

Solve Problems with AI

Everyday AI

VOLUME I

Basics of Artificial Intelligence

NOBEL KHANDAKER

Everyday AI

volume 1

Basics of Artificial Intelligence

Dr. Nobel Khandaker

First Edition

Intramotev Press

Copyright © 2026 Dr. Nobel Khandaker

PUBLISHED BY INTRAMOTEV PRESS

INTRAMOTEV.COM

All rights reserved. No part of this publication may be reproduced, stored or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise without written permission from the publisher. It is illegal to copy this book, post it to a website, or distribute it by any other means without permission.

First edition, 2026

Contents

1	Why AI, Why Now	1
1.1	The Fastest Adoption in History	2
1.2	A New Literacy	3
1.3	The Hype and the Shift	4
1.4	What This Book Promises	5
2	What Is AI and an LLM?	7
2.1	Four Words, Nested Inside Each Other	7
2.2	The One Task: Guess the Next Piece	9
2.3	Why It Sounds So Sure When It Is Wrong	11
2.4	Does It Remember Me?	11
2.5	From Answering to Acting: Agents	12

List of Figures

1.1	How fast each technology reached a mass audience	2
2.1	The buzzwords, nested inside each other	8
2.2	The next-token loop	10
2.3	A chatbot answers; an agent acts	12

*Dedicated to everyone curious enough to ask
what AI can do for them.*

1

Why AI, Why Now

ON A GREY TUESDAY in March, in a high school on the flat farmland south of Columbus, Ohio, a history teacher named Maria Delgado sat at her desk after the last bell and did something that, two years earlier, would have cost her an entire Sunday. She had thirty-one students and one short letter: a complaint written by a coal miner in 1894, four hundred words of cramped, furious nineteenth-century prose. Six of her students read two grades above their level. Four read several grades below it. The letter, as written, would reach perhaps a third of the room.

She typed an instruction into a chat window: rewrite this letter at five reading levels, from fourth grade to college, keep every historical fact identical, and give me a short glossary for the youngest version. Then she watched a small bar slide across her screen. In the time it took to finish her coffee, she had five versions of the letter, a glossary that explained what a “scab” was in 1894, and her Sunday back.

Delgado is not a programmer. She has never written a line of code. What she has is a new kind of literacy—the ability to ask a machine, in plain English, to do a piece of cognitive work and to judge whether it did the work well. That literacy is the subject of this book. It is spreading faster than almost any practical skill before it, and it is already changing what an ordinary person can get done in an afternoon.

How long to reach a mass audience?

Telephone	#####	~75 yrs
Electricity	#####	~30 yrs
Television	#####	~13 yrs
The Web	#####	~7 yrs
Smartphone	#####	~6 yrs
ChatGPT	<-	2 months to 100 million

(illustrative only; bars mix saturation and user counts)

Figure 1.1: An illustrative comparison, not a precise one: the older technologies are measured by how long they took to reach most households, and ChatGPT by raw user counts. The metrics differ, but the shape of the contrast is real—conversational AI spread at the speed of a link shared in a text message.

1.1 The Fastest Adoption in History

It helps to measure the thing. When a new technology spreads, it usually spreads at the speed of factories and salesmen and copper wire. The telephone took roughly seventy-five years to reach most American households. Electricity took decades. Even the smartphone, which felt instantaneous, needed the better part of a decade to saturate the country.

ChatGPT, the chat program that introduced most of the world to artificial intelligence, reached one million users in five days.¹ It reached one hundred million in two months—a pace that made it, at the time, the fastest-growing consumer application ever measured. By the autumn of 2025 it counted eight hundred million people using it every week. By February of 2026, nine hundred million.² In May of 2026, by one widely cited estimate, the app passed a billion users in a single month.³ And ChatGPT is only one product among several; Google’s Gemini and Anthropic’s Claude command audiences of their own.

