

The ghost and the machine: disentangling the fundamental distinctions between human and artificial intelligence

Alberto Oliverio

Sapienza University of Rome, Italy

e-mail: alberto.oliverio@uniroma1.it

The comparison between human intelligence and artificial intelligence (AI) has increasingly captivated scientific and philosophical discourse, raising fundamental questions about the nature of cognition, embodiment, and consciousness. At the heart of this inquiry lies a deceptively simple question: Is AI a synthetic analogue of the human brain -a silicon-based counterpart- or does it constitute a fundamentally novel phenomenon (Chalmers 1995; Boden 2006)? Is it a sophisticated mimic, an ultra-efficient calculator, or a wholly distinct ontological entity?

To explore this question, one must disaggregate the substrates, learning processes, and emergent capacities of both biological and artificial systems. This commentary does so through three analytic dimensions: (1) substrate and construction, (2) learning mechanisms, and (3) emergent properties.

Substrate and construction: wetware versus hardware

One of the most profound differences between human and artificial intelligence lies in the material substrates upon which each system is built.

Human cognition is instantiated in *wetware*—the biological brain. Comprising approximately 86 billion neurons (Azevedo et al. 2009), the human brain is a product of evolutionary selection and adaptive pressures. Neural communication occurs via synapses through electrochemical signaling involving neurotransmitters such as GABA, acetylcholine, dopamine and serotonin. The system is massively parallel, highly redundant, and remarkably energy efficient. Moreover, it is deeply embedded in the body and the environment,

rendering cognition a distributed and embodied process (Clark 2008; Varela et al. 1991).

By contrast, artificial intelligence operates on *hardware* -circuits composed of silicon, copper, and logic gates. Modern AI models, especially deep learning architectures, rely heavily on Graphic Processing Units -GPUs optimized for matrix operations (LeCun et al. 2015). These systems execute binary logic with high precision and computational speed but consume vast amounts of energy, especially during training (Strubell et al. 2019). While inspired by neural processes, artificial neural networks are simplifications: discrete, unidirectional, and devoid of biochemical modulation.

Although artificial neural networks borrow their terminology and high-level architecture from neuroscience, the resemblance is superficial. As Zador (2019) argues, many current AI systems disregard crucial principles of biological learning and neural dynamics. The analogy of a paper airplane to a bird captures the difference appropriately: one is a toy that mimics flight; the other is the evolved product of millions of years of biological adaptation.

Learning mechanisms: sparse embodiment vs. statistical saturation

Another critical distinction concerns how these systems acquire, process, and generalize information (Oliverio 2021).

Humans exhibit a form of *sparse*, embodied learning. A child can recognize a new object after one or two exposures -a phenomenon known as one-shot or few-shot learning (Lake et al. 2015).

Human learning is multimodal and causally structured, grounded in sensorimotor experiences, social interaction, and contextual inference (Gopnik et al. 2004). Additionally, human learners construct internal models of the world, enabling them to infer causality, anticipate consequences, and develop “common sense” (Spelke and Kinzler 2007).

In contrast, AI models -particularly large language models (LLMs)- rely on voracious data consumption. They are trained on massive amounts containing billions of words and learn by identifying statistical regularities (Brown et al. 2020). Their learning is fundamentally correlation-driven rather than causation-aware. For instance, an LLM knows that “cat” often appears near words like “meow” or “fur,” but lacks an experiential or embodied understanding of what a cat *is* (Bender and Koller 2020).

Reinforcement learning (Sutton and Barto 2018) provides a superficial parallel to human learning via reward-based feedback, but it lacks the affective, social, and contextual depth of human experience. Unlike humans -who learn to avoid hot objects through pain and embodied awareness- AI models learn through scalar feedback signals in abstract environments.

Thus, while both systems exhibit forms of generalization, the processes differ totally in architecture, modality, and efficiency. Human learners acquire deep, embodied understanding with sparse data; AI models require immense datasets and still struggle with generalization beyond training distributions (Marcus 2018).

