



You Are Early

How early we still are in the new AI era —
and what to do while almost no one has started.

CODEAUSTRAL

© 2026 CodeAustral LLC. All rights reserved. First edition.

Typeset in Fraunces and Source Serif · youareearly.com

*Written, designed, and typeset with the help of AI agents — which is, in itself, half
the point of this book.*

Contents

Introduction

Part I • Where We Actually Are

- 01 The Curve We're On
- 02 The Deployment Gap
- 03 What Early Looked Like Before
- 04 The Internet Rhyme
- 05 Counting the Adopters
- 06 The Capability Overhang
- 07 Why It Feels Late (But Isn't)

Part II • What Being Early Changes

- 08 Leverage
- 09 Compounding
- 10 The Half-Life of Skills

- 11 Building With Agents
- 12 Distribution in an AI World
- 13 Pricing the Unpriceable
- 14 The Solo Operator Era
- 15 Taste Becomes the Moat
- 16 The New Division of Labor

Part III • What To Do About It

- 17 Positioning Yourself Early
- 18 Mental Models for the Frontier
- 19 The Traps
- 20 A 90-Day Plan
- 21 Building Things That Compound
- 22 The Long Game
- 23 You Are Early

Closing

Notes & Further Reading

PART I

Where We Actually Are

The Curve We're On

Every transformative technology follows an S-curve. The only question that matters is where on it we are standing right now.

The hard part about living through a technological revolution is that it never feels like a revolution at the beginning. It feels like products breaking, people arguing, demos failing, valuations floating away from reality, incumbents pretending not to be worried, and early adopters speaking a private language that sounds half technical and half religious. The future, when seen up close, usually looks less like destiny than inconvenience. It arrives with bugs, missing standards, bad interfaces, and a surplus of confident people being wrong in both directions. Only later, after the curve has climbed, do we redraw the story as if everyone could see it coming.

The S-curve is the simplest map we have for this confusion. A new technology begins slowly, with experiments, amateurs, labs, hobbyists, and skeptics. Then something changes: cost falls, reliability rises, distribution improves, and a use case becomes ordinary enough that adoption compounds. The line steepens. Finally, after the technology saturates the obvious markets, growth slows again and the curve flattens into infrastructure. From a distance this shape is clean. From inside it, each phase is easy to mistake for another.

The beginning feels fake because the absolute numbers are small. The middle feels like the beginning because the rate of change is suddenly visible. The end feels permanent because the technology has disappeared into habit. Electricity, printing, personal

computers, the web, and smartphones all followed some version of this pattern. None moved at the same speed, and none produced only good outcomes. But all punished the same mistake: looking at a young deployment curve and assuming it was a mature technology.

Artificial intelligence, especially large language models, now sits in the most deceptive part of that map. The public has noticed. The tools are useful enough to change real work. The arguments are loud enough to make everyone tired. Yet the actual deployment of the technology into institutions, products, education, medicine, law, government, science, and daily operations is still early. This is the strange position of the mid-2020s: AI feels everywhere in conversation and nowhere near everywhere in practice. That gap is the opportunity.

The Shape Beneath The Noise

An S-curve is not magic. It is what happens when adoption depends on learning, trust, complementary tools, and social proof. At first, a technology is expensive, fragile, strange, and used by people with unusually high tolerance for friction. They do not adopt because the thing is easy; they adopt because they can already imagine the world it points toward. Early printing shops were dangerous businesses. Early electrical systems were local, incompatible, and uneven. Early personal computers required patience that would now be considered a form of punishment.

Then the technology improves enough that people who are not romantics can use it. This is the bend in the curve. The technology is still imperfect, but the balance changes: the pain of using it becomes lower than the pain of not using it. Once that happens, adoption starts feeding on itself. More users produce more investment, more investment produces better tools, better tools produce more users, and eventually the thing that once seemed

exotic becomes the default assumption around which other things are built.

The steep part is often misread because human beings feel change in speed, not level. When adoption goes from one percent to two percent, it has doubled, but most people can ignore it. When it goes from ten percent to twenty percent, it still doubles, but now everyone has a story. A colleague uses it. A competitor announces something. A child asks why the old way exists. The atmosphere changes before the infrastructure does, and the atmosphere is what makes the period feel like a climax.

But the steep part is not one moment. It can last years or decades, depending on what has to change around the technology. Electricity was not simply invented and then adopted. The dynamo mattered, but so did factories, motors, meters, wiring, safety codes, financing, and new ideas about how work should be organized. For a long time, factories used electric motors as if they were merely cleaner steam engines. The larger productivity gains came only when owners redesigned the factory around distributed power instead of one central shaft.