¹TechnologyChecker.io. *ChatGPT Statistics 2026: 900M Users and Beyond*. <https://technologychecker.io/blog/chatgpt-statistics>. Accessed: 2026-06-22. 2026

²TechCrunch. *ChatGPT reaches 900M weekly active users*. <https://techcrunch.com/2026/02/27/chatgpt-reaches-900m-weekly-active-users/>. Accessed: 2026-06-22. 2026

³Sensor Tower. *ChatGPT Crosses 1 Billion Monthly Active Users*. <https://sensortower.com/blog/chatgpt-1-billion-mau>. Accessed: 2026-06-22. 2026

Numbers this large can numb rather than inform. The point is not the bragging rights of a single company. The point is the slope of the line. A technology that reaches a billion people in three years does not wait politely for schools to write a curriculum or for workplaces to draft a policy. It shows up first and asks permission later. The miner's letter reached Delgado's students in March because the tool was already on her laptop, not because anyone had trained her to use it.

Mustafa Suleyman, who helped found two of the most important AI laboratories and now runs Microsoft's AI division, calls this speed a "containment problem."⁴ A technology powerful enough to matter and cheap enough to copy spreads faster than any rule can chase it. He frames the deeper shift in economic terms. The Industrial Revolution made physical power—the work of muscles and horses and water wheels—cheap and abundant. This revolution, he argues, is doing the same thing to *cognitive* power: the work of drafting, summarizing, translating, and reasoning. When the cost of something falls far enough, the people who once rationed it begin to spend it freely, and the world reorganizes around the new abundance.

1.2 A New Literacy

There is a useful precedent. In the middle of the fifteenth century, a goldsmith in Mainz named Johannes Gutenberg assembled a press with movable type, and the price of a book began a long fall toward zero. Before the press, reading was a specialty trade. A handwritten Bible could cost a laborer a year's wages, and literacy was the property of clerics and clerks. After the press, reading slowly became a precondition for ordinary life—for worship, for commerce, for law, for citizenship. The person who could not read was not merely uneducated. They were shut out of the conversation the rest of society was having.

WE ARE LIVING through a comparable shift, and the word that has attached itself to it is *AI literacy*. The analogy is not perfect, and it is worth being honest about where it breaks. The printing press distributed a fixed

⁴Mustafa Suleyman and Michael Bhaskar. *The Coming Wave: Technology, Power, and the Twenty-first Century's Greatest Dilemma*. Crown, 2023

thing—a book said the same words to every reader. An AI system produces something new each time, tuned to the question asked, and it sometimes produces confident nonsense. Reading a printed page rarely misleads you about whether the words are really there. Reading an AI’s answer requires a second skill the medieval reader never needed: the habit of checking.

Still, the core of the analogy holds. A barrier that once stood between an ordinary person and expert-level work is dropping. The owner of a small bakery can draft a marketing plan that would once have required an agency. A nurse can get a dense insurance document translated into plain language in seconds. A student stuck on a proof at eleven at night can summon a tutor who never sleeps and never sighs. None of this makes the agency, the lawyer, or the teacher obsolete. It changes the *entry fee* for cognitive work—the amount of training and money a person needs before they can begin.

This is the promise, and it is real. But a promise is not a guarantee, which brings us to the part of the story the marketing tends to skip.

1.3 The Hype and the Shift

Every powerful technology arrives wrapped in two kinds of exaggeration. The first says it will save us; the second says it will destroy us. Artificial intelligence has attracted both in industrial quantities. It has been sold as a cure for cancer, a tutor for every child, and a replacement for half the workforce, sometimes in the same press release. Some of this will prove true. Much of it will not. The careful reader learns to separate the durable shift from the noise around it.

So what actually changed? For most of computing’s history, a computer did exactly what it was told, in a language built for machines. To get a useful answer you had to translate your messy human intention into rigid code. The shift, beginning in earnest around 2022, is that the translation step is gone. You can now describe what you want in ordinary words—“summarize this contract,” “find the flaw in my argument,” “explain photosynthesis to a ten-year-old”—and a machine will make a genuine, if imperfect, attempt. That is the whole revolution in a sentence: the way we ask a computer for help became a conversation. What it does not mean is that the answer is automatically true. The machine is a fluent assistant, not an

oracle, and much of this book is about telling the difference.