Emergent properties: divergent manifestations of complexity

Emergence -the arising of system-level properties from component-level interactions- offers one of the most intriguing parallels and discontinuities between human and artificial intelligence.

In humans, consciousness is the paradigmatic emergent property. Despite extensive neuroscientific research, no singular “consciousness neuron” has been identified. Rather, subjective experience

arises from the dynamic, distributed activity of the brain (Dehaene and Changeux 2011). Other emergent phenomena -creativity, empathy, morality- are not directly encoded in our genes but emerge from sociocultural interactions and evolved predispositions (Tomasello 2014).

As AI models get bigger, they start showing unexpected abilities. For example, GPT (Generative Pre-trained Transformer) models have learned to reason, translate, and understand abstract ideas, even though they weren’t specifically programmed to do these things (Wei et al. 2022). These skills seem to come naturally as a result of training the models to predict the next word in a sentence. However, the nature of AI emergence is qualitatively different. There is no evidence that AI systems possess *qualia*, intentionality, or self-awareness (Searle 1980; Nagel 1974). Their apparent “reasoning” is better understood as highly refined pattern-matching rather than genuine understanding (Mitchell 2023). In this sense, AI’s emergent capabilities are simulations of intelligence rather than instantiations of it.

To borrow the initial metaphor: in artificial intelligence, the “ghost in the machine” is spectral, a brilliant echo of human cognition, but devoid of interiority. In biological intelligence, the ghost *is* the machine: consciousness is not just an emergent property but a constitutive one.

Conclusion: A mirror of mind, not a mind itself

The relationship between human and artificial intelligence is characterized by profound continuity in conceptual inspiration and emergent complexity, yet also profound discontinuity in embodiment, understanding, and subjective experience.

AI systems mirror the form and function of human cognition but remain ontologically distinct. They are constructed, not evolved; learned from text, not from life; and mimic rather than possess experience. While they may generate artifacts of remarkable sophistication, they do so

without the embodied, emotional, or conscious substrate that defines human life.

The most compelling question may not be when -or whether- AI will attain human-like intelligence, but rather what alien forms of cognition may emerge from non-biological substrates, and what their existence might reveal about the evolutionary, embodied nature of our own.

The interwoven mind: how AI is transforming our human identity

Humans have always been shaped by the instruments they craft. From the lever, which magnified our physical strength, to the book, which expanded our memory, tools have extended our natural capabilities (Clark 2003). But the technologies we develop today are of a different nature. They no longer merely extend our will -they engage with it. These intelligent, responsive systems behave less like tools and more like companions. Our bond with these intelligent machines has shifted from utility to intimacy, fundamentally and subtly altering who we are. This is more than a technological revolution; it's an existential one. It prompts us to ask: what does it mean to be human now?

One of the most immediate transformations is occurring within the inner workings of our own minds, particularly in how we retain and manage information. When did you last memorize a phone number or directions? Increasingly, we are delegating such mental tasks to external systems -search engines, digital assistants- forming what some call an "external brain" (Sparrow et al. 2011). This practice, known as *cognitive offloading*, is strikingly effective. It liberates mental bandwidth for broader, more abstract thinking (Risko and Gilbert 2016). We are no longer expected to hold every detail but to know how and where to locate information when it matters most.

Yet, this convenience carries a cost. Committing knowledge to memory creates neural networks that support creativity and insight (Carr 2010; Oliverio 2008). When we over-rely

on external sources, do we risk diminishing this internal architecture? The very devices that help us remember also relentlessly compete for our attention. Personalized feeds, infinite scrolling, and constant alerts -each algorithmically crafted to engage us- affect our ability to concentrate (Williams 2018). We have grown skilled at skimming and multitasking, but the deeper, undistracted focus needed for reflection and original thought is increasingly rare (Newport 2016).

