understanding need not be assessed through a single format. Oral explanations, annotated revisions, applying a concept to a new case, or comparing several lines of reasoning can each supply complementary evidence. Each has its limits: an oral explanation may be affected by anxiety, and process documentation does not necessarily capture the full extent of a student's thinking. For this reason, the article does not propose an AI-proof assessment. What it suggests

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

instead is the use of varied and accessible evidence so that judgements come closer to the capacities they are meant to capture.

Delegation can bring epistemic benefits as well. Help in finding alternative explanations or simplifying language may free up room for deeper questions. Yet such benefits cannot be assumed merely because a feature exists. Students need the chance to compare output with what they already know, to raise objections, and to change the direction of the discussion. The normative criterion, then, is whether engagement with reasons is sustained, not how much work is done without tools.

Ethical Perspective on Academic Responsibility

Holmes et al. (2022) widen the discussion of AI ethics in education beyond data and computation by attending to the aims and practices of education. Taking this orientation as a starting point, integrity in AI use should be tied to the reasons behind an assignment and to one's obligations toward others. Disclosing AI use supports transparency, but a disclosure alone does not establish that the content is sound. Students still need to account for the claims they make, check their references, and be careful about information that should not be shared with a system.

Responsibility cannot rest on students alone. Lecturers define learning outcomes and assessment rules, while institutions decide on access, support, and the choice of systems. Where rules are vague or access is unequal, it would be unfair to treat every difficulty in using AI as a personal failure. Chan (2023) offers a context for AI governance at the university level by linking pedagogical, governance, and operational considerations. This article reads that linkage as a reason to align what is expected of students with the support that institutions provide.

Accountability should be proportionate to the kind of activity involved. Help with sentence structure does not call for the same explanation as using AI to generate a data analysis or a decision that affects other people. When assistance touches the core reasoning, students may be asked to explain which parts they accepted, modified, or rejected, and why. The documentation required should be sufficient for educational purposes without turning a complete record of private interactions into a routine requirement. Transparency must still respect privacy and workload.

One might object that rules of this kind complicate learning and can easily become a mere formality. The objection is well founded if documentation only adds attachments that never shape feedback. Information about AI use, therefore, needs to feed into conversations about the

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

quality of reasoning. A brief note explaining how a student corrected an AI response may be worth more than a long chat history that nobody reads. The ethical value of documentation lies in making decisions answerable, not in the volume of files produced.

Fairness also requires workable options for students with limited access. Rules that quietly depend on paid services can create unequal opportunities to demonstrate ability. Alternative tasks, shared facilities, or institution-provided tools should be considered according to context. This proposal is normative; the article does not claim to have measured access gaps at any particular university. Its basis is the demand that judgements of ability should not be needlessly entangled with differences in resources.

Humanistic Perspective on Development and Participation

The humanistic orientation adopted here regards students as persons with purposes, experiences, and the potential to grow. Education does more than optimise task output; it also provides opportunities to form judgement and to recognise one's responsibility toward others. The guidance by Miao and Holmes (2023) emphasises protecting human agency, ensuring inclusion, and keeping the use of generative AI pedagogically appropriate. This guidance is treated as a normative policy source, not as proof that its principles have been put into practice in every institution.

Human-centred approaches to technology can support this orientation, although the two are not the same. Shneiderman (2020) separates the level of automation from the level of human control, so the two need not be seen as mutually exclusive choices. This idea is helpful for questioning the assumption that any increase in assistance must reduce agency. Interface control, however, is not enough to secure personal development. Students also need to understand their options, to have room to voice objections, and to be part of educational relationships that encourage reflection.

In their review of learning analytics and AI in education, Alfredo et al. (2024) point to the limited involvement of end users in the design of the systems they studied. This gives reason not to take the label human-centred as a guarantee of participation. In the context of this article, students should be involved when universities interpret the purpose of AI assistance, decide on forms of support, and evaluate learning experiences. Participation does not remove lecturers' professional responsibility, but it brings in information and considerations that are not always visible from an administrator's vantage point.

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

Relationships with lecturers and peers keep their value because education involves encountering the reasons and experiences of others. AI-generated feedback can serve as material for discussion rather than as a final verdict on a student's ability or identity. Students who receive automated assessments should have a route to request an explanation and a human review whenever such systems are used in processes that affect them. This follows normatively from respect for persons; it is not a technical claim about what any particular system can do.

A humanistic lens also changes how accessibility is viewed. Restricting all assistance for the sake of independence may overlook students who need language support or specific aids in order to take part at all. The better question is whether the assistance enables students to carry out the intellectual activity that the task aims at, with fair accommodation. Agency can be strengthened through reliance that is structured and accountable. Protecting agency, in other words, need not mean an obligation to work entirely alone.

A Framework for Judging Justified Delegation

Drawing on the three perspectives, this article proposes five sequential questions. First, what capacity is the task meant to develop? Second, which activity is being handed to AI, and is that activity itself a learning target? Third, how can students check the output in a way that matches their level of competence? Fourth, who is responsible for the resulting decisions and their consequences? Fifth, do students have the access and the opportunity to influence how AI use is regulated? These questions are tools for deliberation, not a psychometric scale with a pass mark.

Consider an argumentative essay. A student might use AI to suggest objections to an initial thesis. Such use is defensible if the student evaluates the objections, checks supporting sources, and explains how the argument was revised. By contrast, if the student hands over the entire choice of thesis and evidence and then merely swaps out a few words, the link between the task and its learning aim becomes weak. This illustration does not automatically establish misconduct; academic decisions still depend on clear task rules and a fair examination of the evidence.