This distinction matters because the first visible wave is usually the shallowest form of adoption. New technologies are first used to do old things a little faster. Printing first reproduced familiar kinds of books before helping create new publics, new professions, and new religious and scientific conflicts. The web first looked like brochures on screens before it became search, marketplaces, social networks, software delivery, and the operating surface of modern commerce. The first use is rarely the final use. It is the bridge by which the old world walks into the new one.

Why The Middle Feels Like The Start

The steep section of an S-curve feels like the beginning because most of the visible disorder is genuinely new. Standards are unsettled. Winners are unclear. Interfaces change monthly. People make grand claims and then ship clumsy products. The serious work is mixed together with theater, and it takes effort to tell one from the other. This is not a sign that a technology is fake. It is often a sign that the technology is leaving the controlled environment where it was easiest to understand.

Consider the early web. By the late 1990s, it was already famous. Magazines had covers about it, companies were adding dot-com to their names, and teenagers were building pages with blinking text. To anyone inside the frenzy, it could feel as if the internet had arrived fully formed. Yet the essential parts of the later web were still missing or immature: broadband was not universal, mobile access was primitive, search was still improving, payments were awkward, cloud infrastructure barely existed in the modern sense, and most businesses had not rebuilt themselves around networks.

The same pattern appeared with the personal computer. By the early 1980s, the PC was culturally visible. It had stores, magazines, clubs, and an expanding mythology. But the world had not yet absorbed what cheap, general-purpose computation on every desk would mean. Spreadsheets were revolutionary, but they were only one layer. Graphical interfaces, local networks, desktop publishing, software ecosystems, email, databases, and eventually the internet all extended the meaning of the machine. The PC looked like a product before it became an environment.

AI is in a similar perceptual trap. Because millions of people have tried chatbots, because executives mention AI on earnings calls, because students use it for essays, because software developers use coding assistants, it is tempting to say that the thing has already

happened. But public awareness is not the same as deep deployment. Trying a tool is not the same as reorganizing a department around it. Adding a chatbot to a website is not the same as redesigning the service model of a company.

The early steep section is where language runs ahead of reality. People say "AI company" before they know what the company does differently. They say "agent" before the agent can be trusted with much responsibility. They say "automation" when what they mean is a clever draft requiring supervision. This imprecision is annoying, but it is also predictable. In the middle of adoption, vocabulary is part of the construction site. Words are scaffolding before they become architecture.

Where Language Models Actually Are

By the mid-2020s, large language models had crossed the threshold from laboratory curiosity to general-purpose tool. That is not a small statement. A technology becomes general-purpose when it can be applied across many domains without being rebuilt from scratch for each one. LLMs can draft, summarize, translate, explain, classify, brainstorm, code, tutor, search through messy text, and serve as a natural-language interface to other systems. They are not equally good at all of these things, and they are not reliable enough for every setting. But their range is already unusual.

Their adoption, however, is uneven. Individuals are often ahead of institutions. A designer uses an image model quietly. A lawyer asks a model to summarize a contract but does not tell the firm. A programmer uses a coding assistant for boilerplate and tests. A teacher experiments on Sunday night, then teaches on Monday inside a school system that has not yet decided what the rules are. This is how many technologies enter the world: first as shadow practice, then as sanctioned workflow, then as infrastructure.

Organizations are slower because organizations are made of risk. They need permissions, budgets, procurement, compliance, security reviews, training, integrations, and someone to be responsible when the model says something wrong. The gap between a stunning demo and a production system is not a failure of imagination. It is the difference between possibility and liability. A model that can impress a room for five minutes may still be unsuitable for a regulated process that runs ten thousand times a day and must leave an audit trail.

This is why the claim "AI is already everywhere" and the claim "AI has barely begun" can both be true. AI is everywhere as an object of attention. It is not everywhere as a redesigned operating layer. Most companies have not rebuilt their org charts, software stacks, data flows, customer experiences, training programs, and measurement systems around it. Most schools have not rebuilt assessment around a world where fluent text is cheap. Most governments have not rebuilt services around conversational access to expertise. Most professionals are still deciding which parts of their work are judgment, which parts are pattern, and which parts were only habit.

The current moment is therefore not the end of the AI transition. It is closer to the moment when electricity had left the fairground and entered the factory, but before the factory had been redesigned. It is closer to the moment when the web had left the university and entered the office, but before commerce, media, and social life had been rearranged around it. The tools are visible. The deeper changes are still negotiating their way through budgets, institutions, norms, law, and fear.

