ARTIFICIAL INTELLIGENCE AND THE RISK OF COGNITIVE DORMANCY IN EDUCATION: A PHILOSOPHICAL CRITIQUE OF TOTAL MACHINE DEPENDENCE

Dr. Oyemade Mark OBAFEMI
Department: Arts & Social Science Education
Institution: Lead City University, Ibadan, Nigeria
Email: oyemademark2@gmail.com
Phone: 0803 638 4275

Kawthar Sulaiman JAMIU
Department: Educational Management
Institution: University of Abuja, Abuja, Nigeria
Email: kawtharsjaishat@gmail.com
Phone: 08036979916

Oyenike Ayobami OYEBIYI
Department: Educational Management
Institution: University of Ilorin, Ilorin, Nigeria
Email: oyebiyioyenike@gmail.com
Phone: 0805 233 2864

Abstract

This paper examines how habitual dependence on artificial intelligence (AI) may affect learners' capacity for critical thinking, independent reasoning and unaided problem-solving. It calls this potential condition cognitive dormancy. Drawing on Aristotelian virtue epistemology, Kantian theories of autonomy and contemporary philosophy of mind, the paper argues that uncritical adoption of AI can replace rather than support human cognition, thereby undermining important educational aims. It distinguishes cognitive dormancy from cognitive offloading, intellectual laziness and cognitive atrophy, then examines the epistemic status of AI-mediated knowledge, the educational value of productive struggle and the ethics of outsourcing intellectual work to non-conscious systems. The paper addresses two objections: that AI can democratise access to cognitive development, and that AI-assisted thinking extends rather than threatens the mind. It concludes by advocating mindful integration of AI-an approach intended to preserve learners' cognitive agency while making productive use of AI tools.

Keywords: artificial intelligence, cognitive dormancy, education, epistemic autonomy, machine dependence, philosophy of mind

Introduction

Major tools have repeatedly changed what people expect of their own minds. The abacus, printing press, calculator, and internet expanded how people calculate, store, and retrieve information,

while also raising concerns about their effects on human capacities. Artificial intelligence (AI) in education is sometimes presented as the latest example of this pattern. However, its ability to generate and organise content raises distinct questions about how learners engage with knowledge.

Large language models, adaptive tutoring platforms, automated grading tools and recommendation systems are increasingly used in education (Holmes et al., 2022; Luckin et al., 2016). Advocates argue that these tools can extend instructional support to under-resourced schools and assist learners with needs such as dyslexia. These potential benefits deserve consideration alongside questions about how AI affects learning.

The history of educational technology offers reasons to examine such claims carefully. Cuban (2001) and Reich (2020) discuss earlier predictions that motion pictures, radio, television, computers and the internet would transform schooling. Thomas Edison, for example, predicted in 1913 that books would soon become obsolete in public schools as students learned through motion pictures. Such forecasts did not fully account for the complexity of teaching and learning. The relevant question is therefore not only what AI can do for education, but also how its use shapes learners' cognitive engagement. Contemporary reporting from 1913 supports the historical account of Edison's prediction.

This paper proposes the term *cognitive dormancy* for a possible decline in learners' active cognitive engagement when they habitually rely on AI instead of thinking through tasks themselves. It distinguishes this concept from cognitive offloading, which involves transferring a task to an external resource and does not necessarily imply a decline in ability; intellectual laziness, which describes a disposition or choice; and cognitive atrophy, which refers to the loss of a previously developed capacity. Cognitive dormancy, as used here, refers to a capacity that remains latent because it is insufficiently exercised or developed. The paper treats this as a conceptual proposal requiring further examination, rather than an established empirical diagnosis.

The paper uses a conceptual and normative method. It defines cognitive dormancy, considers the concept through virtue epistemology, Kantian autonomy, the extended mind thesis and care ethics, and evaluates its implications for education. It does not claim to replace empirical research on learning outcomes. Where empirical findings are available, they are considered as evidence relevant to the argument; philosophical analysis is not treated as a substitute for such evidence.

