

Undoing the Knower? Education, the Problem of Truth, and Artificial Intelligence

Greg Seals
College of Staten Island/CUNY
and
Mark Garrison
West Texas A&M University

Abstract: AI researchers admit to not knowing how AI works; yet the pudding is made and consumed. In the process, AI reduces knowledge and knowing to doing and saying, to stating results of performance in accord with programmed algorithms. Our paper aims to counter reductionist epistemologies of the sort presented by AI. Conducting an epistemologically focused content analysis of a purposive sample of AI-related news and commentary, we conclude that neither metaphysics nor epistemology supports reduction of knowing to doing. Both doing and understanding are required for an adequate account of knowledge and an optimized capacity to work wisely in the world. Education requires an ontology of ideas *and* things. AI fails for the same reason metaphysical idealisms typically fail: Reason is not all of reality.

What are educators to make of the growing push for adoption of artificial intelligence (AI) technologies? While the AI industry talks endlessly of “innovation,” growing concern about AI in education is also evident. Education technology scholar Neil Selwyn identified five themes that should be attended to as educators are pressured to adopt the new, untested technologies:

1. Pay less attention to speculative AI technologies in favor of exciting AI.
2. Admit limited ability to model social contexts/simulate human emotions.
3. Taking seriously social harms associated with AI use.
4. Acknowledging values inherent in AI advocacy.
5. Environmental sustainability of AI development/implementation.¹

Selwyn’s concerns echo popular discourse regarding AI (academic integrity, bias, “deepfakes,” transparency). But we worry a significant issue with AI in education is being missed: a possible epistemological shift that eviscerates the knowing subject. AI can “do” (demonstrate competency) without there being a knower. Since AI “does all the doing” without any of the knowing, is AI “undoing” the knower? Is AI breaking down the distinction between performative and perspicacious knowledge? In this paper we explore the reach and depth of Selwyn’s second concern. We do this by conducting an epistemologically focused content analysis of a purposive sample of AI-related news and commentary. Four questions guide the analysis:

1. What epistemological arguments are evident?
2. What epistemological arguments are absent (but assumed)?
3. What ontological assumptions are made about teaching and learning?
4. How do ontological assumptions relate to evident/absent epistemic arguments?

¹ Neil Selwyn, “The Future of AI and Education: Some Cautionary Notes,” *European Journal of Education*, (17 October 2022), <https://doi.org/10.1111/ejed.12532>

About the Source Material

The news and commentary to which we applied our epistemically oriented content analysis comes from an *Education Week* “Spotlight” on artificial intelligence.² *Education Week*, in operation since the 1980s, is a production of Editorial Projects in Education (EPE) and, in accord with EPE editorial intent, describes itself as more than a source of news about K-12 education:

Because at *Education Week*, we don’t merely inform. We generate critical conversations that shine a spotlight on our schools. We hold education’s leaders accountable — equipping them with the right tools, so they can become instruments of change. And we spark transformation, helping our nation support its students — and build a stronger tomorrow.³

This frames the purpose of *Education Week* “Spotlights,” which are typically a collection of news articles and commentary about issues in K-12 education, such as AI. Spotlights tend toward advocacy, promoting the science of reading, specific approaches to social-emotional learning, or, as in this case, the rapid adoption of artificial intelligence in schools.⁴

Education Week receives support from major philanthropic entities. The Board that oversees operations includes executives from Discovery Education (digital learning platform), Boston Consulting Group (neoliberal school reform), Allovue (educational technology), WNET (Boston public television), Cooley LLP (mergers and acquisitions), Stanton Chase (consultancy), as well as superintendents from large school districts operated on neoliberal ideas about education. From sources such as these, EPE has received substantial and sustained funding for general operations, a variety of initiatives, and as an incentive to cover specific education topics. According to EPE’s *Education Week* website, this private funding has come from the following 25 funders.⁵ Many funders have a focus that includes technology. Listed in alphabetical order, they are:

Bill & Melinda Gates Foundation
Carnegie Corporation of New York
The Chan Zuckerberg Initiative
Charles and Lynne Schusterman Family Philanthropies
The Charles Butt Foundation
The Charles Stewart Mott Foundation
The Field Foundation of Illinois
The Jack Kent Cooke Foundation
The Joyce Foundation
The Kern Family Foundation
The Lemelson Foundation
Meyer Memorial Trust
The NEA Foundation
The Nellie Mae Education Foundation

² “Education Week Spotlight: Artificial Intelligence (AI),” *Education Week* (5 January 2024): 1-15 <https://epe.brightspotcdn.com/35/2a/cdd0047540b29d8ffd0ca897c491/1-5-24-aispotlight-sponsored.pdf>

³ See *Education Week* website, “About Us: Our Organization” <https://www.edweek.org/about/our-organization>

⁴ See, for example, <https://epe.brightspotcdn.com/8a/50/18f484b74230ac4e5639a59f472a/spotlight-science-of-reading-sponsored.pdf> and <https://epe.brightspotcdn.com/09/9f/6748386e4b71a3e65d16afee6e1c/6-15-socialemotionalllearning-sponsored.pdf>

⁵ <https://www.edweek.org/about/our-supporters>. For board membership, see <https://www.edweek.org/about/our-board#board-of-trustees>

NewSchools Venture Fund
 NoVo Foundation
 Oak Foundation
 The Robert R. McCormick Foundation
 Siegel Family Endowment
 Spencer Foundation
 The Susan Crown Exchange
 The Wallace Foundation
 The Walton Family Foundation
 William and Flora Hewlett Foundation
 William E. Simon Foundation

Considering its self-professed and well-funded role in holding educators, as quoted above, “accountable” by “equipping them with the right tools,” EPE and its spinoff endeavors should be considered first and foremost a political and ideological effort, and as we discuss here, one with significant epistemological and ontological implications.

The articles collected in the 2024 *Education Week* AI Spotlight we analyze were published between August and November 2023. Included were six “news” articles written by *Ed Week* staff followed by two “opinion” pieces, one by Stacie Marvin, a math educator turned congressional aide, and one by Rachna Nath, an award-winning science teacher with fifteen years’ experience in K-14 settings. The news articles included quotes from teachers, educational administrators, and technology company executives. As a convenience sample, clearly the contributors to the AI Spotlight we quote from reflected editorial intent to promote rapid adoption of AI. This reveals both significant contradictions and challenges faced by *Education Week*’s EPE-directed approach to AI in schools. The educators quoted in these Spotlight articles are spread across the United States from Kansas (including a fully online “academy”) to Maryland, New York, Pennsylvania, Texas, and Tennessee. The articles also quote education faculty from Stanford University and executives from organizations like Code.org and Edthena, as well as a former *Time* magazine teacher of the year gone national consultant, keynoter and grant writer.

