



© Cristian-AdobeStock

LIGHTS - Leadership

The Fog of Intelligence: AI education in an era of shifting definitions

ESCP Impact Paper No.2025-05-EN

Kenneth KRUSHEL
ESCP Business School

[The Fog of Intelligence: AI education in an era of shifting definitions

Kenneth Krushel*

ESCP Business School

Abstract

How do business school programs meaningfully teach artificial intelligence when the subject itself defies stable definition? Amid rapid technological transformation, we want to measure AI's growth, but we are challenged by what to measure. This is a problem not of data, but of definitions. The goal is not definitional certainty but definitional literacy, a cartography of competing understandings equipping students to navigate a landscape where meaning itself is contested and constantly evolving. The goal is to prepare future business leaders for a world in which the meaning of (artificial) intelligence is constructed on shifting ground.

"Fog of intelligence" seems contradictory - an oxymoron in the literal sense: pointedly foolish. Yet sometimes the tension between opposing words reveals what simpler language cannot. What appears nonsensical - seeking clarity by accepting obscurity - may be precisely the wisdom needed for teaching a subject that defies definition. The Fog of Intelligence points to a meta-question about teaching the unknowable, about preparing for undefined futures, about guiding MBA students through emergence.

Keywords: artificial intelligence, definitional fluidity, business education, interdisciplinary teaching, technological uncertainty

*Professor, NYU Stern School of Business

ESCP Impact Papers are in draft form. This paper is circulated for the purposes of comment and discussion only. Hence, it does not preclude simultaneous or subsequent publication elsewhere. ESCP Impact Papers are not refereed. The form and content of papers are the responsibility of individual authors. ESCP Business School does not bear any responsibility for the views expressed in the articles. Copyright for the paper is held by the individual authors.


The Fog of Intelligence: AI education in an era of shifting definitions

The objective in teaching artificial intelligence, at least as explained in course descriptions populating MBA program curricula, is to not only introduce students to the subject but prepare them for a post-industrial economy, what economists, sociologists, AI pundits, and politicians/regulators cite as a migration from a fourth to a fifth industrial revolution. Yet teaching these technological shifts reveals a fundamental challenge: the very terms we use to describe them resist stable definition.

Simply explained, the **Fourth Industrial Revolution** involves the rise of "intelligent machines". For example, AI, machine learning, automation, LLMs, and robotics: societal disruption is possible as machines replace human beings and algorithms make decisions. Society grapples with surveillance, data privacy, disinformation, and growing economic inequality. Technology grows "smarter", but we grow more anxious.

The Fifth Industrial Revolution, in an optimistic scenario, suggest that humans and advanced digital technology collaborate; but the divide between those who design and those who live with the consequences deepens. Artificial General Intelligence (AGI) becomes increasingly prominent, manifesting, for example, as agentic AI and digital avatars - where AI can act *autonomously* with minimal if any human supervision. This progression from AGI points toward the ultimate goal: artificial superintelligence (ASI). The Fifth Industrial Revolution is speculative and aspirational, helping us - in the words of Open AI's CEO Sam Altman, "elevate humanity by increasing abundance, turbo charging the global economy, and aiding in the discovery of new scientific knowledge that changes the limits of possibility."

These category headings are more than academic constructs. The terms are captivating investors, unlocking buckets of capital previously unavailable. Evidence of how companies are exploiting the concept of AI, even if ill-defined and lacking profitable use cases, is demonstrated by how frequently AGI is mentioned, according to the *Financial Times* report on AGI (Jun 19, 2025): 53% more times in companies' earnings calls in the first quarter of 2025 compared to Q1, 2025. The motivation for investors is not just the perceived opportunity to earn hundreds of millions, but altering what drives GDP, which in investors' projections, will potentially create trillions of dollars in value.



Source: Alpha Sense – Q2 data to June 16, 2025

Both META and OpenAI are making significant investments in advanced AI research, with a stated purpose of achieving ASI. META has invested \$14.3 billion in data labeling and data infrastructure company *Scale* (June 2025), and reportedly offered hundreds of millions of dollars to attract “AI all-stars” and superintelligence researchers to establish a “superintelligence lab”, aiming to develop AI systems that surpass human intelligence.

OpenAI has raised \$40 billion in a recent funding round (at a \$300 billion valuation) explicitly targeting superintelligence. OpenAI CEO Sam Altman has declared that the era of superintelligence has begun and has outlined ambitious timelines for AI agents performing cognitive work and generating original scientific discoveries. Imagine: superintelligence could tell scientists what to discover, proposing hypotheses, designing experiments, and analyzing results. In essence, ASI would surpass human intelligence across all cognitive domains, ranging from problem solving to creativity.