The evidence that this matters is not in the marketing; it is in the behavior. In Britain, a 2026 survey of full-time undergraduates found that ninety-five percent used AI in some way, and nearly as many used it to help with graded work.⁵ American students are not far behind. In the workforce, the numbers are lower but climbing fast: by August of 2025, roughly thirty-seven percent of employed Americans reported using AI tools on the job.⁶ The use split sharply along lines of education and income—college graduates used these tools at more than twice the rate of workers without degrees. Stanford’s annual index of the field documents the same story from every angle—capability, investment, and use all climbing steeply at once.⁷ That split is the quiet warning inside the good news. A tool that multiplies the ability of whoever uses it tends to widen the gap between those who use it well and those who never learn to.

1.4 What This Book Promises

This book will not make you a programmer or a data scientist. It will not explain the calculus inside a neural network, and it will not pretend that these tools are either miracles or menaces. Its aim is narrower and more useful: to make you fluent enough to fold AI into the ordinary work of your life and to stay in charge while you do it.

Fluency, here, rests on four skills, and the chapters that follow are built around them. You will learn to *read* better—to turn a forty-page document into a conversation, and to do it without being fooled by a confident summary of a page that does not exist. You will learn to *write* better—to use the machine as a drafting partner and an editor without surrendering your own voice. You will learn to *learn* better—to turn a chatbot into a patient tutor that meets you where you are. And you will learn to *solve problems*

⁵Higher Education Policy Institute. *Student Generative Artificial Intelligence Survey 2026*. <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2026/>. Accessed: 2026-06-22. 2026

⁶Federal Reserve Bank of St. Louis. *The State of Generative AI Adoption in 2025*. <https://www.stlouisfed.org/on-the-economy/2025/nov/state-generative-ai-adoption-2025>. Accessed: 2026-06-22. 2025

⁷Stanford Institute for Human-Centered AI. *The 2026 AI Index Report*. <https://hai.stanford.edu/ai-index/2026-ai-index-report>. Accessed: 2026-06-22. 2026

better—to break a hard task into pieces a machine can help with and to check its work before you trust it.

Two habits will recur on every page. The first is *prompting*: the small craft of asking well, which turns out to be most of the game. The second is *verifying*: the discipline of never taking a confident answer on faith. A tool that is right ninety percent of the time is dangerous precisely because it is so often right; the skill is knowing how to catch the other ten percent. We will build both habits slowly, starting in earnest in Chapter 3, and they will carry you long after the particular products in this book have been replaced by their successors.

Maria Delgado did not get her Sunday back because she understood the mathematics of a language model. She got it back because she knew how to ask, and how to check what came out. Before she handed the rewritten letter to her youngest readers, she read it against the original and confirmed that the machine had kept the facts straight and invented nothing. That combination—the ask and the check—is the new literacy. The rest of this book is about learning to do both without thinking about it, the way you once learned to read.

2

What Is AI and an LLM?

THERE IS A PARLOR TRICK that the comedian and computer scientist alike have enjoyed for years. You type the first half of a familiar phrase into your phone—“the early bird gets the”—and the phone, without being asked, offers the word “worm.” Nobody believes the phone knows anything about birds. It has simply seen the phrase often enough to guess what tends to come next. Hold that humble image in your mind. It is, with a great deal more scale and subtlety, the whole secret of the technology this book is about.

The phrase you hear everywhere—*artificial intelligence*—is doing a lot of work, and most of the confusion around it comes from using one big word for several different things. So it is worth slowing down and taking the terms apart, one at a time, because once the words are clear the machine stops feeling like magic and starts feeling like something you can actually use.

2.1 Four Words, Nested Inside Each Other

The terms you will hear—artificial intelligence, machine learning, deep learning, large language models—are not rivals. They are nested, like wooden dolls, each one a special case of the one before it.