AI: Two Models of Science

As our work aims to show that AI has implications for our most basic understanding of how knowledge is acquired and what it means to know, some initial observations are warranted. To set the stage for our analysis, an analogy may be found in the fact that science, the systematic development of a body of knowledge, has two, sometimes competing, models.⁶ The first, most common, especially in relation to technology development, focuses on “prediction and control.” Efficacious prediction about modification of the natural or social world is equated with knowledge. Call this the “doing” model of science. The proof is in the pudding. The second model stipulates science advances by way of “explanation and understanding.” Call this the “seeing” model of science. On this model, knowledge is enacted as perspicacious insight into the workings of the world. An adequate understanding of science as scientific knowledge requires an account of both doing and seeing but, importantly and to its detriment, the “doing” model is evident in AI discourse to the exclusion of the “seeing” model.

⁶ Peter T. Manicas, *A Realist Philosophy of Social Science: Explanation and Understanding* (New York: Cambridge University Press, 2006).

In his book, *Why Things Matter to People*, Sayer argues reason is reduced to rationality on the doing model of science. Reason is thus “attenuated.” This attenuation occurs when a mode of reasoning is abstracted from what the reasoning is about. This separation licenses the view that there are no limits to technical rationality and the increasing displacement of *practical reason*. This latter form of reasoning is concerned with objects and is consciously context bound. While much critique of this techno-scientific rationality comes from the Frankfurt School, locating the problem as one of *instrumental reason* (to paraphrase Adorno: the means of thought have become independent of the purposes of thought),⁷ Sayer argues a more fundamental problem exists, originating in the increasing domination of *formal reason* (that is, procedure over content). Instrumental rationality need not be formal, he points out, and it is of obvious value for achieving practical goals. While instrumental reason “is concerned with means rather than ends in themselves, it requires a close focus on the objects which it uses as means, and on the end, which it is trying to meet.” Thus, importantly, and unlike formal reason, instrumental reason does not turn “away from the object.” Formal reasoning does turn away, becomes inwardly focused, losing its object: *the logic of things becomes things of logic*. Most importantly for our work here, “when instrumental reason is also formal *as in the case of general technologies*, it becomes more dangerous and morally indiscriminate, for to the prioritization of means over ends is added the detachment of rationality from its object.”⁸

What Epistemological Arguments Are Evident?

AI is reliabilist in its epistemic orientation. Reliabilism holds that “a belief produced by a reliable belief-producing process is justified.” AI is (presumed to be) such a reliable belief-forming process. Justified belief is the result of a *performance* of a process that may or may not be associated with the giving of reasons for the belief. That is, beliefs formed by reliable processes may be justified but not necessarily true. Reliabilists attempt to skirt this epistemological issue by setting parameters to shore up the reliability of belief-forming processes.⁹ A main style of parameter setting is to set a rule that, among humans at any rate, true beliefs are, *ceteris paribus*, to be preferred to false beliefs. This truth-preference parameter has given rise to virtue or value epistemologies;¹⁰ but, whatever the usefulness of this attempt to shore up reliabilism, such moves are unavailable to AI. AI is a sort of *raw* reliabilism in the sense that not even AI engineers can say exactly how it works. It is as likely that AI operates on what we might call the willy-nilly algorithm as it is that AI results can be traced to some patterned processes of reasoning on the part of AI which resonate with human thinking about a topic. The mysterious nature of the production of AI results creates two epistemological issues related to AI, one an issue internal to AI and another issue external to AI but affecting users of AI. Thomas Kelly points out that since reliabilism lacks relevant parameters, such as a preference for true beliefs over false ones, then for a knower like AI “*nothing* would be either

⁷ See Theodor W. Adorno, “The Subject and the Object,” in *The Essential Frankfurt School Reader*, eds., Andrew Arato and Eike Gebhardt (New York: Continuum, 1985), 497-511 and Yvonne Sherratt, “Instrumental Reason’s Unreason,” *Philosophy and Social Criticism* 24, no. 4 (July, 1999): 23-42.

⁸ Andrew Sayer, *Why Things Matter to People: Social Science, Values and Ethical Life*. (Cambridge University Press, 2011), 67-69, emphasis added.

⁹ Matthew Frise, “The Reliability Problem for Reliabilism,” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 175, no. 4 (April, 2018): 923-945. See 923 for the quote.

¹⁰ For example, Ernest Sosa, “The Place of Truth in Epistemology,” in Michael DePaul and Linda Zagzebski, eds., *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford: Oxford University Press, 2002), 1-13.

epistemically rational or irrational.”¹¹ The untoward consequence of this nonrational approach to the generation of answers when questions are posed to it engages AI in a sort of epistemic pathology in which we are told, instrumentally speaking, what or how to do, but are not told, theoretically speaking, why we should believe what AI has to say.¹² By not being able to say why its answer is to be believed, AI creates a situation that undermines the reliabilist epistemology in terms of which it operates. Users may begin to second guess AI answers in much the same way that ancient peoples wondered what the oracle may have meant by saying that. When externalist processes originally proposed to be sufficient for knowledge need to be supported by processes of thought internal to AI users, internalist epistemology subverts rather than supplements reliabilism.¹³

Most, if not all, of the more common, popularly discussed concerns about AI — its potential for biased results; its propensity to reproduce rather than produce knowledge; its high probability (like any meta-analytic exercise) to include error as well as fact in answers it develops; its seeming hallucinatory relation to reality — appear traceable to its reliance on raw reliabilism.¹⁴ The epistemic crisis produced by AI is that even though we may grant (for the sake of argument) that AI is a reliable process of belief formation, in the sense that AI will give the same or a similar answer every time the same or a similar question is put to it, we cannot say that the beliefs so formed are true. This epistemic crisis contributes to an ethical crisis, as well, in two ways. First, uncritical use of AI’s reliabilist assumptions about the nature of knowledge requires socialization of teachers and students to assume/accept a limited role in understanding and explaining courses of action undertaken in the world. Unpredictable human agency surrenders to algorithmic machine certainty. This surrender raises a second ethical issue described in Gilbert’s discussion of happiness.¹⁵ On Gilbert’s view, intelligence is not what makes humans unique. Gilbert notes that all living things are intelligent in the sense that they all behave on the basis of

¹¹ Thomas Kelly. “Epistemic Rationality as Instrumental Rationality: A Critique,” *Philosophy and Phenomenological Research* 66, no. 3 (May, 2003): 612-640, 627 for the quote. Emphasis in original.