Ilya Sutskever, previously co-founder of OpenAI and one of the world’s leading AI researchers, founded “Safe Superintelligence” in 2024. Based on an April 2025 funding round, the startup is worth \$32 billion.

The challenge when introducing students to the lexicon of AI is how to teach a subject where the actual meanings of the acronyms - AI, AGI, ASI, etc. - are both elusive and illusive. It's not that teachers don't understand; it's that our understanding may be outdated by the time we develop our lectures and PowerPoints. We prepare today's lesson for tomorrow's technology using yesterday's definitions.

The category heading shifts not with some natural evolution of understanding but with strategic necessity, market positioning, investor relationships, and capital raising requirements. If the unraveling of the term “AI” is any indication, it seems unlikely the industry will coalesce around one definition, and will continue to debate the meaning of various aspects of innovation such as agentic AI, physical intelligence, and AI safety. This leads to misaligned expectations and difficulties in measuring the value of teaching the subject to students who have already adapted to and embraced digital technology as a utility, something they take for granted. Adding to the challenge: based on recent research, many aspects of AI use, especially writing code and using chatbots as surrogate essayists, stunt critical thinking.

The elusiveness of reliable definition

To teach is to define. But define what, why and how? As John Dewey wrote, “If we teach today's learners as we taught yesterday's, we rob them of tomorrow. When it comes to artificial intelligence, definition itself becomes a moving target, evolving with each technological breakthrough, paradigm shift, Silicon Valley manifesto, and political realignment. How do we teach what we cannot definitively define?

It's like trying to photograph fog. Stand in one place, and you see a particular configuration of mist. Move slightly, and the entire landscape shifts. The photographer doesn't capture the fog itself but merely one momentary arrangement of its particles - just as our definitions of AI capture not its essence but a temporary configuration of understanding.

Geoffrey Hinton, often called the “godfather of deep learning,” offers perhaps the most technically grounded definition. In his Royal Society lecture, Hinton described AI as “computational systems that can learn from data to perform tasks that typically require human intelligence.” Yet he later admitted in an interview with *The Guardian* that “the definition keeps changing because as soon as we figure out how to make computers do

something that previously required intelligence, it stops being an AI problem." Hinton also warns that without AI safety regulation, AGI could lead to human extinction. (Note: Ilya Sutskever completed his PhD in computer science under Hinton's supervision at the University of Toronto).

This definitional paradox - what philosopher Hubert Dreyfus called the "AI effect" in "Alchemy and AI" (1965) - points to a fundamental challenge in teaching AI. The goalposts continuously move. What was once considered the pinnacle of artificial intelligence becomes, upon achievement, merely sophisticated software.

Kate Crawford, author of "Atlas of AI," takes a markedly different approach, defining AI not as a technical achievement but as "a registry of power." For Crawford, AI is "neither artificial nor intelligent." Rather, artificial intelligence is both embodied and material, made from natural resources, fuel, human labor, infrastructures, logistics, histories, and classifications.

These divergent perspectives illustrate our first dilemma. Teaching AI becomes akin to the ancient parable of the blind men and the elephant - each authority grasping one part of a larger reality they cannot fully perceive.

The gentle singularity: a new digital species

Sam Altman has been particularly influential - and particularly fluid - in defining the field's horizons. In Congressional testimony, Altman defined AGI as "a system that is generally smarter than humans." By 2024, he had shifted to discussing "superintelligence," which he described as "AI systems that are quantitatively and qualitatively far more capable than those we have today."

Recently, Altman introduced the term "the gentle singularity." The singularity refers to a hypothetical future point at which technological growth becomes uncontrollable and irreversible. Altman uses the term to explain that "OpenAI is a lot of things now, but before anything else, we are a superintelligence research company."

Altman is not just preparing for the singularity; he argues that the early stages of the process have already begun, with OpenAI at the forefront of the research and development.

This definitional drift is not merely semantic. It reflects strategic positioning in an industry where definitions shape public perception, regulatory frameworks, and corporate valuations.

Contrast Altman's expansive vision with that of Mustafa Suleyman, co-founder of DeepMind, former CEO of Inflection AI, and now CEO of Microsoft AI. In his book "The Coming Wave," Suleyman offers a more grounded definition: "AI is simply the science of making machines capable of tackling tasks that would require intelligence if performed by humans." His definition - focused on task performance rather than general capability - views AI as a "digital specie", as if we are creating a new civilization. He is both optimistic and cautious, highlighting the importance of "symbiotic relationships" between humans and AI, while stressing the need for regulatory structures to manage AI development.

