

Manufactured Intelligence: Why AI Is Not a Technology Story

Shay O'Kelly · August 2026

For all of human history, intelligence has been the one economic input you could not buy more of. Energy scaled, first with wood, then coal, then oil, then the grid. Land scaled with conquest and agriculture. Labor scaled with population. Capital scaled with finance. But cognition, the thing that actually invents everything else, only ever scaled at the speed of biology. If you wanted more thinking in the world, you needed more people, raised over decades and educated over decades more. Every civilization, company, and research lab has operated under that constraint, so completely that we stopped noticing it was a constraint at all. AI turns intelligence into a manufactured good, with a price, a supply curve, and a capex line. Framed that way, the story is bigger than technology. It removes what has always been the binding constraint on growth, which is why I would put it with industrialization, maybe even agriculture, on the short list of true economic regime changes.

To see why, it helps to notice that this has happened twice before, in slower motion. Intelligence has gone through three regimes. The first was biological. For millions of years, cognitive capacity improved at the pace of evolution, and knowledge died with the brain that held it. The second regime was cultural. Language, then writing, then printing let knowledge outlive its host and accumulate across generations. This is why economic growth for most of history tracked population. More brains meant more ideas, and better information technology meant less rediscovery of what someone had already figured out. The printing press did not make any individual smarter, but it changed the replication economics of knowledge, and within a few centuries came the scientific revolution and then the industrial one. Each regime change was not an intelligence upgrade. It was a bandwidth upgrade. Which is what makes the third regime so violent as a discontinuity. A trained model is cognition itself, not a description of it, and it copies at essentially zero marginal cost. Culture let us copy what a mind had learned. Synthetic intelligence lets us copy the mind.

The standard history of AI, from the 1956 Dartmouth workshop through the winters of the 70s and 80s to the deep learning breakthrough of 2012¹, is usually told as a story of scientific dead ends and eventual insight. I see it differently. The core ideas won almost nothing on novelty. Neural networks date to the 1940s and 50s, backpropagation was published in usable form in 1986², and the researchers who carried connectionism through the winters were mostly right the whole time. What they were missing was nine orders of magnitude of compute cost decline. The AI winters look less like failures of imagination than long stretches where the hardware economics had not caught up to hypotheses sitting in the literature for decades. When compute got cheap enough, in GPU clusters originally built to render video games, the old ideas started working almost immediately. That reading matters because it tells you what kind of problem

intelligence turned out to be. It was never a mystery waiting for a genius. It was an industrial input waiting for its cost curve.

The scaling laws made that explicit. Starting around 2020, researchers showed that model capability improves as a smooth, predictable function of compute, data, and parameters, holding across many orders of magnitude³. It is hard to overstate how strange and how important that is. Capability in every prior technology came from invention, which is lumpy and unschedulable. Capability in AI comes, to a first approximation, from spend. The scaling laws turned artificial intelligence from a research problem into a procurement problem, and the industry noticed. The defining questions in AI today are megawatts, fab allocation, and financing structures, which is why the org charts of AI labs increasingly resemble energy companies. When the smartest thing on earth improves predictably with dollars, the interesting questions all become questions about dollars, land, power, and silicon.

So where is this heading? My view, and this is the claim the rest of these essays build on, is that both capability and diffusion will outrun what consensus assumes, for two separate reasons. On capability, the scaling laws have not bent, and the newer ingredient, letting models think longer and act as agents rather than answer single prompts, compounds on top of pretraining rather than replacing it. Nothing in the data says we are near a ceiling. On diffusion, I think Wall Street is using the wrong reference class. Consensus models AI adoption like enterprise software, a slow S-curve gated by IT budgets and change management. But when what you are buying is labor rather than a tool, adoption is not gated by budget cycles, it is gated by the cost delta, and the cost delta between a knowledge worker and an API call is not twenty percent, it is often a hundred to one. Even if capability froze today, current models are so under-deployed relative to what they can already do that diffusion alone locks in years of growth. The gap between what already exists and what is actually used may be the largest unpriced fact in the market. And capability is not freezing.

None of this requires believing in imminent superintelligence. It requires believing three things that I think are now well-evidenced. Intelligence has become manufacturable, its quality improves predictably with investment, and its output is a substitute for the most expensive input in the economy. If those hold, then the constraint on growth shifts from how many smart people exist to how much intelligence-producing capital we can build. Which turns the most important question in economics into an embarrassingly physical one, what does it actually cost to build the factories? That is the subject of the next essay, and the numbers are bigger than almost anyone is ready for.

SOURCES

1. The Dartmouth Summer Research Project on Artificial Intelligence (1956), and Krizhevsky, Sutskever, and Hinton, ImageNet classification with deep convolutional neural networks (2012).
2. Rumelhart, Hinton, and Williams, "Learning representations by back-propagating errors," *Nature*, 1986.
3. Kaplan et al., "Scaling Laws for Neural Language Models" (2020), and Hoffmann et al., "Training Compute-Optimal Large Language Models" (2022).

No Offer or Solicitation. The only purpose of this site is to provide general background information on RAXIA. This is not an offer to sell or the solicitation of any offer to buy any securities. Any such offer would be made only to qualified investors after distribution and return of the Private Placement Memorandum and governing agreement for the applicable Fund (the "Offering Materials"). Nothing on this site is investment advice or a recommendation of any security. The author may hold positions in companies discussed. Past performance is not indicative of future results.