¹² Kelly. “Epistemic Rationality as Instrumental Rationality,” 635. The epistemic pathology Kelly describes cuts the other way, too. Merely knowing what to believe does not, in and of itself, tell us what to do.

¹³ Kristoffer Ahlstrom-Vij, “Esoteric Reliabilism,” *Episteme: A Journal of Individual and Social Epistemology* 18, no. 4 (2021): 603-623. The same may be said of the addition to validity checks, confidence intervals, and probabilistic reconfiguration of AI results, a problem that equally haunts results from standardized tests, etc. When responsibilist considerations are added by human interpreters of AI, standardized test results, etc., the reliabilism of the process *per se* is called into question not merely refined. For an account and assessment of the effect on our understanding of knowledge and of knowing of the probabilistic turn in multiple phases of contemporary culture see Gert Gigerenzer, Zeno Swijtink, T. Porter, Lorraine Daston, John Beatty, and Lorenz Kruger, *The Empire of Chance: How Probability Changed Science and Everyday Life* (Cambridge: Cambridge University Press, 1989).

¹⁴ Pertinent discussions include Eamon Duede, “Instruments, Agents, and Artificial Intelligence: Novel Epistemic Categories of Reliability,” *Synthese* 200, no. 491 (2022): file:///C:/Users/Owner/Downloads/Duede%202022%20-%20Instruments,%20agents,%20and%20artificial%20intelligence%20%20novel%20epistemic%20categories%20of%20reliability-1.pdf ; Hiroaki Kitano, “Nobel Turing Challenge: Creating the Engine for Scientific Discovery. *npj (Nature Partner Journal) Systems Biology and Applications* 7, no. 29 (2021): <https://doi.org/10.1038/s41540-021-00189-3> ; Lisa Messeri and M. J. Crockett, “Artificial Intelligence and Illusions of Understanding in Scientific Research,” *Nature* 627 (2024): 49–58 <https://www.nature.com/articles/s41586-024-07146-0#citeas> ; Frederica Russo, Eric Schliesser, and Jean Wagemans, “Connecting Ethics and Epistemology of AI,” *AI and Society* (17 January 2023): <https://doi.org/10.1007/s00146-022-01617-6> ; and Andrew Smart, Larry James, Ben Hutchinson, Simone Wu, and Shannon Vallor, “Why Reliabilism Is Not Enough: Epistemic and Moral Justification in Machine Learning,” *Proceedings of the AIES (Artificial Intelligence, Ethics, and Society) Conference* (2020): 372-377 <https://dl.acm.org/doi/10.1145/3375627.3375866>

¹⁵ Daniel Gilbert, *Stumbling on Happiness* (New York: Alfred A. Knopf, 2006).

reinforced predictions, whether or not they can explain or understand why they are doing what they are doing. William James makes the same point with greater panache when he says, “[C]reatures extremely low in the intellectual scale may have conception. All that is required is that they should recognize the same experience again. A polyp would be a conceptual thinker if a feeling of ‘hello! *thingumbob* again!’ ever flitted through its mind.”¹⁶ The unique achievement of humans, Gilbert asserts, is not in the things they do or make, say, the pyramids, etc. Machines can build these. Rather, conscious experience and imagined future worlds are the stuff of human achievement. AI’s raw reliabilism, its thoroughgoing externalist orientation to knowledge, taken to the extreme, threatens to ignore/degrade/deny the conscious human experience of creating our own futures. In the name of instrumentality, efficiency, and expediency, AI robs users of the core of their humanity and education of much of its *raison d’être*.

What Epistemological Arguments Are Absent (but Assumed)?

Teachers appear to endorse a robust epistemology that includes elements of doing (performance) and of seeing (perspicacity). The title of the *Education Week* Spotlight discussing AI and schooling, “What Teachers Need to Know About AI, But Don’t” (which appears in the Table of Contents and gets repeated in the title of the first article) seems somewhat epistemologically at odds with the teachers’ point of view. Use of “need” in the title makes the title ill-structured in two ways. First, the phrase “need to know” casts shades of “esoteric reliabilism,” an optimally functional but presumably undesirable reliabilist epistemic universe in which exists “a form of two-level epistemic consequentialism, where an esoteric commitment to reliabilism will be appropriate for an enlightened few, while a form of epistemic fetishism — on which some heuristics are treated as fundamental epistemic norms — is appropriate for the rest of us.”¹⁷ Second, deciding what teachers “need to know” could be done externally to teacher thinking or could arise internally from teacher thinking. On the externalist approach teachers would be *told* what they need to know to be efficient at work. On the internalist approach teachers would be *asked* what they *want* to know to become adept at educationally energetic use of AI. Taking contributors to the *Education Week* Spotlight feature on AI as a convenience sample of teachers gives us some insight into what teachers want to know about AI. Perhaps contrary to the wishes of their publisher and its sponsors, what teachers really want to know is how to use AI to help achieve among their students the epistemic ideal of learned fluency, a sort of confident familiarity with the subjects they study in school.

True, teachers do have concerns about what they do not know about AI: but these comments are at least as much declarations of curiosity as they are confessions of ignorance. *Education Week* staff writer Lauraine Langreo in the lead article of the AI Spotlight summarizes concerns expressed by Chad Towarnicki, an 8th grade English teacher in Wissahickon, Pennsylvania, that teachers “are already so far behind on AI knowledge and skill that students who have grown up

¹⁶ William James, *The Principles of Psychology*, in ed. Robert Maynard Hutchins, *Great Books of the Western World*, Volume 53 (Chicago: Encyclopaedia Britannica, Inc., 1952), 301. Emphasis in original.