Our reliance on AI now reaches into the realm of decision-making. From the directions we follow to the films we watch, algorithmic guidance shapes our choices (Mittelstadt et al. 2016). These systems ease decision fatigue and help us navigate, but excessive dependence on them may erode our judgment. Trusting opaque, "black box" algorithms, can mean surrendering agency without realizing it (Burrell 2016). We risk delegating not just choices, but our capacity for critical thinking.

However, this evolving partnership is not solely about what we might lose, it also offers remarkable gains, particularly in the creative domain. For the first time, we have intelligent systems that can collaborate with us artistically. Generative AI can offer prompts to writers, sketch ideas for visual artists, or suggest melodies for composers (McCormack et al. 2019). It is an unlimited creative partner. This changes the human role from being the sole creator to serving as initiator, curator, and editor. We become visionary directors working with powerful creative engines.

As we watch machines excel in domains, we once thought uniquely human -reasoning, composing, designing- we are forced to re-examine ourselves. If machines can imitate these capacities, what distinguishes us? The answer lies in experience. Machines may replicate output, but they cannot feel. They cannot sense the warmth of sunlight, mourn a loss, or appreciate joy. Our lived, embodied experience -subjective and emotional- is our most vital distinction (Damasio 1999; Nagel 1974).

This recognition is reshaping the traits we most value. As AI takes over analytical and logical tasks, the uniquely human qualities emerge:

empathy, emotional intelligence, moral clarity, and authentic connection (Turkle 2011). The ability to understand and care for others, to make ethical decisions, to lead with compassion - these are the strengths that machines cannot mimic and that our future may increasingly depend upon.

Ultimately, the boundary between human and machine is dissolving. Our devices are not just tools; they are cognitive extensions of ourselves, a kind of “extended Mind” (Clark and Chalmers 1998). In a real sense, we are becoming hybrid beings—cyborgs—whose thoughts and memories are partially digital.

So, are intelligent machines changing us? Undoubtedly. They are reshaping our inner lives and compelling us to rethink our role in the world. But this is not a tale of diminishing humanity. It is a narrative of transformation. We are moving beyond intellect as our defining feature and rediscovering what lies at the heart of being human: consciousness, empathy, and the capacity to connect. In the end, being human is less about capability, and more about identity—about who we choose to become.

References

- Azevedo FAC, Carvalho LRB, Grinberg LT, et al (2009) Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *J Comp Neurol* 513: 532–541. <https://doi.org/10.1002/cne.21974>
- Bender EM, Koller A (2020) Climbing towards NLU: On meaning, form, and understanding in the age of data. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, p. 5185–5198. <https://aclanthology.org/2020.acl-main.463/>
- Boden MA (2006) *Mind as Machine: A History of Cognitive Science*, Oxford University Press, Oxford.
- Brown T, Mann B, Ryder N, et al (2020) Language models are few-shot learners. *Adv Neural Inf Process Syst* 33:1877–1901. <https://doi.org/0.3389/fcomp.2023.1118182>
- Burrell J (2016) How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data Soc* 3. <https://doi.org/10.1177/2053951715622512>
- Carr N (2010) *The shallows: What the Internet is doing to our brains*, W. W. Norton & Company, New York.
- Chalmers DJ (1995) Facing up to the problem of consciousness. *J Conscious Stud* 2:200–219.
- Clark A (2003) *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford University Press, Oxford.
- Clark A (2008) *Supersizing the mind: embodiment, action, and cognitive extension*. Oxford University Press, Oxford.
- Clark A, Chalmers D (1998) The extended mind. *Analysis* 58:7–19. <https://www.jstor.org/stable/3328150>
- Damasio A (1999) *The feeling of what happens: Body and emotion in the making of consciousness*, A Harvest Book, Harcourt.
- Dehaene S, Changeux JP (2011) Experimental and theoretical approaches to conscious processing. *Neuron* 70:200–227. <https://doi.org/10.1016/j.neuron.2011.03.018>
- Gopnik A, Meltzoff AF, Kuhl PK (2004) *The scientist in the crib: Minds, brains, and how children learn*, Harper Perennial.
- Lake BM, Salakhutdinov R, Tenenbaum JB (2015) Human-level concept learning through probabilistic program induction. *Science* 350:1332–1338. <https://doi.org/10.1126/science.aab3050>
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521: 436–444. <https://doi.org/10.1038/nature14539>
- Marcus G (2018) Deep learning: A critical appraisal. *arXiv:1801.00631*. <https://doi.org/10.48550/arXiv.1801.00631>
- McCormack J, Gifford T, Hutchings P (2019) Autonomy, authenticity, authorship and intention in computer generated art. In: *Computational Creativity: the philosophy and engineering of autonomously creative systems*, *Proceedings 8th International Conference, EvoMUSART 2019*, Leipzig, Germany, Springer, p. 57–76.