The empirical findings on AI and cognition are mixed. Gerlich's (2025) mixed-methods study of 666 participants reported a negative association between frequent AI tool use and critical-thinking scores, with cognitive offloading as a mediating factor. Younger participants reported more frequent AI use and offloading and had lower critical-thinking scores, although these findings do not establish that AI use caused lower scores. Gerlich's study is observational in its quantitative component. Essel et al. (2024), by contrast, reported gains in students' critical, reflective and creative thinking following structured ChatGPT-supported activities in a university setting. Their study suggests that the way AI is incorporated into learning may matter. Together, these studies motivate a closer examination of the conditions under which AI supports or displaces learners' cognitive effort.

The paper is organised as follows. Section 2 develops the concept of cognitive dormancy and distinguishes it from related concepts. Section 3 examines AI-mediated knowledge and its epistemic limits. Section 4 considers productive struggle as an educational virtue. Section 5 examines the ethics of outsourcing intellectual labour, including the justice-based case for AI adoption. Section 6 proposes mindful integration of AI in education. The conclusion considers what these arguments imply for education's purpose.

Cognitive Dormancy: A Philosophical Account

Defining Cognitive Dormancy and Related Concepts

This paper uses *cognitive dormancy* to describe the underdevelopment of capacities such as critical analysis, creative synthesis, analogical reasoning, epistemic self-correction and intellectual perseverance when machine-mediated processes habitually perform the work through which those capacities might otherwise develop. The term differs from three related concepts often conflated in discussions of AI and cognition.

Cognitive offloading is the delegation of cognitive tasks to external resources, such as notebooks, calculators or AI systems. It is value-neutral and not inherently harmful; people have long used external aids to support memory and reasoning (Clark & Chalmers, 1998). This paper focuses on a narrower concern: offloading that occurs before a capacity has developed, or that displaces opportunities to practise it so extensively that it may not develop robustly.

Intellectual laziness refers to an unwillingness to exert cognitive effort when it is appropriate. It describes a disposition or choice, rather than a developmental outcome. It can also moralise the issue by placing responsibility on the learner while overlooking the tools, incentives and learning environments that shape their behaviour. *Cognitive dormancy*, by contrast, directs attention to conditions that may prevent a capacity from developing, regardless of a learner's effort or intention.

Cognitive atrophy refers to the weakening of a previously developed capacity that is then lost or diminished through disuse. Dormancy is intended to include both the underdevelopment of a capacity and the reduced exercise of one that remains latent. A student who has consistently relied on AI to construct arguments, for example, may not have lost an established ability; the ability may not yet have developed through practice. Aristotle's distinction between possessing a capacity potentially and actualising it through habituated practice helps clarify this point. In the *Nicomachean Ethics*, his account of *hexis* distinguishes mere potential from a disposition developed through practice (Aristotle, trans. 1998, Book II). On this account, a capacity that is systematically bypassed may remain underdeveloped in a young learner or become less exercised in an adult. This distinction also affects the appropriate response: support may need to cultivate an undeveloped capacity, restore a weakened one, or both.

Historical Philosophical Antecedents

Concerns about technology's effects on cognition predate AI. In Plato's *Phaedrus*, the Egyptian god Theuth presents writing as an invention, while King Thamus warns that reliance on written words could weaken memory. People, he suggests, may trust external marks rather than remember for themselves, acquiring the appearance rather than the substance of wisdom (Plato, trans. 2002, 275a–275b). The analogy with AI is not exact, but it raises a relevant concern: external signs of learning, such as submitted work or high grades, may not always reflect the development of the understanding they are meant to demonstrate.