Teaching through the fog

The definitional chaos creates a "fog of intelligence," a condition where clear vision becomes impossible, yet navigation remains necessary. The solution lies not in dispersing this fog but in teaching students to navigate within it.

Consider the field of philosophy, which has thrived for millennia despite, or perhaps because of, its resistance to definitive answers. Philosophy professors don't begin by offering final definitions of "truth," "justice," or "beauty." Instead, they introduce competing definitions, exploring how each illuminates different aspects of these concepts.

Similarly, effective AI education might begin not with a single definition but with a map of the definitional landscape, exploring how different conceptions of AI reflect different values, priorities, and disciplinary perspectives. We might approach teaching AI not as delivering certainty but as cultivating comfort with uncertainty - what poet John Keats called "negative capability," the capacity to remain in "uncertainties, mysteries, doubts, without any irritable reaching after fact and reason." This literary concept offers a crucial insight for AI education: the ability to dwell productively within questions rather than rushing toward premature answers. When we teach students to inhabit uncertainty rather than flee from it, we prepare them for a field where the most important insights often emerge from sustained questioning rather than quick resolution.

The cartography of understanding

In the fall of 2024, I assigned MBA students a simple task: have a conversation with ChatGPT about surveillance capitalism, then write about what happened. I expected analysis. What I received was revelation.

One student, a marketing manager, found herself in a forty-minute exchange about data privacy with a system that demonstrated intimate knowledge of her company's customer acquisition strategies. "It knew things about my industry I didn't know it knew," she wrote. "But it also couldn't understand why I felt creepy about our targeting methods. It was like talking to a brilliant colleague who had no soul."

Another student, working in financial services, tried to get the system to explain why algorithmic trading might be ethically problematic. The AI generated sophisticated arguments about market fairness, then casually offered to help him design better trading algorithms. "I realized," he wrote, "that I was debating ethics with something that had no stake in the outcome."

These weren't academic exercises. They were encounters with the fundamental instability of AI as a category. Students discovered that the technology's capabilities exceed our definitional frameworks, and that this gap matters for business strategy. More importantly, they experienced firsthand how AI's definitional ambiguity creates practical challenges for decision-making in real organizational contexts.

Definitions shape expectations

Effective AI education starts with phenomena students can observe. In our Epic Games case study, students analyze how the company's definition of the metaverse as a "digital ecosystem" enabled specific business strategies while constraining others. They see definition as strategy in action.

When we study ChatGPT's launch, students don't start with technical specifications. They explore the gap between prediction and reality. Headlines promised superintelligent productivity tools. Almost three years later, those productivity gains remain underwhelming, while something unanticipated occurred: people formed relationships with AI systems, talking to chatbots as friends, mentors, therapists.

Students realize they're not just learning about technology; they're witnessing how definitions shape expectations, investments, and ultimately, reality itself.

I've started asking students to design prompts that reveal AI's limitations. One team tried to get ChatGPT to experience buyer's remorse. Another asked it to procrastinate. The results illuminate something essential about AI's definitional problem: the system can simulate these experiences perfectly while claiming it cannot experience them at all.

"It's like method acting," one student observed. "But who's the actor? And if there's no actor, what does the performance mean?" This question cuts to the heart of AI's definitional instability. We're watching sophisticated simulation without clear understanding of what, if anything, lies beneath the performance. The student's insight reveals how AI challenges not just our technical understanding but our fundamental categories for distinguishing authentic experience from convincing imitation.

The interdisciplinary imperative

Breakthroughs in AI understanding often occur at disciplinary boundaries. Our course structure reflects this reality. Students encounter AI through multiple disciplines: economic (surveillance capitalism), cultural (synthetic media), philosophical (machine consciousness), and strategic (platform competition).

When we read chapters from Walter Benjamin's "The Work of Art in the Age of Mechanical Reproduction" (1935) alongside discussions of current deepfake examples, students discover how 1930s media theory illuminates current AI challenges. Benjamin's insight about mechanical reproduction - that new technologies don't just change what we make but how we understand authenticity itself - helps students recognize that AI doesn't merely generate content; it challenges our assumptions about originality, creativity, and human uniqueness. The philosophical framework reveals what's genuinely new about AI-generated content versus what extends older technological disruptions.

This interdisciplinary approach serves a practical purpose. MBA graduates will evaluate, implement, and manage AI systems - not build them. They need to assess vendor claims, navigate regulatory proposals, and make strategic decisions where the technology's meaning remains contested.

From cognitive science, we gain insights into the nature of intelligence itself. From philosophy, we explore questions of mind, consciousness, and meaning. From sociology and anthropology, we examine how AI technologies reflect and reshape social structures. From economics, we analyze AI's impact on markets, labor, and value creation. From ethics, we confront questions of responsibility, fairness, and harm. The intangibility of AI encourages people to make meaning according to individual perspective: for some it's a scientific goal, a religion, or even a marketing term.