The Capability Curve And The Deployment Curve

One of the most important distinctions in this book is the difference between the capability curve and the deployment curve.

The capability curve describes what the technology can do under favorable conditions. The deployment curve describes where those capabilities have been reliably installed in the world. The two curves are related, but they are not the same. Confusing them is the source of much bad thinking about AI.

Capabilities can improve in sudden jumps. A new model appears and writes better code, handles longer context, reasons more consistently, sees images, hears audio, speaks naturally, or calls tools. The improvement is easy to feel because the interface is direct: you ask, it answers. This intimacy creates the impression that the world itself has changed at the same speed. But the world is not an interface. The world contains hospitals, courts, payroll systems, legacy databases, union contracts, insurance policies, reputations, and people who have been burned by software migrations before.

Deployment moves through a thicker medium. It requires trust, and trust is earned differently in different domains. A marketing team can tolerate a model that produces ten mediocre headlines to get one good one. A radiology workflow cannot tolerate that same looseness. A software team may accept a coding assistant that speeds up routine work while still reviewing every line. An aviation maintenance system requires another level of assurance entirely. The same underlying technology can be revolutionary in one setting, premature in another, and irrelevant in a third.

Historical analogies help because they remind us that capability often arrives before transformation. The dynamo made electric power possible, but factories did not instantly become modern. The printing press made replication cheaper, but literacy, distribution networks, censorship battles, religious conflict, scientific correspondence, and commercial publishing evolved over long periods. The transistor made smaller electronics possible, but decades of manufacturing, design, software, and

supply chains turned that possibility into personal computers and smartphones. A capability is a seed. Deployment is the ecosystem that lets it grow.

AI's capability curve is moving quickly enough to disorient everyone. Its deployment curve is moving more slowly, but probably faster than many past general-purpose technologies because software distributes quickly and the interface is language. That combination produces the central weirdness of the present: the tool can be copied instantly, but the habits required to use it well cannot. You can give every employee access to a model in an afternoon. You cannot give every employee judgment, taste, context, security discipline, and process redesign by Thursday.

The Insider's Mistake

Insiders are usually wrong in a particular way. They overestimate the next six months and underestimate the next ten years. This is not because they are foolish. It is because they are too close to the capability curve. They see what works in the lab, what almost works in the demo, what a better version might do, and what the current system could become if the remaining problems were solved. Their imagination is trained on potential.

Short-term overestimation comes from compressing all the hard parts into the phrase "just integrate it." Anyone who has built software knows how much trouble can hide inside that word. Integrating AI into a real workflow means connecting data sources, permissions, prompts, evaluations, user interfaces, escalation paths, monitoring, billing, latency budgets, and failure modes. It means deciding what the model is allowed to know, what it is allowed to do, and who is accountable when it does something unexpected. These are solvable problems, but solvable is not the same as solved.

Long-term underestimation comes from the opposite error: assuming that current constraints are permanent. People look at today's hallucinations, costs, latency, awkward interfaces, and governance problems and conclude that the technology will remain trapped inside them. That is usually a bad bet with a general-purpose technology. Early systems are not merely worse versions of later systems; they invite the creation of complementary inventions. The web did not become useful to billions only because HTML improved. It became useful because browsers, search engines, payments, cloud servers, smartphones, advertising systems, and social habits evolved around it.

AI will have its own complements. Some will be technical: better retrieval, better evaluation, better memory, better tool use, better security, better hardware, better model routing, better ways to represent uncertainty. Some will be organizational: new job designs, new training norms, new procurement standards, new liability frameworks, new forms of supervision. Some will be cultural: a better understanding of when to ask a model, when to distrust it, and how to combine machine fluency with human responsibility. The mature AI era will not be today's chatbot with a nicer logo. It will be an environment of tools and assumptions built around cheap cognition.

This is why both evangelists and skeptics can be directionally useful and temporally wrong. The evangelist sees the long-term slope but expects tomorrow to behave like the destination. The skeptic sees today's failures but mistakes them for fixed boundaries. The practical person studies the gap. The gap is where businesses are built, skills are acquired, institutions are redesigned, and fortunes of competence are made.

Chaos Is Not Maturity

The present AI moment is chaotic in ways that can make it look more mature than it is. There are already conferences, consultants, lawsuits, enterprise platforms, policy debates, billion-dollar infrastructure plans, open-source factions, benchmarks, safety arguments, and workplace anxieties. The conversation is so large that it feels late. But volume is not maturity. A crowded argument at the beginning of an era can sound like the final debate only because everyone has arrived before the building is finished.