Kant offers a second point of reference. In "An Answer to the Question: What Is Enlightenment?" he describes enlightenment as emergence from "self-incurred immaturity," understood as an inability to use one's understanding without another's guidance (Kant, trans. 1996, p. 17). His injunction *sapere aude*-dare to know-places the independent use of reason at the centre of intellectual autonomy. Habitual reliance on an external authority, including an AI system, may

limit opportunities to practise independent judgement. In educational settings, however, such reliance should not automatically be attributed to a learner's choice: institutional expectations and learning-environment design may also encourage it.

Cognitive Dormancy and the Extended Mind

Any account of AI and cognition must consider the extended mind thesis. Clark and Chalmers (1998) argue that cognitive processes can extend beyond the brain into external resources, such as a notebook for storing information or a calculator for performing arithmetic. From this perspective, a tool may be part of a cognitive process rather than merely an external aid. This raises the question of whether using an AI tutor extends cognition or constitutes a distinct educational concern.

This paper argues that the extended mind thesis does not, by itself, resolve the developmental question. An adult using a calculator may extend an established mathematical ability. A learner who has not yet developed a foundational capacity, however, may not be extending that capacity when an AI system performs the relevant reasoning on their behalf. Learning to acquire a capacity and using a tool to extend an existing one are related but distinct questions.

There is also a normative question. Even if AI-assisted activity counts as cognitive extension in some cases, that does not mean every form of extension is equally beneficial for a developing learner. Using AI to support practice may differ educationally from using it to replace the practice through which a capacity develops. The central question for this paper is therefore not only whether AI use can count as cognition, but also which forms of use support capacity development and which may displace it. The extended mind thesis helps frame that question but does not settle it.

The Epistemic Status of AI-Mediated Knowledge

Knowledge, Understanding and Information

Discussions of AI in education can blur distinctions among information, knowledge and understanding. Information consists of content conveyed through a medium, such as the value of pi, the date of a battle or a chemical formula. In many classical accounts, knowledge involves more than true belief: it also requires justification or a reasoned account that connects a belief to evidence (Plato, *Meno* and *Theaetetus*). Understanding involves grasping relationships among ideas, applying knowledge to unfamiliar situations and recognising the limits of one's own knowledge.

Current AI systems can generate and organise information and produce text that resembles a knowledge claim (Holmes et al., 2022). Whether they possess understanding in the human sense is more contested. This paper treats understanding as involving a learner's active engagement with a subject, including working through confusion and developing insight. When a student receives an AI-generated explanation of Kantian ethics, for example, the explanation may convey information and possible reasons. Whether the student comes to understand the material depends partly on what they do with it: question it, test it against counterexamples and connect it to what they already know. An explanation's fluent or authoritative tone does not, by itself, establish the learner's understanding.

Testimony and Epistemic Dependence

Much of what people know comes from testimony rather than personal verification. People routinely rely on others for information, such as the location of a country's capital. Testimony is not necessarily an inferior source of knowledge; it is essential to social learning. However, epistemologists have debated when reliance on another source is justified and what grounds that reliance requires.

C. A. J. Coady (1992) discusses this debate through the "reduction problem." Reductionists, following Hume, argue that testimonial beliefs require independent grounds for trusting the source. Anti-reductionists maintain that testimony can be trusted by default unless there is a reason to doubt it. AI-generated content complicates this distinction. It may be fluent and coherent while still containing errors, biases or superficial reasoning that are difficult to identify from the output alone. A student who lacks the knowledge or critical skills to evaluate an AI-generated claim may therefore be poorly equipped to assess its reliability.

Miranda Fricker's account of epistemic injustice concerns credibility wrongs between people (Fricker, 2007). Applying the concept to a person's relationship with their own judgement extends Fricker's framework beyond its original scope. Still, the analogy raises a concern: repeated deference to AI output may lead learners to undervalue their own developing judgement. This paper considers that possibility as a form of epistemic harm relevant to education, while recognising that it uses the concept by analogy rather than as a direct application of Fricker's account.