Artificial intelligence is the oldest and broadest term. It dates to the nineteen-fifties, and it means, loosely, getting a machine to do something we would call intelligent if a person did it—play chess, recognize a face,

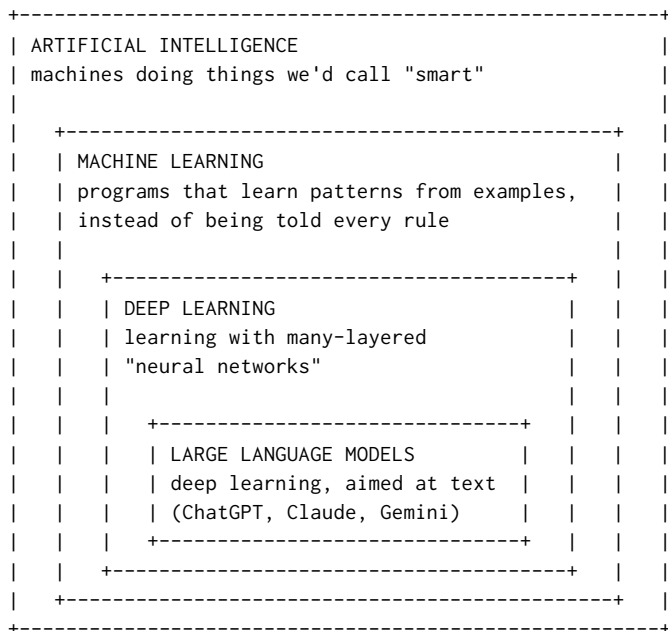


Figure 2.1: The buzzwords are not competitors; they are nested. Every large language model is a kind of deep learning, which is a kind of machine learning, which is a kind of artificial intelligence.

plan a route. For decades, the dominant approach was to write the rules out by hand. A programmer who wanted a machine to spot spam email would sit down and list the telltale signs.

Machine learning flipped that arrangement around. Instead of writing the rules, you show the machine thousands of examples—spam and not-spam—and let it work out the rules for itself. Nobody tells it what to look for. It infers the pattern from the pile of examples. This is why these systems are “trained” rather than “programmed,” and it is the single most important idea to carry forward: modern AI learns from examples rather than from instructions.

Deep learning is machine learning done with a particular tool called a *neural network*—a web of simple mathematical units, loosely inspired by the neurons in a brain, stacked in many layers. The word “deep” just means “many layers.” Each layer passes its output to the next, and with enough layers and enough examples, the network can learn strikingly complex patterns. The brain metaphor is convenient but misleading; a neural network is no more a brain than an airplane is a sparrow. It is mathematics that happens to learn.

A *large language model*, finally, is a very large neural network trained on an enormous quantity of text. “Large” is not a boast; it is a description. These models contain billions of internal numbers, called *parameters*, that get tuned during training—and the largest models contain far more. Feed in a slice of the written world—books, articles, code, conversation—and the model adjusts those parameters until it becomes extraordinarily good at one narrow task.

2.2 The One Task: Guess the Next Piece

That narrow task is the autocomplete trick from the start of this chapter, scaled up almost beyond recognition. A large language model does exactly one thing: given some text, it predicts what piece of text is likely to come next. That is all. Everything else—answering questions, writing essays, translating French, drafting code—is that single ability, dressed in different clothes.

To do this, the model first chops text into *tokens*: small chunks, each roughly a short word or a piece of a word. “Basketball” might become two

```

YOU TYPE: "The capital of France is"
      |
      v
+-----+
| break into tokens |
+-----+
      |
      v
+-----+
| the model scores | Paris .... 91%
| every possible   | the .... 3%
| next token       | a .... 1%
+-----+ ...
      |
      v
pick a likely token: "Paris"
      |
      v
GROWS TO: "The capital of France is Paris"
      |
      +---> feed the whole thing back in,
              and guess the NEXT token...
              (repeat, token by token)

```

Figure 2.2: The entire machine, one loop. The model guesses one token, adds it to the text, and feeds the longer text back in to guess the next. A whole essay is this step, run hundreds of times.

tokens, basket and ball; a comma is its own token. The model works in tokens the way you and I work in syllables. When you talk about how much a model can read at once, or what it costs to use, you are really talking about tokens.

Here is the engine, stripped to its frame:

The model does not write a sentence and then check it. It builds the answer one token at a time, each guess shaped by everything written so far. The breakthrough that made this work at scale arrived in 2017, in a research paper with the immodest title “Attention Is All You Need.”¹ It introduced an architecture called the *transformer*, whose key trick is letting the model weigh how strongly each token should relate to the earlier to-

¹Ashish Vaswani et al. “Attention Is All You Need.” In: *Advances in Neural Information Processing Systems*. <https://arxiv.org/abs/1706.03762>. 2017

kens in the passage—which words, in other words, deserve its “attention.” The letter T in ChatGPT and in countless other products stands for that transformer. You do not need the mathematics. You need only the consequence: a machine that, having absorbed a vast slice of digitized text, has become an uncanny predictor of what comes next.