In a task that assesses conceptual analysis, help with language editing may be justified when linguistic barriers prevent a student from showing what they understand. If the course is instead assessing editing skills, the same assistance should be limited or used differently. This shift in judgement is not inconsistency; it follows from a difference in aims. The relevant unit of

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

judgement is thus the relationship among student, task, aim, and assistance, rather than the technology taken in isolation.

A further objection holds that human thinking is always distributed across tools and communities, so insisting that certain intellectual activities be retained is too individualistic. This article accepts that learning is relational. Yet when a university awards a qualification to a particular person, it still needs grounds for attributing achievement to that student. Attribution does not require that every piece of work be done alone. It does require a fair opportunity to show one's contribution, understanding, and responsibility within work supported by others.

In practical terms, learning outcomes, permissions for AI use, available support, and assessment evidence need to be aligned. Lecturers can explain, before a task begins, what kinds of assistance are allowed and why certain limits apply. Students then receive feedback on both the content of their work and the reasoning behind their choices. Institutions support this alignment through adequate access, professional development for teaching staff, and procedures for reviewing decisions. The details of implementation will need to be adapted to each discipline and evaluated empirically, which is why the framework is not offered as a one-size-fits-all policy package.

Limitations and Directions for Future Research

A purposive corpus of twelve sources cannot represent the whole of research on educational AI published between 2020 and 2025. The availability of full texts and the predominance of English-language publications may narrow the range of perspectives. Some sources address educational AI in general, whereas others deal specifically with generative AI; this difference has been kept in view throughout the interpretation. The study also did not comprehensively survey classical philosophical works outside the set period, so the use of the term humanistic is confined to the working sense explained earlier.

The proposed framework has yet to be tested against students' experiences or learning outcomes. Future studies could explore how students explain their AI-assisted decisions, how lecturers judge evidence of understanding, and how access shapes participation. Comparative studies across disciplines might examine whether the boundaries of acceptable delegation shift with the nature of knowledge and professional responsibility in each field. Evaluations of this kind should distinguish product quality, demonstrable understanding, and the experience of autonomy so that the three are not treated as a single measure.

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

CONCLUSION

In the era of generative AI, the aims of higher education can be reconstructed around a form of student agency that brings together understanding, responsibility, and personal development. Epistemologically, academic output should be accompanied by evidence that students can explain and revise their reasoning. Ethically, delegation calls for proportionate accountability and a fair division of obligations among students, lecturers, and institutions. From a humanistic standpoint, the use of technology should open up opportunities for participation and maintain respect for students as developing persons.

The answer to the research question lies in the alignment between learning aims and forms of delegation. AI assistance is justified when it supports access and learning activities while still allowing students to demonstrate the capacities being assessed. The five considerations put forward here (aims, delegated activities, verification, accountability, and participation) provide a basis for judging that alignment. The article's contribution is a philosophical argument and a deliberative framework, not a demonstration of impact or a validated assessment instrument. Empirical testing and contextual adaptation are still needed before the framework is adopted as operational policy.

REFERENCES

Alfredo, R., Echeverria, V., Jin, Y., Yan, L., Swiecki, Z., Gašević, D., & Martinez-Maldonado, R. (2024). Human-centred learning analytics and AI in education: A systematic literature review. *Computers and Education: Artificial Intelligence, 6*, 100215. https://doi.org/10.1016/j.caeai.2024.100215

Bond, M., Khosravi, H., De Laat, M., Bergdahl, N., Negrea, V., Oxley, E., Pham, P., Chong, S. W., & Siemens, G. (2024). A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour. International Journal of Educational Technology in Higher Education, 21, 4. https://doi.org/10.1186/s41239-023-00436-z

Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. International Journal of Educational Technology in Higher Education, 20, 38. https://doi.org/10.1186/s41239-023-00408-3

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, 43. https://doi.org/10.1186/s41239-023-00411-8

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

Herman, J., & Lara-Steidel, H. (2025). Artificial intelligence on campus: Revisiting understanding as an aim of higher education. Educational Theory, 75(4), 603–625. https://doi.org/10.1111/edth.70026

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Buckingham Shum, S., Santos, O. C., Rodrigo, M. T., Cukurova, M., Bittencourt, I. I., & Koedinger, K. R. (2022). Ethics of AI in education: Towards a community-wide framework. International Journal of Artificial Intelligence in Education, 32, 504–526. https://doi.org/10.1007/s40593-021-00239-1

Holmes, W., & Tuomi, I. (2022). State of the art and practice in AI in education. European Journal of Education, 57(4), 542–570. https://doi.org/10.1111/ejed.12533

Jaakkola, E. (2020). Designing conceptual articles: Four approaches. AMS Review, 10, 18–26. https://doi.org/10.1007/s13162-020-00161-0

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., … Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, 102274. https://doi.org/10.1016/j.lindif.2023.102274

Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. https://doi.org/10.54675/EWZM9535

Mouta, A., Pinto-Llorente, A. M., & Torrecilla-Sánchez, E. M. (2024). Uncovering blind spots in education ethics: Insights from a systematic literature review on artificial intelligence in education. International Journal of Artificial Intelligence in Education, 34, 1166–1205. https://doi.org/10.1007/s40593-023-00384-9

Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human–Computer Interaction, 36(6), 495–504. https://doi.org/10.1080/10447318.2020.1741118

Copyright at authors some right reserved this work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License