The Phenomenology of Understanding

Michael Polanyi's distinction between explicit and tacit knowledge is also relevant. Explicit knowledge can be stated in propositions, whereas tacit knowledge is embodied in practice and may be difficult to articulate fully. Polanyi expresses this distinction in the claim that "we can know more than we can tell" (Polanyi, 1966, p. 4). Mathematical intuition, literary judgement, and historical understanding may involve more than possessing correct propositions; they can develop through sustained engagement with a subject.

AI can explain concepts and support practice, but an explanation alone cannot give a learner the experience and practical judgement that develop through their own engagement. These capacities may require repeated attempts, feedback and correction. If AI routinely performs the analysis, synthesis or judgement that learners need to practise, they may receive correct information while having fewer opportunities to develop understanding. Their performance on tests of factual knowledge may not, by itself, show whether that deeper learning has occurred.

Productive Struggle as a Pedagogical Virtue

The Value of Difficulty in Learning

The growing use of AI in classrooms makes it important to examine the educational value of *productive struggle*: sustained engagement with a challenging task through which learners develop understanding. Research on "desirable difficulties" suggests that some conditions that make learning feel harder in the short term can support more durable learning. Research on productive failure likewise suggests that, under structured conditions, attempting a problem before receiving instruction can support conceptual understanding and transfer (Bjork & Bjork, 2011; Kapur, 2016). These findings do not mean that all difficulty is beneficial. Struggle is productive when it is appropriately designed and followed by guidance that helps learners make sense of their attempts.

The value of struggle may also extend beyond its effects on retention or performance. In *Democracy and Education*, John Dewey describes learning as a process of encountering a problem, forming a hypothesis, testing it and reflecting on the result (Dewey, 1916). When students receive an answer without engaging in this process, they may complete a task without developing the reasoning that the task was meant to cultivate.

Hannah Arendt distinguishes thinking from knowing in *The Life of the Mind* (Arendt, 1978). Thinking involves active reflection, while knowing can involve possessing an answer or fact. An educational environment that prioritises immediate answers may encourage recall without sufficient opportunity for reflection. If AI tools routinely provide answers before learners have worked through a question, they may reduce opportunities for the kind of thinking Arendt describes.

Virtue Epistemology and Intellectual Character

Virtue epistemologists such as Ernest Sosa, Linda Zagzebski and Jason Baehr argue that epistemology should consider not only whether beliefs are justified and true, but also the intellectual traits that contribute to good knowing (Sosa, 1991; Zagzebski, 1996; Baehr, 2011). These traits can include intellectual courage, humility, open-mindedness, perseverance and epistemic autonomy.

Such traits develop through opportunities to practise them. If learners routinely delegate analysis, argument construction and self-correction to AI, they may have fewer opportunities to exercise these capacities. The concern, then, is not only whether students arrive at correct answers, but whether they develop the intellectual habits involved in reaching and evaluating them. Zagzebski's account of the "admirable epistemic agent" emphasises that intellectual evaluation concerns the way a person arrives at a belief, not only whether the belief is true (Zagzebski, 1996). This paper argues that education should consider these broader aims alongside conventional measures of achievement.

The Paradox of Frictionless Learning

Educational technology is often presented as a way to make learning seamless and reduce friction. Yet confusion and difficulty can play a role in learning when they are manageable and supported. Socrates' use of *elenchus*, or questioning, illustrates how exposing uncertainty can prompt further inquiry. An AI system designed to resolve confusion immediately may therefore remove opportunities for reflection if learners receive explanations before attempting to reason through a problem.

Nel Noddings' ethics of care also suggests that supporting learners does not always mean shielding them from difficulty (Noddings, 1984, 2003). An educator's care for a student's development may

sometimes involve allowing the student to work through a challenge while remaining available to guide and support. Accordingly, an educational tool optimised only for immediate satisfaction or ease may not serve longer-term learning. The aim should not be to preserve difficulty for its own sake, but to distinguish useful challenge from unnecessary frustration.