In young markets, confusion is information. It tells you that the old categories no longer fit. Is an AI assistant software, labor, search, a tutor, a colleague, a database interface, a creative partner, or a compliance risk? The honest answer is often yes. New technologies blur boundaries before they create new ones. The early automobile was a carriage without a horse, a mechanical toy, a danger to pedestrians, a luxury object, a farm tool, and eventually the basis for suburbs, highways, logistics, motels, fast food, policing patterns, and oil geopolitics.

AI is producing the same category trouble. A language model is not a person, though it can sound like one. It is not a database, though it can retrieve and synthesize information when connected to one. It is not a traditional program, though it can execute instructions through tools. It is not a search engine, though many people use it as one. It sits awkwardly across familiar boundaries, which is precisely why it matters.

The early steep section also creates a strange social pressure. People feel behind because the headlines move faster than their own lives. They imagine that everyone else has already mastered the tools, automated their business, retrained their team, and discovered secret workflows. In reality, most people are experimenting, exaggerating, or waiting. Even sophisticated

organizations are often running pilots that would look fragile if shown plainly. The surface area of adoption is broad, but the depth is still limited.

This matters for morale as much as strategy. If you believe the market is mature, every confusion feels like personal failure. If you understand that the market is early, confusion becomes part of the work. You are not late because the tools are changing. You are standing in the period when learning has unusually high leverage. The people who benefit most from new curves are rarely those who predicted them perfectly. They are the ones who stayed close enough to the technology to notice when experiments hardened into habits.

The First Uses Are Not The Final Uses

The early uses of a general-purpose technology tend to be imitative. We use the new thing to reproduce the old thing because the old thing is what we understand. The first printed books imitated manuscripts. Early movies borrowed the grammar of theater. Early websites borrowed the structure of brochures, catalogs, and magazines. Early AI products borrow the shape of chat, search boxes, help desks, writing tools, and coding assistants.

This is not a criticism. Imitation is how adoption begins. People need familiar handles. A blank new medium is too much to ask of ordinary users, so the first products translate novelty into an existing mental model. The danger is assuming that the translation is the destination. A chatbot is one interface to AI, not the final form of AI. The keyboard and monitor were not the final form of computing. The browser page was not the final form of the internet.

The more interesting AI applications will often be less theatrical than today's demos. They may not look like a person talking. They

may look like a medical practice that handles documentation without consuming the physician's evening. They may look like a manufacturer that turns maintenance logs into preventive action. They may look like a small law firm that can serve clients previously priced out of careful work. They may look like software that adapts to the user instead of forcing the user to adapt to menus designed years earlier.

Some of the largest changes may be invisible because they will occur inside workflows. A model drafts the first version of a report, checks it against source material, routes exceptions to a human, updates a system of record, and learns from corrections. No single moment feels like science fiction. Yet a process that once required days may take hours, and a service that once required expensive coordination may become available to smaller firms or poorer customers. Transformations that matter economically are not always transformations that look dramatic on stage.

This is one reason the AI era is easy to underestimate. The public imagination searches for robots, superintelligence, and cinematic rupture. Much of the real impact will be duller and larger: paperwork compressed, expertise distributed, interfaces simplified, software made more malleable, small teams given leverage, and ordinary workers surrounded by tools that reduce the penalty for not knowing where to begin. The future often enters through the back office before it changes the skyline.

The Price Of Being Early

Being early is not the same as being right. This distinction matters because early markets are full of intelligent failure. Many people in the first wave of railroads, electricity, automobiles, radio, aviation, personal computing, and the internet lost money, time, or reputation. They saw something real and still chose the wrong company, the wrong timing, the wrong architecture, or the wrong

customer. The existence of a curve does not eliminate execution risk. It concentrates it.

The early AI market will be no different. Some models will disappoint. Some startups will vanish. Some enterprise deployments will produce more slide decks than savings. Some jobs will be changed in painful ways, and some predictions will age badly. There will be fraud, waste, regulatory backlash, security failures, and a great deal of software that should never have been shipped. A serious argument for being early has to include these facts, not step around them.

But the existence of waste does not disprove the curve. The dot-com bubble did not mean the internet was overhyped in the long run. It meant prices, expectations, and business models had run ahead of deployment. After the crash, the boring work continued: broadband spread, logistics improved, cloud infrastructure emerged, search got better, online payments normalized, and software moved into the browser. The bubble was real. So was the future.