¹⁷ Ahlstrom-Vij, “Esoteric Reliabilism,” 613. The teacher ignorance asserted in the title of the *Education Week* Spotlight feature is reinforced in Stacie Marvin’s (former teacher/current congressional aide) claims implying a steep, chaotic learning curve for teachers relative to AI. : “...two months of reading about AI, listening to podcasts, and even attending hearings and briefings about AI on Capitol Hill, I’ve come away with a great big “who knows?” shrug of the shoulders [a gesture indicating uncertainty, indifference, or lack of knowledge].” from p. 13 of the *Education Week* AI Spotlight, “I’m a Math Teacher: Graphing Calculators Taught Me a Thing or Two About AI: Artificial Intelligence Will Enhance Learning If We Use It Right,” 12-13. (Originally published 16 October 2023).

learning how to use all kinds of digital technologies will be running circles around them.” But the concern prompted the follow-up observation based on comments from Glenn Kleiman, whose research at the Stanford Graduate School of Education focuses on the potential of AI to enhance teaching and learning, that “It’s critical that schools of education act quickly to ensure prospective teachers have a foundational understanding of AI, know how to use it effectively in instruction, and are able to infuse AI literacy into every subject so that they’re ready when they get in the classroom.”¹⁸ In her opinion piece included in the Spotlight, Rachna Nath, a veteran science teacher, concurs with Kleiman: “And we need training on theoretical and practical levels: training to understand what artificial intelligence actually is and where it stands in the development timeline and training about how to integrate it into our classes.” Nath also professes a broad lack of understanding of AI among teachers but sees that lack as a removable obstacle to idea-sharing among members of the teaching profession: “The problem is that my fellow teachers and I cannot have those productive conversations about generative AI if we do not understand it. We need models of what is or isn’t acceptable when using this technology. At this very moment — and completely on the fly — we’re setting the standard of what is acceptable, and each one of us has a different perspective.”¹⁹ From an epistemological perspective, what emerges from uses of “know” and its cognates in analysis of the quotes above is a two-fold theory of knowledge. Knowledge has at least two parts, alternatively described as knowledge *and* skill, foundational understanding *and* know how, memorization *and* employment, responsible *and* effective use, theoretical *and* practical aspects. That is, knowledge comprises both a science of perspicacious seeing and a science of effective doing with the performance informed by the perspicacity.

Further, for teachers, knowledge comprises some standard of truth, some process of verification to meet the truth standard, and personal confidence in assertion, via word or deed, of belief. These epistemological ideas are expressed in the *Education Week* Spotlight issue on AI via quotes that express teachers’ distrust of AI. Teachers distrust information provided by AI because AI is a doing. AI is not a knowing; and, so, tells only half of the full epistemological story. Nonetheless, the doing is often mistaken for and supplanting of the knowing. What *is* to be done? For the teachers in the sample, understanding is the antidote to mere believing: “Of course, there are potential downsides to the new technology. It can produce inaccurate or biased responses based on faulty data it draws from... Experts have cautioned that when using these tools, it’s important to know how they [the AI programs in use] were trained and what datasets were used. It’s also important to be skeptical about any information these tools provide and to double-check it with a trusted source.”²⁰ At the heart of this skepticism about AI is recognition that: “There is no standard of truth on the internet, and generative AI has absolutely no ability to confirm the veracity of what it is presenting as fact.” This observation leads to the general recommendation that students may be turned into skeptical readers and set to the task of fact checking material generated by AI: “...my students are learning how to be cynical readers, perhaps better now than ever. So, my advice to teachers is to use any and all the generative AI you can get your hands on. Then experience — for yourself — verification of the information.

¹⁸ Langreo, “What Teachers Need to Know About AI, But Don’t: Students Must Be Critical AI Users, 3 Steps for Teachers to Prepare,” 1-2. Both quotes may be found on 2. (Originally published 31 August 2023.)

¹⁹ Nath, “AI Can Teach Students a Powerful Lesson about the Truth: How I’m Harnessing ChatGPT in the Classroom,” 14-15, 14 for both quotes. (Originally published 7 September 2023.)

²⁰ Lauraine Langreo, “Beyond ChatGPT: The Other AI Tools Teachers Are Using,” p. 10 of the *Education Week* AI Spotlight (Originally published 15 August 2023).

Track it back to the source because in doing so, you'll land on the adjustments you need to make in your classes next year."²¹

Without evidential support, knowledge is mere belief. One use of each of the belief cognates, 'believes' and 'believe' indicate this. Mere belief is held counter to or without consideration of fact and therefore fails to reach epistemic fulfillment as understanding: "There are so many things that a teacher does that AI would never fully be able to replicate a wonderful teacher," Edwards [April Edwards, a 6th-grade Social Studies teacher from Texas] said. Still, Edwards believes that AI is the future of education."²² Stacie Marvin expresses a similar conception of belief as epistemically under-formed knowing when she says, "Knowing most students will pull out their phones to perform even the smallest of calculations, we've had to rely on our relationships of trust to teach them to have confidence in their own abilities. We help them understand that they are more capable without the technology than they believe."²³ These comments express, perhaps unsurprisingly, teachers' tendency to think of knowledge in traditionalistic terms as justified, true belief. Beliefs are turned into knowledge by achievement of evidentiary standards matched to rigorous canons of truth. Thus, the epistemic problem presented to teachers by AI is that AI does not fulfill any of the requirements for knowledge endorsed by teachers: AI has no standard of truth; AI is incapable of offering verification of its claims; and AI has no confidence whatsoever in anything it has to say.

What teachers want for their students, namely, fluency with subject matter, is also what teachers want for themselves in learning AI. The object of knowledge is fluent use of AI for classroom instruction: "...preparing prospective teachers to integrate emerging technologies into instruction." April Edwards, who uses AI for lesson plan development, "...has not introduced AI to her students, because she wants to fully understand it before allowing students to use it in the classroom." Says Edwards, "I have not let my students use AI yet, but I plan to incorporate it this school year.... "My goal is to show them how to use AI responsibly and effectively and model that for my students."²⁴ Part of the *Education Week* Spotlight on AI is "A toolkit of principles to think through when crafting AI guidance." The toolkit includes the heading "Knowledge: How Can Schools Advance AI Literacy?"²⁵ The toolkit equates "knowledge" to "literacy." Ideally, AI literacy culminates in AI fluency. Fluency with AI gets described in some detail in a discussion about use of calculators in math class. Fluency in use of graphing calculators goes beyond mere competency in carrying out computations and expresses understanding of which calculations are relevant to skillful development of correct answers to present problems: "We've worked to define what computational 'fluency' means over and above precision of calculations. Fluency includes drawing on mathematical understanding to find flexible and efficient pathways to correct answers."²⁶ By the same token, teachers who know AI know how to use it to find solutions to pedagogical issues such as tailoring lessons to student interest and/or need. Dyane Smokorowski, coordinator of digital literacy and citizenship for Wichita Kansas Public Schools, explained to *Ed Week* staff writers that "Because she does not have knowledge on every learning

²¹ Both quotes from Nath, "AI Can Teach Students a Powerful Lesson about the Truth," 14 and 15, respectively.