Ethical Dimensions of Intellectual Outsourcing

Autonomy, Dignity and the Educated Self

Kant's account of moral autonomy-self-governance through reason-has an educational parallel in the development of learners who can form judgements and reach conclusions through their own reasoning. From this perspective, habitual reliance on AI to perform intellectual work may limit opportunities to develop autonomy. This raises an ethical concern because education has a responsibility to support learners as developing rational agents.

Joseph Raz distinguishes the capacity for autonomy from the conditions that allow people to exercise it (Raz, 1986). Autonomy requires more than freedom from coercion; it also depends on developing the skills and self-understanding needed to make choices. An educational environment that routinely replaces learners' reasoning with machine-generated responses may weaken these conditions, even without direct coercion.

The Ethics of Intellectual Labour

Marx's account of alienated labour offers another lens for considering intellectual outsourcing. In his analysis, workers can become separated from the products and processes of their labour and from the exercise of their creative capacities (Marx, trans. 1988). By analogy, learners who habitually outsource the reasoning involved in learning may become detached from the process through which they build understanding and develop judgement.

This argument does not treat effort as inherently virtuous. Rather, it holds that using one's cognitive capacities is part of the learning through which those capacities develop. A craftsperson who delegates every stage of a task may produce an object without developing the judgement gained through practice. Arendt's distinction between work, which produces durable objects, and labour, which produces things consumed in the process (Arendt, 1978), offers a further analogy. Learning focused primarily on generating, submitting and discarding AI-produced work may leave

less durable intellectual development. This is a concern to be examined, not an inevitable consequence of AI use.

Justice, Access and the Democratisation Argument

A serious ethical discussion of AI in education must consider the argument that AI can broaden access to learning. Reich (2020) and Luckin et al. (2016) discuss how AI tools may extend educational support beyond institutions with abundant resources. In its strongest form, this argument holds that AI can provide valuable learning opportunities to students who lack access to expert teachers, well-stocked libraries, small seminars or individual mentorship. A defence of productive struggle must therefore avoid romanticising unequal access to educational support or treating hardship as a virtue.

The access argument does not, however, settle how AI should be used. Cuban (2001) and Selwyn (2019) describe patterns in which educational technologies are adopted differently across institutions and socioeconomic contexts. This paper proposes the term *substitution asymmetry* for the possibility that well-resourced institutions may use AI to supplement strong teaching. In contrast, under-resourced institutions may use it to replace instructional support they lack. This proposal requires direct empirical testing. If the pattern holds, it could lead to a two-tier system: some learners would use AI alongside teaching that develops independent reasoning. In contrast, others would receive AI in place of those learning opportunities. Such an outcome would fall short of the equity that the democratisation argument seeks to advance.

A commitment to educational justice therefore requires more than equal access to AI tools. It also requires attention to how those tools shape learning in different institutional settings. Where AI risks replacing rather than supporting instruction, safeguards and human support are especially important. The aim is to make AI an aid to learning while preserving opportunities for learners to practise the reasoning and judgement education is meant to develop.

Towards Mindful Integration: A Philosophical Proposal

Principles for Philosophically Grounded AI Integration

This argument does not call for excluding AI from education. AI tools can provide personalised feedback, accessibility support, information retrieval and adaptive scaffolding. When used

deliberately, they may support learning rather than displace it (Holmes et al., 2022). The following principles draw on the preceding philosophical discussion and broadly complement practical guidance for designing AI-assisted assignments (Mollick & Mollick, 2023).

1. **Cognitive subsidiarity:** AI should perform a cognitive task when the learner cannot yet perform it independently, and the task itself is not a learning objective. For example, AI might help locate sources for a student's analysis, but it should not replace the analysis when the goal is to develop analytical skill.