2.3 Why It Sounds So Sure When It Is Wrong

This design explains the technology’s most important flaw, and we meet it early because everything later depends on understanding it. The model is built to produce *plausible* text—text that fits the pattern—not *true* text. Most of the time, plausible and true travel together, because the writing it learned from was mostly trying to be true. But not always. Its basic tendency is to continue plausibly, not to check whether the continuation is correct. Chat products are now trained to hedge—to say “I’m not sure”—more often than the bare model would, but the underlying pull is always toward the fluent guess, which can be a confident, fluent, entirely invented fact.

This behavior has a name—*hallucination*—and it is not a bug that a future update will simply delete. It is the shadow side of the very thing that makes the model useful. A system that always guesses the most fitting next words will sometimes guess wrong while sounding exactly as assured as when it is right. Chapter 7 returns to this in detail. For now, hold one rule close: confidence and correctness are different things, and the model controls only the first.

2.4 Does It Remember Me?

A natural question, the first time a chatbot answers you by name, is whether it knows you. By default, the raw model does not. Each conversation begins as a blank slate; the model does not carry memories from one chat to the next, and it learns nothing permanent from talking to you. Its knowledge was fixed when its training ended.

That said, the honest answer today is “it depends on the product.” The companies have wrapped the forgetful core in features that imitate memory: a setting that lets the tool keep notes about you across conversations,

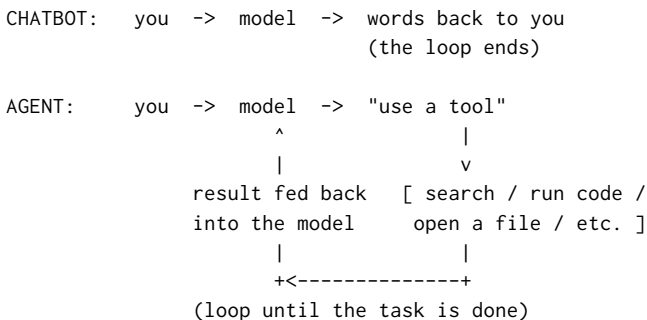


Figure 2.3: A chatbot talks. An agent acts: it can call a tool, read what comes back, and keep going on its own until the job is finished.

“projects” that hold a standing set of documents and instructions, and accounts that quietly carry your preferences from session to session.² These are conveniences bolted on the outside, not a change in the nature of the model underneath, and—as Chapter 7 will press—they are also a privacy decision worth making on purpose rather than by accident.

2.5 From Answering to Acting: Agents

For its first few years, the public face of this technology was the chatbot: you type, it types back, and nothing happens in the world. The newer and faster-moving frontier is the *agent*—a language model given hands.

An agent is an LLM connected to *tools* and allowed to use them in a loop. Rather than only describing how to do something, it can search the web, open and read a file, run a calculation, fill out a form, or write and execute a small program—then look at the result and decide what to do next. The model still only predicts text; the trick is that some of the text it predicts is an instruction to a tool, and the system carries that instruction out and feeds the result back in.

This is a genuine leap in usefulness, and a genuine leap in risk. A chatbot that invents a fact wastes your time. An agent that acts on an invented fact—deletes the wrong file, sends the wrong email, trusts the

²Anthropic. *Claude’s memory and projects*. <https://www.anthropic.com/news>. Accessed: 2026-06-22. 2026

wrong website—can do something you cannot take back. The power and the hazard rise together, which is why this book introduces agents here, returns to them practically in Chapter 3, and treats their dangers squarely in Chapter 7.

For now, the mental model is enough, and it is smaller than the jargon suggests. Strip away the words and you are left with one machine that has read a great deal and learned to guess what comes next, sometimes wrapped in a memory, sometimes handed a set of tools. Everything you will do with it in the chapters ahead—reading, writing, learning, solving problems—is a way of pointing that guess at something you care about, and then checking where it lands.