²² Lauraine Langreo, Lydia McFarlane, and Caitlyn Meisner, "Can AI Improve Instruction? Teachers Share How They Use It," *Education Week* AI Spotlight, 3-4. (Originally published 10 November 2023.) Edwards is quoted on 3.

²³ Marvin, "I'm a Math Teacher," 12.

²⁴ Langreo, McFarlane, and Meisner, "Can AI Improve Instruction?" 2 and 3, respectively.

²⁵ Alyson Klein, "The 93 Questions Schools Should Ask about AI," *Education Week* AI Spotlight, 5. (Originally published 17 October 2023.)

²⁶ Marvin, "I'm a Math Teacher," 12.

difference or middle school pop-culture trend, she can ask ChatGPT to provide a writing sample for students to interact with, which makes personalized lesson plans much easier to craft.”²⁷

Another use of ‘understand’ and three uses of its cognate ‘understanding’ give insight into what teachers hope from use of AI, namely, deepening mastery of subject matter for themselves and their students, adept use of AI for finding solutions to subject matter questions, and development of ways to deal wisely with emergent technologies. For example, the *Education Week* Spotlight on AI says of Amanda Young, principal of the online Education Imagine Academy that, “In her experience . . . using AI does not harm the critical thinking abilities of students, but it can foster a deeper understanding of the new technologies that come into fruition in their own lifetimes.”²⁸ Similarly, Stacie Marvin asserts that after using AI via graphing calculators for a month “...my understanding of trigonometric functions likewise deepened as I took over the class to learn to teach students using those same graphing calculators. Suddenly, as I watched the sine and cosine waves unspool from the unit circle on the calculator’s screen, I grasped for the first time why they were called “cyclic functions.” Marvin goes on to say, “...we continue, as a community of math educators, to generally move in the direction of capitalizing on the calculator’s ability to enhance conceptual understanding.” She adds as a final thought, “When our students have access to all the information that they once spent much brain power memorizing, they are able to apply, analyze, evaluate, and create. The miracle here is that brain research tells us that when they are doing these magical things, they remember and understand *at the same time and more permanently.*”²⁹

Teacher approach to AI expresses the epistemological debate between reliabilist and responsibilists. While reliabilists see the stuff of knowledge as the result of a reliable belief-forming process, responsibilists argue that knowledge is better estimated by a look at the intellectual virtues involved in the processes involved in producing it.³⁰ Intellectual virtues typically endorsed by responsibilists include courage, open-mindedness, curiosity, creativity, and humility, and are conceived as “stable, excellent dispositions of cognitive character.”³¹ Apparent contradictions in discussions of AI for education map well onto the reliabilist-responsibilist framework: AI makes the teacher’s job easier vs. Teachers must always check AI because it gets things wrong; Students already know way more about AI and teachers better get caught up vs. Students need to prepare for a world with lots of AI; AI is empowering and a great tool vs. AI is very hard to use responsibly without lots of support to figure out what to do with it in classrooms; Students need AI because it is sure to be a super-important part of the future vs. We need to carefully police what students do with AI because of all the issues with AI.

While reliabilists may wish to assume that intellectual virtues are in play for belief-forming processes legitimately labeled “reliable,” there is no accountability in this regard when AI is the

²⁷ Langreo, McFarlane, and Meisner, “Can AI Improve Instruction?” 4. Of course, if she doesn’t know these differences or sub-cultures, she will be challenged, at least in terms of the time required, to verify AI-provided answers or, simply take their truthfulness and usefulness for granted.

²⁸ Caitlyn Meisner, “ChatGPT Is Everywhere in This District: Here’s What It Looks Like,” *Education Week* Spotlight on AI, 7-8. See 8 for the quote. (Originally published 29 August 2023.)

²⁹ Marvin, “I’m a Math Teacher,” 12 for the first two quotes and 13 for the final one, emphasis in original.

³⁰ See Lorraine Code, “Toward a ‘Responsibilist’ Epistemology,” *Philosophy and Phenomenological Research*, 45 (September, 1984): 29-50 for the original proposal of epistemic responsibilism.

³¹ Nathan L. King, “Responsibilist Virtue Epistemology: A Reply to the Situationist Challenge,” *The Philosophical Quarterly* 64, no. 255, (April, 2014): 243-253, 243 for the quote.

process guiding belief formation.³² Therefore, from the educator’s point of view, both reliabilist processes and epistemic virtues must be taught in AI classrooms. That is, the AI “fact-checking” recommended by teachers goes beyond just making sure AI got the right answer. Considerations of how AI arrived at that answer are of paramount importance. There is a great difference between merely being able to do something and knowing how to do something. Knowing how to do something involves critical thinking about the processes involved in doing things.³³ AI can do things without knowing how to do things. Teachers are not willing to let their students off so easily. Responsible use of AI is as important to teachers as is the ability to use AI to generate answers. Even if it is objected that reliabilism is more fundamental than responsibilism because responsibilism is not so much a philosophical position as it is a commitment to removing obstacles to discovery of truth, teachers’ interest in discovering truth seems firmly rooted in the undeniable responsibilist axiom: “If a belief is responsibly arrived at, it is, therefore, likely to be true.”³⁴ That is to say, in their epistemology, teachers are committed to some version of evidentialism, the epistemological position that “a person is justified in believing a proposition just in case the person’s evidence sufficiently supports the proposition and the person responsibly acquired and sustained the evidence that supports the proposition.”³⁵ This is likely part of what teachers *mean* by calling a belief-forming process “reliable.” Commitment to evidentialism carries with it the epistemological mission of showing how reliabilism and responsibilism relate to one another, but the black box of AI reliabilism makes articulating that relation a difficult task.

What Ontological Assumptions Are Made about Teaching and Learning?