2. **Transparent limitations:** Students should be taught about AI's limitations, including its tendency to produce errors and its lack of human understanding. This can help them treat AI as a tool whose output requires evaluation rather than as an unquestionable authority.

3. **Productive discomfort:** Teaching should preserve a level of challenge appropriate to the learner's development. AI can serve as a scaffold that helps a learner work beyond what they can currently do alone, while still leaving them meaningful intellectual work to complete (Vygotsky, 1978).

4. **Epistemic humility before the machine:** Students should practise critically evaluating AI-generated claims, just as they should evaluate information from other sources. The purpose is not to reject AI but to exercise the critical judgement education seeks to develop.

These principles align with emerging research, but the available evidence does not show that guided AI use always produces gains or that unstructured use necessarily causes harm. Essel et al. (2024) reported gains following structured ChatGPT-supported activities in a university setting, while Gerlich (2025) found an association between frequent AI use and lower critical-thinking scores. The studies differ in design and do not directly compare the four principles proposed here.

The Role of the Teacher in an AI-Integrated Environment

Philosophy of education often emphasises the relational dimensions of learning. Noddings's ethics of care, Buber's account of the I–Thou relation and Freire's dialogical pedagogy each stress the importance of recognising learners as persons rather than treating them solely as objects of instruction (Noddings, 1984; Buber, 1970; Freire, 1970).

If AI takes on routine instructional tasks, the teacher's role may shift to supporting relationships and development. Teachers can model intellectual virtues, sustain students' motivation through challenging work and respond to learners as individuals (Williamson et al., 2020). A concern, therefore, is not only whether AI could replace particular teaching tasks, but whether institutions might use it to justify reducing the human support learners need. Selwyn (2022) raises this concern in relation to pressures to reduce costs and scale. The extent of this risk depends on how institutions choose to implement AI.

Reconceiving the Purpose of Education in the Age of AI

If AI can perform tasks traditionally developed through education-such as retrieving information, conducting routine analysis or drafting an initial response-the reasons for teaching these capacities need to be considered beyond their economic value. Aristotle's concept of *eudaimonia*, or human flourishing through the exercise of one's capacities, offers one way to frame this purpose (Aristotle, trans. 1998). From this perspective, education develops reason, inquiry and understanding as parts of a flourishing life, rather than valuing them only for the answers or assignments they produce.

This perspective changes the question institutions should ask about a particular use of AI. Whether AI can perform a task is often less important than whether its use supports students' development or substitutes for the learning they need to undertake. An institution focused on human flourishing may answer this question differently from one focused primarily on efficiency, throughput or credentialing. This paper therefore argues for evaluating AI use in light of education's purposes, rather than relying on a single, general policy for all uses (Biesta, 2022; Stommel, 2023).

Conclusion: Thinking Against the Machine

AI is now part of educational practice, but its presence is not, by itself, the problem. The central question is whether its use supports education's purposes or displaces them. This paper proposed *cognitive dormancy* to describe the underdevelopment or limited exercise of intellectual capacities-not permanent incapacity, laziness or simply skill loss. Such underdevelopment may not show up in conventional measures such as grades, completion rates, and credentials, which do not necessarily show how much independent reasoning or inquiry a learner has developed (Biesta, 2010; Zhao, 2018). The paper considered three challenges to its argument: the extended mind

thesis, the claim that AI can democratise access to learning, and the appeal of frictionless instruction. Each raises important points, but none settles whether particular uses of AI support learners' development. The four proposed principles-cognitive subsidiarity, transparent limitations, productive discomfort and epistemic humility-are guides for judgement, not a fixed procedure; applying them requires attention to the learner, the task and the institutional context. To "think against the machine" is not to reject AI, but to question the assumption that every difficulty in learning should be removed. Institutions can benefit from AI while preserving opportunities for learners to reason, question and develop judgement. If efficiency becomes the primary measure of success, educational activity may appear productive even as learners have fewer opportunities to develop the capacities that learning is meant to cultivate.