- Mitchell M (2023) Artificial intelligence: a guide for thinking humans, Penguin, UK.
- Mittelstadt B. D, Allo P, Taddeo M, et al (2016) The ethics of algorithms: Mapping the debate. *Big Data Soc* 3. <https://doi.org/10.1177/2053951716679679>
- Nagel T (1974) What is it like to be a bat? *Philos Rev* 83:435–450. <https://doi.org/10.1093/oso/9780197752791.001.0001>
- Newport C (2016) Deep work: Rules for focused success in a distracted world, Grand Central Publishing.
- Oliverio A (2021) Rôle essentiel de la mémoire dans la formation de toute représentation. In: Belaubre G, Chenin E, Mastrangelo V, et al (eds) *Les signatures neurobiologiques de la conscience*. Académie Européenne Interdisciplinaire des Sciences, Edp Sciences, Paris, p. 31–41. <https://www.edp-open.org/books/edp-open-books/421-les-signatures-neurobiologiques-de-la-conscience>
- Oliverio A (2008) Brain and creativity. *Progress of Theoretical Physics* 173:66–78. <https://doi.org/10.1143/PTPS.173.66>
- Risko E, Gilbert SJ (2016) Cognitive offloading. *Trends Cogn Sci* 20:676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Searle JR (1980) Minds, brains, and programs. *Behav Brain Sci* 3:417–457. <https://doi.org/10.1017/S0140525X00005756>
- Sparrow B, Liu J, Wegner DM (2011) Google effects on memory: Cognitive consequences of having information at our fingertips. *Science* 333:76–778. <https://doi.org/10.1126/science.1207745>
- Spelke ES, Kinzler KD (2007) Core knowledge. *Dev Sci* 10:89–96. <https://doi.org/10.1111/j.1467-7687.2007.00569.x>
- Strubell E, Ganesh A, McCallum A (2019) Energy and policy considerations for deep learning in NLP, ACL.
- Sutton RS, Barto AG (2018) Reinforcement learning: An introduction, adaptive computation and machine learning series, MIT Press.
- Tomasello M (2014) A natural history of human thinking, Harvard University Press.
- Turkle S (2011) Alone together: Why we expect more from technology and less from each other. Basic Books.
- Varela FJ, Thompson E, Rosch E (1991) The embodied mind: cognitive science and human experience, MIT Press.
- Wei J, Tay Y, Bommasani R, et al (2022) Emergent abilities of large language models. [arXiv:2206.07682](https://arxiv.org/abs/2206.07682). <https://doi.org/10.48550/arXiv.2206.07682>
- Williams J (2018) Stand out of our light: Freedom and resistance in the attention economy, Cambridge University Press.
- Zador AM (2019) A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Commun* 10:3770. <https://doi.org/10.1038/s41467-019-11786-6>



This work is distributed under the terms of a Creative Commons Attribution-NonCommercial 4.0 Unported License <http://creativecommons.org/licenses/by-nc/4.0/>