That AI drives a wedge between reliabilism and responsibilism is of main concern to use of AI as an educational tool. AI’s close-mindedness about its processes, instantiated in human inability to check those processes for reliability, violates responsibilist demand for open-mindedness in the search for truth. But, perhaps more importantly, what gets closed out from consideration by AI’s solipsism is the local contexts in which AI and its recommendations will be used. A report issued in March 2024 by the National Education Policy Center (NEPC) at the University of Colorado, Boulder argues that AI risks “undermining teachers’ ability to make professional judgments about their students and devaluing teachers’ subjective experience, subject matter expertise, classroom interactions, and contextual knowledge of a given child, class, or social setting.” As AI comes to be seen as a critical factor in teaching, teaching may be reduced to a recommended set of programed steps to be taken to get curricular material across to AI numerically profiled students. AI also, the report continues, “poses a real threat to the accuracy of school knowledge and thus to the validity of curriculum materials.” Curriculum designed, developed, and disseminated using AI risks inability to distinguish merely plausible

³² Jason Baehr, “Responsibilist Virtues and the ‘Charmed Inner Circle’ of Traditional Epistemology,” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 174, no. 10 (October, 2017): 2557-2569.

³³ Heather Battaly, “Teaching Intellectual Virtues: Applying Virtue Epistemology in the Classroom,” *Teaching Philosophy* 29, no. 3 (September, 2006): 191-222; J. Adam Carter, Ben Kotzee, and Harvey Siegel, “Education for Intellectual Virtue: A Critique from Action Guidance,” *Episteme: A Journal of Individual and Social Epistemology* 18, no. 2 (2021): 177-199; and Ben Kotzee, “Learning How,” *Journal of Philosophy of Education* 50, no. 2 (May, 2016): 218-232.

³⁴ For the objection see Guy Axtell, “Expanding Epistemology: A Responsibilist Approach,” *Philosophical Papers* 37, no. 1 (March, 2008): 51-87 and for the axiom see Joel Katzav, “Riggs on Strong Justification,” *Australasian Journal of Philosophy* 76, no. 4 (December, 1998): 631-639, 635 for the quote.

³⁵ Christopher Michael Cloos, “Responsibilist Evidentialism,” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 172, no. 11 (November, 2015): 2999-3016. Find the quote on 3013.

sources from proven ones. But the biggest epistemological threat to education is that as decisions made about students along the arc of their educational histories become increasingly machine-guided, "...neither teachers nor administrators will be able to understand, explain, or justify the conclusions the programs reach, much less audit or document their validity." Disturbingly, on this scenario from an epistemic point of view, education becomes guided by ignorance, a process the NEPC report says serves to "replace human discretion with automated suspicion."³⁶ Thus, a central negative epistemic consequence of AI's close-mindedness is irresponsible revocation of human freedom to decide what to do in a particular case in favor of inscrutable machine calculation that may possibly be true. In a worst-case scenario, AI could make recommendations about a student that educators familiar with the student know to be unwarranted. Here we have a clear case of what it means to say AI's doing risks undoing human knowing.³⁷

Considerations along these lines prompted NEPC to call for a pause in use of AI while public agencies and elected officials get their act together with answers to questions like: What is acceptable use of AI? Who can use AI? What are state and district approved best-practices? What's to be done in cases of policy violations?³⁸ However, a pause in AI use and development seems unlikely. First, AI use is already rampant. Stacie Marvin, in her opinion piece in the *Education Week* Spotlight on AI sees the role of teachers changing as the educator's duties increasingly include the need to "prepare our students to thrive in a world that has AI embedded into nearly everything they are going to be asked to do."³⁹ Keith Krueger, executive director of the Consortium for School Networking (CoSN), points out that the problem of AI's virtually universal presence is particularly pressing for schools: "Now, many district leaders understand that AI is "probably in almost every product that you already purchase."⁴⁰ Second, economic

³⁶ Ben Williamson, Alex Molnar, and Faith Boninger, *Time for a Pause: Without Effective Public Oversight, AI in Schools Will Do More Harm Than Good* (Boulder, Colorado: National Education Policy Center, 2024). For the quotes see 13, 14, 15, and 16, respectively. NEPC use of the word 'suspicion' has rhetorical power in casting doubt on uncritical acceptance of AI recommendation. However, it may mislead as to what is going on in processes leading to AI output. "Thinking" or "feeling" are ill suited to describe algorithmic outputs. The outputs of AI are probabilistic in nature, and believed to be reliable on the grounds of probabilistic reasoning. On the chance they are true, we are to believe them. The calculator (i.e., the instrument, AI) which does the calculating need not nor can think or feel anything about the calculations or results of calculations. There is no feeling or thought (i.e., evaluation) anywhere in this system — feelings and beliefs would be "external" to it (or about it), but not in it. We tend to impute thought and feeling to machines and machine processes because in our own experience, thoughts and feelings are part and parcel (attendant to, and emergent properties) of our psychological-biological-social systems. That is, in the view of interpersonal neurobiology, thoughts and feeling are not epiphenomenon of the brain or society, but rather central features of energy and information flow that govern our entire life world. Thoughts and feelings exist as evaluations of how things are going for us, i.e., we are by nature evaluative beings. On this view, emotions are not the antithesis of reason; emotions have a rationality and rationality has a feeling for us. See Dan Siegel, *The Developing Mind: How Relationships and the Brain Interact to Shape Who We Are*, 2nd ed. (New York: The Guilford Press, 2012) and Sayer, *Why Things Matter to People*.

³⁷ It should not go unmentioned that the *Education Week* Spotlight on AI tends to downplay the difficulty concerning AI's undoing of human knowing by treating the trouble as a possibility rather than an actuality. Anything "negative" about AI is given as a "potential downside" (4 instances); the "positive potential" of AI is also discussed (6 instances). The "positives" offered in the *Education Week* Spotlight are not well-established, and some are simply presented as ideas or things yet to be tested that teachers might try. But we know AI has the downsides mentioned. These are not potential; they are currently existent. The *Education Week* Spotlight on AI takes a propagandistic turn with this possibilistic rhetoric.

³⁸ Williamson, Molnar, and Boninger, *Time for a Pause*, 3-5.

³⁹ Marvin, "I'm a Math Teacher," 13.