References

Arendt, H. (1978). *The life of the mind*. Harcourt Brace Jovanovich.

Aristotle. (1998). *Nicomachean ethics* (D. Ross, Trans.). Oxford University Press. (Original work published ca. 350 BCE)

Baehr, J. (2011). *The inquiring mind: On intellectual virtues and virtue epistemology*. Oxford University Press.

Biesta, G. J. J. (2010). *Good education in an age of measurement: Ethics, politics, democracy*. Paradigm Publishers.

Biesta, G. J. J. (2022). *World-centred education: A view for the present*. Routledge.

Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R. Pomerantz (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society* (pp. 56–64). Worth Publishers.

Buber, M. (1970). *I and thou* (W. Kaufmann, Trans.). Scribner. (Original work published 1923)

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19. https://doi.org/10.1093/analys/58.1.7

Coady, C. A. J. (1992). *Testimony: A philosophical study*. Oxford University Press.

Cuban, L. (2001). *Oversold and underused: Computers in the classroom*. Harvard University Press.

Dewey, J. (1916). *Democracy and education: An introduction to the philosophy of education*. Macmillan.

Essel, H. B., Vlachopoulos, D., Essuman, A. B., & Amankwa, J. O. (2024). ChatGPT effects on cognitive skills of undergraduate students: Receiving instant responses from AI-based conversational large language models (LLMs). *Computers & Education: Artificial Intelligence, 6*, Article 100198. https://doi.org/10.1016/j.caeai.2023.100198

Freire, P. (1970). *Pedagogy of the oppressed* (M. B. Ramos, Trans.). Herder and Herder.

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies, 15*(1), Article 6. https://doi.org/10.3390/soc15010006

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S. B., 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*(3), 504–526. https://doi.org/10.1007/s40593-021-00239-1

Kant, I. (1996). An answer to the question: What is enlightenment? (M. J. Gregor, Trans.). In M. J. Gregor (Ed.), *Practical philosophy* (pp. 11–22). Cambridge University Press. (Original work published 1784)

Kapur, M. (2016). Examining productive failure, productive success, unproductive failure, and unproductive success in learning. *Educational Psychologist, 51*(2), 289–299. https://doi.org/10.1080/00461520.2016.1155457

Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). *Intelligence unleashed: An argument for AI in education*. Pearson.

Marx, K. (1988). *Economic and philosophic manuscripts of 1844* (M. Milligan, Trans.). Prometheus Books. (Original work written 1844)

Mollick, E. R., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts. *SSRN Working Paper*. https://doi.org/10.2139/ssrn.4475995

Noddings, N. (1984). *Caring: A feminine approach to ethics and moral education*. University of California Press.

Noddings, N. (2003). *Happiness and education*. Cambridge University Press.

Plato. (2002). *Phaedrus* (R. Waterfield, Trans.). Oxford University Press. (Original work published ca. 370 BCE)

Polanyi, M. (1966). *The tacit dimension*. Doubleday.

Raz, J. (1986). *The morality of freedom*. Oxford University Press.

Reich, J. (2020). *Failure to disrupt: Why technology alone can't transform education*. Harvard University Press.

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

Selwyn, N. (2022). The future of AI and education: Some cautionary notes. *European Journal of Education, 57*(4), 620–631. https://doi.org/10.1111/ejed.12532

Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge University Press.

Stommel, J. (2023). *Undoing the grade: Why we grade, and how to stop*. Hybrid Pedagogy.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Harvard University Press.

Williamson, B., Bayne, S., & Shay, S. (2020). The datafication of teaching in higher education: Critical issues and perspectives. *Teaching in Higher Education, 25*(4), 351–365. https://doi.org/10.1080/13562517.2020.1748811

Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.

Zhao, Y. (2018). *What works may hurt: Side effects in education*. Teachers College Press.