The productive nature of uncertainty

Consider how business schools teach strategy. We don't begin with a definitive definition of "competitive advantage" because the concept means different things in different contexts. Instead, we explore multiple frameworks, helping students recognize which applies when. AI education works similarly. When students encounter conflicting definitions in business contexts, they learn to ask: Which definition serves whose interests? What assumptions does this definition embed? What possibilities does it foreclose?

These questions prove more valuable than any single definition. They prepare students for a world where, as one recent graduate told me, "Every vendor claims their software uses AI, every startup promises AI transformation, and every consultant peddles AI strategy, but nobody agrees on what AI actually is."

Building upon the student's insight, definitions shape business models, regulatory responses, and user experiences. Students begin to see that definitional choices have consequences, that how we define AI influences what we build, how we regulate it, and ultimately, how it shapes society.

One class exercise asks students to rewrite their company's mission statement as if it were an AI company. A pharmaceutical student reimagined drug discovery as "intelligence augmentation for molecular research." A retail student described inventory management as "predictive commerce optimization." The results reveal how AI terminology reshapes business identity while often obscuring rather than clarifying actual capabilities and limitations.

This approach aligns with what philosopher of technology Albert Borgmann calls "focal practices" - activities that demand engagement, attention, and discernment rather than mere consumption. Teaching AI as a focal practice means inviting students to actively engage with its definitional complexities rather than passively consuming simplified explanations. Without such engagement, students increasingly default to academic shortcuts, using AI tools to complete assignments without genuine learning - precisely the opposite of the critical thinking skills they need to navigate AI's professional implications.

This is not the type of mastery we intend. True mastery lies not in memorizing definitions but in understanding why definitions matter - and how they shape not just our understanding of AI but the very development of the technology itself.

Conclusion: embracing the Fog

The Fog of Intelligence becomes not an obstacle to understanding but a medium through which new understandings emerge - if not treated as a barrier to clarity but as a field of possibility. In a domain where definitions evolve rapidly, teaching students to navigate definitional multiplicity may be the most practical preparation we can provide.

What seems like a limitation becomes, paradoxically, a source of freedom. We equip students to navigate a landscape where the terrain itself constantly shifts, where the meaning of intelligence - artificial or otherwise - remains productively unsettled. At the risk of sounding pedantic or performative, we borrow upon Walter Benjamin's interpretation of technology, recognizing that AI challenges society to reconstruct definitions of authentic experience. AI changes not just how we make things but what we think those things fundamentally are.

Within five years, much of what we currently teach about AI will be automated. Systems will generate business cases, market analyses, and strategic frameworks endlessly. But that's not the point of education anyway.

The real work, the work no machine can inherit, is the craft of judgment: knowing when a definition matters, when it conceals, and when to ask the unsettling question. It is the art of navigating shifting meanings without losing one's footing.

Students need to understand that in AI, as in business, the map is never the territory. Definitions are instruments, not truths. To move skillfully among definitions, to sense their limits and shift frameworks as the landscape requires - may be the most enduring skill we can offer.

References

Time Magazine, January 8, 2025 : "How OpenAI's Sam Altman is Thinking About AGI and Superintelligence in."

Dewey, John. (1916), "Democracy and Education: An Introduction to the Philosophy of Education", MacMillian.

Vision Sciences Society. "2022 Keynote." *Vision Sciences Society*, www.visionsciences.org/2022-keynote.

"Geoffrey Hinton: 'godfather of AI' quits Google and warns over dangers of machine learning." The Guardian, 2 May 2023, www.theguardian.com/technology/2023/may/02/geoffrey-hinton-godfather-of-ai-quits-google-warns-dangers-of-machine-learning

Dreyfus, Hubert, (1965) "Alchemy and AI". Rand Corporation.

Crawford, Kate. (2021) "Atlas of AI". Yale University Press.

U.S. Senate Judiciary Committee. "Oversight of AI: Rules for Artificial Intelligence." 118th Congress, 1st Session, 2023.

Altman, Sam. "The Intelligence Age." *Medium*, 23 Sept. 2024, medium.com/@padamasalilavanya/the-intelligence-age-by-sam-altman-unveiling-a-new-era-of-human-potential-3cee77a490ee

Altman, Sam. June 10, 2025, <https://blog.samaltman.com/the-gentle-singularity>

Suleyman, Mustafa. (2023) "The Coming Wave," Crown.

Herman, John. "Everyone is Cheating Their Way Through College." *New York Magazine*, 7 May 2025.