⁴⁰ Alyson Klein, "180 Degree Turn: NYC District Goes From Banning ChatGPT to Exploring AI's Potential," *Education Week* Spotlight on AI, 6. (Originally published 5 October 2023.)

development arguments for AI are beginning to take shape and gain strength. *Education Week* reports, “Data science is rapidly becoming a critical skill, and the United States has a lot to learn from global peers in equipping students for the coming economic shift.” And the same article quotes Zarek Drozda, executive director of the University of Chicago’s Data Science 4 Everyone program, a working group monitoring developments throughout the world of methods of teaching data science, as saying: “Artificial intelligence has triggered a global talent race, and whichever country is able to find the talent to not only build AI tools, but more importantly, effectively implement the technology economy-wide, will quickly shift the economic pecking order” and “I think we [in the United States] are setting ourselves up to quickly build technologies our broader population does not understand, nor will be able to effectively leverage. To maintain U.S. competitiveness, we need to create broad, population-level data literacy.”⁴¹ The point is well-taken. Literacy in recognizing and complying with legitimate use of AI would be a strong source of support for NEPC’s call for AI governance. Governance via policy statement takes us away from understanding AI. The type of procedural knowledge involved in policy setting (tell me what steps to take and what I can and cannot do) itself invokes programmed, computer-like logic. Given the ubiquity and fluidity of AI with respect to any procedural steps governing its use, a lacquer of legislation and regulation sprayed over AI is unlikely to seal off irresponsible use of AI without an extant wish, widespread among users, to be responsible in their use of AI. Regulation, on its own, will likely set up an anxiety-ridden ontology of policing. Regulators will be hard-pressed to keep everyone in line unless regulators deputize all users to keep an eye on AI use. Educating towards AI fluency as a continuously applied, inquiry-based policing of AI data analysis and AI data analytics makes good sense because undergirding the epistemic issues we are discussing is the metaphysical issue of the relation of ideas to things.

How Do Ontological Assumptions Relate to Evident/Absent Epistemological Arguments?

Bernardo Kastrup, a philosopher and computer science theorist, has recently revived doctrines of metaphysical idealism under the new label analytical idealism. Kastrup lays out the main points of the theory using slogans such as “Everything is mind.” (The First Principle of Metaphysics.) “Matter is what presents itself on the screen of perception.” “Matter is what conscious inner life looks like from a perspective.” These slogans aim to articulate the core belief of metaphysical/analytical idealism, namely, that there is no physical world, instead all existence is a mental construct given the appearance of physical form through mental processes of perceptions.⁴² Kastrup’s view is as vulnerable as any form of idealism that reduces existence to

⁴¹ Sarah D. Sparks, “Ready or Not for an AI Economy: How U.S. Students Stack Up,” *Education Week* (18 March 2024): https://www.edweek.org/teaching-learning/ready-or-not-for-an-ai-economy-how-u-s-students-stack-up/2024/03?utm_source=nl&utm_medium=email&utm_campaign=eu&M=9329214&UUID=6e33ca36297c0e666ef3dac56b2ffdeb&T=12334777

⁴² See Bernardo Kastrup, “Conflating Abstraction with Empirical Observation: The False Mind-Matter Dichotomy,” and “Informing Metaphysical Choices with Epistemic Considerations,” 341-347 and 357-360, respectively, in *Constructivist Foundations: An Interdisciplinary Journal* 13, no. 3 (July, 2018): 357-360; “The Universe in Consciousness,” *Journal of Consciousness Studies* 25, nos. 5-6 (May-June, 2018): 125-155; “The Next Paradigm,” *Future Human Image* 9 (2018): 41-51; and “On the Plausibility of Idealism: Refuting Criticisms,” *Disputatio: International Journal of Philosophy* 9, no. 44 (May, 2017): 13-34. Also see two podcasts: Bernardo Kastrup, “The Nature of Reality and AI,” *Singularity Podcast*, Episode 57 (6 June 2022) host, Steven Parton <https://www.su.org/feedback-loop/the-nature-of-reality-ai> and *Simulation Podcast*, #676 (25 July 2020) host, Allen Saakayan <https://www.facebook.com/simulationseries/videos/simulation-676-dr-bernardo-kastrup-metaphysical-idealism/292649535303977/>

mind in an effort to avoid reliance on a material substrate to describe the make-up and structure of the world. Mind is as much a substrate, an abstracted metaphysical entity or theoretical postulate, as ever was any material world. Science seems to deal with more than ideas in understanding the world. And the very dichotomy of mind and matter misstates the philosophical issue at hand, namely, the relation of mind *to* matter. That relation is unlikely to be explained by any theory aimed at eliminating one of the relata from the discussion.⁴³

However, whether metaphysical idealism accurately and adequately describes the world in which we live, it quite correctly describes the world of Artificial Intelligence. Part of the *artificiality* of AI's "intelligence" is that AI operates in a world in which only ideas exist. Even Kastrup the metaphysical idealist says AI will never develop (self-)consciousness. There is no substrate, either mental or physical, against which AI is required to assess the validity and adequacy of its ideas. This makes the ideas very atypical in the way they are formed. AI does its thinking independently of, rather than in terms of, real-world solutions to real-world problems. This is precisely what makes AI troublesome for use in a world that seems to demand a metaphysics of both ideas and things. AI's ideas are being thought up in a world of momentary phenomenal awareness and then being applied to a world that may be constitutively different from the world in which those ideas arose. The schism between the world of pure ideas and the world of real events ineradicably imbues AI advice with the aura of incompetence.

Seeing AI's intelligence as involving consciousness rests on a significant ontological and epistemological mistake. Intelligence does not involve or entail consciousness. Intelligence is different in kind than consciousness. It is easy to make the invalid assumption that as the intelligence of a being increases so too does its consciousness. Consciousness is an achievement of our species distinct from intelligence. While consciousness is dependent upon some level of intelligence and some level of science and technology, consciousness is neither reducible to nor singularly derived from intelligence understood as the ability to recognize rules and follow patterns, to "next," to predict. AI is already much "smarter" than we are in some respects, but this has not made it 1% more conscious of itself (or anything else). In fact, culture-wide deployment of machine intelligence may work against self- and social-consciousness because for humans to know and be aware requires action on and interaction with the world, including one's inner world, that is, understanding requires a conscious participation and finding out. Absent expenditure or release of this energy, human power (which includes the essential power to decide based on our evaluation of what is good for us) cannot be developed. Yes, humans can be trained to be part of larger AI-governed processes, even including fact-checkers, or other deputized guardrail monitors along the assembly line of AI-controlled socialization. But like calculators, the understanding (or consciousness) does not inhere in or emerge from the process of calculation (intelligence) itself. AI makes no evaluation because for AI there is only behavioral output on the basis of probabilities. What AI does is meaningful to humans, but it is not meaningful to AI because AI does not produce its output on the basis of meaning, understanding, or knowledge. Its output is the result of predicting the sequence of words based on data mining of huge repositories of human-generated narrative.