AI will likely produce similar confusion. Some periods will feel overheated; others will feel disappointing. Progress will not arrive as a smooth line of better press releases. It will come through strange combinations of breakthroughs, plateaus, product refinements, cost reductions, legal decisions, customer learning, and integration work. A person who requires the curve to move cleanly will quit too early. A person who expects chaos and studies it carefully has a better chance of remaining useful.

The price of being early is that you must act without full social confirmation. By the time everyone agrees, the easy terrain is gone. This does not mean rushing blindly. It means building the habit of small, serious experiments before the playbook is obvious. Use the tools. Measure what improves. Notice where they fail.

Learn the new vocabulary without being captured by it. The goal is not to believe every claim. The goal is to develop judgment before judgment becomes expensive.

The Work Of Locating Yourself

The practical question is not whether AI will change everything. That phrase is too large to be useful. The practical question is where your work, your organization, your customers, and your skills sit relative to the curve. Some domains are already in the steep section. Software development, marketing drafts, customer support triage, research assistance, design exploration, translation, and internal knowledge work have all felt real pressure. Other domains are still waiting on reliability, regulation, cost, or trust.

To locate yourself, ask what part of your world depends on language, pattern recognition, routine judgment, or the transformation of messy information into usable output. That is a large territory. Much white-collar work is not the production of final truth but the conversion of ambiguity into drafts, summaries, recommendations, decisions, and next actions. LLMs are built for the early stages of that conversion. They are less like a finished expert and more like a tireless, uneven apprentice with access to a strange amount of memory.

Then ask where deployment is blocked. Is the bottleneck model quality, or is it data access? Is the problem accuracy, or is it that no one has rewritten the workflow? Is the real obstacle regulation, customer trust, internal politics, procurement, security, or simply lack of imagination? Many people use "AI is not ready" as a single sentence for several different problems. Some are real limits of the technology. Others are limits of organizational design.

This distinction separates action from commentary. If the model cannot do the task reliably, wait, test, or narrow the task. If the

model can help but the workflow has not changed, redesign the workflow. If employees are using tools secretly, create rules that make good use easier than bad use. If customers distrust automation, make the system transparent and give humans clear authority. If the economics do not work today, watch the cost curve. The right response depends on the true location of the bottleneck.

Being early, in this sense, is not a mood. It is a diagnostic. It means the gap between capability and deployment is still wide enough for individuals and small teams to cross it before institutions fully standardize the route. It means learning still compounds. It means taste still matters. It means the field has not hardened into credentials, procurement checklists, and best practices guarded by committees. There is inconvenience in that, but also freedom.

The Beginning After The Beginning

Every era has a moment when people confuse familiarity with completion. They have heard enough about the technology to become bored by it, but they have not yet lived in the world it will create. This is a dangerous moment for attention. The novelty fades before the consequences arrive. The words become clichés before the institutions change. People stop looking precisely when looking becomes useful.

AI is approaching that moment. The first shock has passed. Many people have seen the magic trick. Some have been impressed, some irritated, some threatened, and some disappointed. The second phase is less glamorous and more important. It is the phase of changing workflows, rebuilding software, training people, testing reliability, discovering new business models, and deciding what humans should remain responsible for. This is where the curve begins to matter economically and morally.

The printing press did not merely make books cheaper. It changed who could argue, who could learn, who could organize, and who could challenge authority. Electricity did not merely replace candles and steam shafts. It changed the architecture of work, the rhythm of cities, the design of homes, and the expectations of modern life. The web did not merely publish documents online. It reorganized media, commerce, friendships, politics, and the distribution of software. General-purpose technologies begin as tools and end as assumptions.

Large language models are not guaranteed to follow the most dramatic version of that path. History does not hand out destiny. There are real limits, and there will be corrections. But the structure is visible: a broad capability, rapidly improving, cheap enough to spread, flexible enough to touch many fields, and early enough in deployment that most institutions have not yet absorbed it. That is not the top of the curve. It is the beginning after the beginning.

The best news about this moment is that it is still legible. You can still see the machinery of change. You can still separate capability from deployment, hype from use, and theater from infrastructure. You can still learn faster than the rules are written. The curve we are on will not wait for perfect certainty, because no important curve ever has. The only question that matters is where we are standing right now, and the honest answer is this: far enough along to know AI is real, early enough that what you do next can still matter.

That was chapter one of twenty-three.

The rest of the argument — the whole case for why it's still early and exactly what to do about it — is in the full book.

459 pages · \$19 · youareearly.com