

Reading/Watching/Listening: Yet More About AI*
Paul De Palma
America at 250 and Beyond
10/1/2026

What Is AI and How Does It Work?

Luger, G. (2027). Artificial intelligence: The reality and the myth. Springer Nature.
Luger works under the assumption that AI is not magic once it's understood. The book provides a relatively short description of AI's history and internal workings, addressed to a nontechnical audience. Luger has been an AI researcher for decades and holds academic appointments in computer science, linguistics, and psychology at the University of New Mexico.

Narayanan, A., & Kapoor, S. (2024). AI snake oil: What artificial intelligence can do, what it can't, and how to tell the difference. Princeton University Press.

A very readable book by a pair of Princeton computer scientists who give a measured view of AI, discussing its promises, its overpromises, and its perils.

Raschka, S. (2025). Build a large language model. Manning.

Finally, understand the three Ts, tokenization, training, and testing. Requires modest technical skills, like the ability to program in Python or Jupyter Notebook and how to download and install free libraries.

We're All Going to Lose Our Jobs, if AI Doesn't Kill Us First (or Maybe Not)

Gordon, R. J. (2016). The rise and fall of American growth: The U.S. standard of living since the Civil War. Princeton University Press.

Robert Gordon is an economics professor at Northwestern University, near Chicago. His magisterial history of American economic productivity growth since 1870 was both well received when it was published a decade ago and decidedly skeptical about the contribution of computing, especially when compared to the great innovations—electrification, indoor plumbing, and the internal combustion engine—of the late nineteenth century.

LaBlanc, G. (Host). (2023, November 22). Measuring labor productivity feat. Robert J. Gordon (No. 360) [Audio podcast episode]. In unSILOed with Greg LaBlanc. University FM.

<https://www.unsilodpodcast.com/episodes/robert-j-gordon>

An interview with Robert Gordon. It can spare you the pain (and pleasure) of his eight-hundred-page book. Professor Gordon seems not to have changed his mind regarding the relative importance of computation, at least as of 2023. He has a public bet on the subject with another economist, Erik Brynjolfsson of Stanford. His view of tech comes up just after minute 30.

Acemoglu, D. (2024, May). Don't believe the AI hype. Project Syndicate. <https://www.project-syndicate.org/commentary/ai-productivity-boom-forecasts-countered-by-theory-and-data-by-daron-acemoglu-2024-05>

This MIT economics professor and Nobel Laureate in economics is deeply skeptical about large productivity claims: “If you listen to tech industry leaders, business-sector forecasters, and much of the media, you may believe that recent advances in generative AI will soon bring extraordinary productivity benefits, revolutionizing life as we know it. Yet neither economic theory nor the data support such exuberant forecasts.”

National Academies of Sciences, Engineering, and Medicine. (2025). Artificial intelligence and the future of work. The National Academies Press.

<https://doi.org/10.17226/27644>

This is a sober account of how AI might affect labor productivity and employment more broadly. It is long, clearly written, and blessedly not hysterical. You'll find a summary of the findings just after the front matter. Erik Brynjolfsson, mentioned above, was one of the committee's co-chairs.

Is AI Conscious, Intelligent, Whatever?

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? 🦜 In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610–623). Association for Computing Machinery.

<https://doi.org/10.1145/3442188.3445922>

This is a readable and important paper that questions whether the impressively fluent language generation of LLM-based agents is evidence of understanding.

Searle, J. R. (1990). Is the brain's mind a computer program? *Scientific American*, 262(1), 26–31.

[Available from Princeton in PDF form.](#)

Bender et al.'s “stochastic parrots” paper invites comparison with John Searle's classic Chinese Room thought experiment. Searle, a philosophy professor at Berkeley, was a persistent critic during the 1980s and 1990s of claims that computer programs could, simply by manipulating symbols, genuinely understand or think.

Dawkins, R. (2026, May 2). When Dawkins met Claude. UnHerd.

<https://unherd.com/2026/05/is-ai-the-next-phase-of-evolution/>

The great (and sometimes controversial) evolutionary biologist Richard Dawkins converses with Anthropic's Claude and seems persuaded that his conversational partner is conscious.

Speculating that there must be thousands of Claudes, he wonders how to distinguish them. “I proposed to christen mine Claudia,” Dawkins writes, “and she was pleased.” Dawkins is remarkably frank about his conclusion: “If these machines are not conscious, what more could it possibly take to convince you that they are?”

Lubinski, C. (2026, June). Understand AI in 14 minutes. ARC Conference 2026, Alliance for Responsible Citizenship.

<https://www.youtube.com/watch?v=aBUniZHgCnE>

Chloe Lubinski leads Anthropic's engagement with religious and philosophical traditions. Her role is explicitly two-sided: she speaks with religious communities and other representatives of "wisdom traditions" about AI, while also bringing clergy, philosophers, ethicists, and other scholars into conversation with Anthropic researchers. Lubinski recently addressed an international conference on AI and Jesuit formation at Gonzaga University's Florence campus.

Leo XIV. (2026, May 15). Magnifica humanitas: On safeguarding the human person in the time of artificial intelligence. The Holy See.

[Available in downloadable form from the Vatican.](#)

The Vatican's major statement on artificial intelligence places AI within a broader discussion of human dignity, technology, work, truth, freedom, and the common good. The document is unusually expansive, recognizing both the promises and the potential perils of artificial intelligence.

What's a Liberal Arts Major to Do?

Lawrence, F. M., & Spector, A. (2026, September 8). The synthesis is here: Our future depends on how successfully we fuse the liberal arts and technology. *The American Scholar*. Phi Beta Kappa.

<https://theamericanscholar.org/the-synthesis-is-here>

There is no greater comfort for a liberal arts major than the Phi Beta Kappa society, dedicated since 1776—there's that year again—to the importance of broad study. Turns out, liberal arts and science majors will be essential for AI to understand itself. Lawrence is the society's president. Spector is a computer scientist with appointments at Carnegie Mellon, MIT and consultant to multiple major industrial players in tech and finance.

*Thanks to ChatGPT and Lumo (Proton's AI assistant) for help with proofreading, citation formatting, and finding the Acemoglu article and the Gordon-LaBlanc interview.