⁴³ For a collection of criticisms of Kastrup's view see Chris Fields, "Mind Is an Abstraction," 350-352; Sebastian Kletzl, "Concepts, Intuitions, and Hypotheses," 352-354; Robert Prentner, "Epistemology, Metaphysics and the Preconditions of Science," 354-355; and Konrad Werner, "Is Speaking of Mind or Matter a Matter of Choice," 355-356. all in *Constructivist Foundations: An Interdisciplinary Journal* 13, no. 3 (July, 2018).

Technological “solutionism,” the position that technology is the key to solving all problems and, further, that all problems should be put in technological terms, empirically illustrates the metaphysical problem presented by AI. Certainly, technology can solve problems put in terms it can understand, but that merely restates rather than resolves the issue of AI’s incompetence. Not all problems are translatable into techno-language. For instance, bias and violations of privacy seem persistently to be involved in AI’s reductive description of students only in terms of leveraged, quantifiable data points. Technological solutionism, in requiring problems be put in terms of techno logic so that they may be solved via technology, creates an ideological position, an unbreakable logic of techno dependence that borders on religious devotion. Like any other deified voice, however otherworldly, to its devotees, AI is right because it has spoken. AI is an intelligence unto itself, a causal force that acts independently of human concerns and cognition. The performances of AI, its epistemological doings, become pronouncements of truth by their very utterance.⁴⁴ From a techno determinist point of view, AI doing becomes the pre-eminent form of knowing, that on which all other knowing is to be based, modelled, and justified. In reducing knowing to doing, AI becomes a doer without also being a knower, and knowledge becomes a performance rather than an achievement. Contrary to techno solutionism, the solution to educational problems is to integrate educational thinking into AI systems, not the other way around. Ultimately, techno solutionism is an authoritarian position wrapped up in the language of techno determinism: AI is here to stay whether you like it or not, so you better get used to it and do what it tells you. Socialization to acceptance of the pronouncements of AI becomes an epistemological norm. To become familiar with AI is to fall under its sway. AI is not so much a tool in the classroom as it is a presence to which teachers and students must adapt if they wish to be understood as efficient teachers and competent students.

But in the world in which techno determinists wish AI to do or to guide work, there are not only ideas. There are things, as well. And ideas almost always roughly and, typically only with effort, map onto things. The philosophy of the situation, in its metaphysics and in its epistemology, suggests the only way to live amicably with AI is to get everyone in the habit of criticality towards AI. We must subject pronouncements of AI, like any other claim to know, to a process of responsible human inquiry as litmus paper with which to test the reliability of the processes supporting the claims AI makes. This may be the strongest meaning of suggestions to use AI as a tool in the classroom. *Tool* talk is symbolic of technical power, something that empowers. The power is obtained through proper use of the tool. If used properly, the tool enhances and extends human power. In particular, then, AI assertions must be assessed in terms of morally responsible epistemology. AI’s reliable processes must also pass tests of responsibility. Especially, AI recommendations must be given evaluation in terms of aims beyond efficiency and techno logics that describe and express what type of lives we might want for ourselves, a set of imaginings which lie outside the parameters of AI’s sometimes unfathomable reliabilism. Epistemologically speaking, ideas go sour, sometimes disastrously, when applied in nonideal circumstances of their use. The corrective is to treat AI as just another data source and one that needs to be checked against reality. At both the level of student use of AI for paper writing and at the level of professional use of AI for policymaking, we’re asking the same thing: make critical use of the info and answers AI provides. Do not take AI advice at face

⁴⁴ For more on the performative theory of truth see Peter F. Strawson, “Truth” in Margaret MacDonald, ed., *Philosophy and Analysis* (New York: Oxford University Press, 1955): 260-277. Strawson’s deflationary theory finds its source in the redundancy theory of truth presented in Frank P. Ramsey, “Facts and Propositions,” *Proceedings of the Aristotelian Society*, Supplementary Volume 7 (1927): 153-170.

value. Sprinkle AI opinion liberally with grains of salt before consuming. The best way to make this happen is to socialize students and teachers in criticality towards curriculum, especially curriculum generated by AI. John Dewey articulates this sort of epistemic socialization in terms of the emotions students critically engaged in learning will develop. Dewey observes, “There is nothing in the inherent nature of habit that prevents intelligent method from becoming itself habitual; and there is nothing in the nature of emotion to prevent the development of intense emotional allegiance to the method.”⁴⁵ Emotional attachment to responsibilist epistemology adds a component to “understanding” AI — one richer than that found in much of the narrative about uncritical acceptance of AI. Beyond mere epistemic understanding of AI as a reliable belief-forming process, we need to add ‘understanding’ as a meta-epistemic term meaning something akin to ‘taking as a parameter of discussion’ that AI results must constantly be checked against reality as we perceive reality to be and as we wish to see reality shaped by us.

A. N. Whitehead has said, “Civilization advances by extending the number of important operations which we can perform without thinking of them.”⁴⁶ But the *type* or *character* of the civilization that develops when more and more operations are automated is not identified. We add to Whitehead’s postulate this thesis: “The self-consciousness, or lack thereof, of a civilization is dependent upon how it develops, views, adopts and understands technology.” We might further suggest that there is no necessary relationship between how technologically advanced a civilization is and how advanced that civilization’s *consciousness* is. Simply put, in the case of AI, technological advance does not automatically entail advanced cultural forms, in the sense of conscious awareness of self and others, including our most deeply held moral convictions and their evaluation. That we *can* does not imply we *must*. As Rollo May observes, “Human freedom involves our capacity to pause between the stimulus and response and, in that pause, to choose the one response toward which we wish to throw our weight. The capacity to create ourselves, based upon this freedom, is inseparable from consciousness or self-awareness.”⁴⁷ Self-awareness is exactly what’s missing from AI; and AI must rely on us to fill its epistemological emptiness.

⁴⁵ John Dewey, “Experience and Education,” in John Dewey, *The Later Works, 1925-1953*, Volume 13, ed. Jo Ann Boydston (Carbondale: Southern Illinois University Press, 1988), 58.

⁴⁶ Alfred North Whitehead, *An Introduction to Mathematics* (London: Williams & Norgate, 1912), 46.

⁴⁷ Rollo May, *The Courage to Create* (New York: W. W. Norton & Company, reprint edition, 1994), 100.