

AI Is Easy: AI Models

A free chapter — the whole set at ai-is-easy.com

ALEX POOLMAN · THOUGHTS IN ACTION

Contents

Welcome

Chapter 1: What an AI model really is

Chapter 2: The different kinds of model

What happens next

Welcome

What is this little book about?

This book is about the AI helpers everyone keeps talking about. The ones you type a question into, and they answer you back. You may have heard names like ChatGPT or Claude or Gemini. You may have nodded along while someone explained them, and quietly thought, I have no idea what any of this means. Good. You are in the right place, and you have missed nothing.

Do I need to know anything first?

Not one thing. You do not need to be young. You do not need to be clever with computers. If you can read this page, you have everything you need. I will not use a single big word without explaining it in the very same breath.

Why does everyone make this sound so hard?

Because it is easy to sound clever and hard to sound clear. Most books and videos about AI are written by people who already understand it, for people who already understand it. That leaves out almost everybody. This book does the opposite. If a ten-year-old could not follow a sentence, that sentence does not belong here.

Is everything in here the whole truth?

It is the honest short version, which is a different thing, and I would rather say so on the first page. Every explanation in this book is a simplification chosen to be useful rather than complete. Where a simplification could mislead you, I say so in the same breath. Where a claim comes from somewhere, you will find the source at the back.

How should I read it?

However you like. The book is built out of questions. Real ones, the kind you would ask out loud at the kitchen table. Each question gets a short, plain answer. Read it front to back, or hunt for the one question that has been bothering you and read-only that. Both are fine.

Is there anything to do, or only to read?

There is something to do, and I would push you towards it. Most chapters end with one small thing to try, in a shaded box. Reading about this teaches you surprisingly little. Ten minutes of trying teaches you a great deal.

What will I be able to do at the end?

You will know what these helpers actually are, why they sometimes get things wrong, and how to pick the right one without asking anybody. That is the whole promise. Nothing more, and nothing less.

Chapter 1: What an AI model really is

What is the sort of AI model this book is about?

Almost every helper you can type a question into is built around something called a large language model. That is a computer program fed an enormous amount of writing. It has become very good at one particular trick. It works out what small piece of text should come next.

Not a word?

Not quite a word, and the difference matters later. These programs work in small chunks called tokens. A common word is usually one chunk. A long or unusual word gets broken into two or three. For now, picture it guessing the next word, and remember that the real unit is a little smaller.

That sounds too simple. Is that really all it does?

It is remarkably close to all it does, and I want to show you why that is enough. Imagine you find a page from a film script. It is a scene between a person and their helpful assistant. You can read what the person says. But the assistant's reply has been torn off the bottom of the page. Now imagine you also own a strange machine. You feed it any piece of writing, and it tells you which word would sensibly come next. One word. That is all it gives you.

How would that machine finish the script?

You would feed it the torn page. It would offer you one word. You would add that word to the page and feed the whole thing back in. It

would offer you the next word. You would do that again, and again, and again, and slowly a full reply would appear on the page. Not copied from anywhere. Built up, one piece at a time, each piece chosen because it fitted what came before.

And that is what happens when I use one of these helpers?

That is essentially what happens. When you type a question into one of these helpers, the computer is holding a script. Your question is on it. And it is filling in what a helpful assistant would say next, one piece at a time, faster than you can blink. Hold on to that picture. It explains a great deal of what these helpers do, and we will come back to it.

So it is not looking up the answer somewhere?

Not by itself, and this is the part almost nobody tells you. It is not searching a library. It is not checking a list of facts. It is continuing a piece of writing in the way writing usually continues. Many of these helpers can now go and search the internet as well, and when they do they will normally tell you. But left on its own, it is producing the shape of an answer rather than fetching one.

Then why is it right so often?

Because it has absorbed a great many patterns from a great deal of human writing, and much of that writing is sound. That is why the answers are often useful. It is not a guarantee. The same pile of writing also contains errors, out of date facts, fiction, arguments and nonsense, and the machine has no way of separating those. Sounding right and being right are different things, and it is very good at the first.

How much writing did it read?

More than a person could read in many lifetimes. You will see dramatic figures quoted for this, and most of them depend on assumptions nobody states, so I will not print one. The honest version is enough: you are not talking to something that skimmed a few books.

Is the model thinking while it answers me?

Not in the way you and I think. There is no little person in there weighing things up. There is no worry, no opinion, no wanting. There is a very large amount of arithmetic, working out which piece of text fits best. It feels like thinking from the outside, because the result reads like a person wrote it.

Why do people call it a "model" and not a program?

Because it was not written the way normal programs are written. Nobody sat down and typed out rules for it to follow. Instead it was shown mountains of writing and slowly shaped itself around what it saw. A bit like a violinist who never learned to read music, but has heard ten thousand hours of it. A model here means something shaped by example rather than by instruction.

Is this the same thing people call AI?

It is one part of it. AI is the big outer circle: the whole idea of machines doing things that used to need a person. Inside that circle is a smaller one, machine learning, which means machines that get better by looking at examples instead of being given rules. Inside that is a smaller circle again, deep learning, which is the method behind most of what you hear about today. The helpers in this book sit inside that innermost cir-

cle. Now you know how the four words nest, like bowls stacked in a cupboard.

Do I need to remember all of that?

Honestly, no. Remember one sentence from this chapter. The helpers in this book are very well-read machines that predict what text comes next. Much of this book grows out of that.

Try this now. Ask one of these helpers to explain how it produces an answer, in words a ten-year-old would follow. Then ask it what its explanation left out. The second answer is usually the more interesting one.

Chapter 2: The different kinds of model

Is every AI model one of these text machines?

No, and this is worth clearing up early, because the word model gets used for very different things. What they share is that all of them learned from examples rather than from rules. What they do with that learning varies enormously.

What are the main kinds I will actually meet?

Five, and you will bump into all five without necessarily noticing.

Text models.

The subject of this book. You give them words, they give you words back. Everything you would call a chatbot or an assistant is one of these with a doorway attached.

Image models.

You describe a picture and one appears. These work quite differently underneath. Most of them do not predict the next piece of text at all. They start from visual noise and tidy it, step by step, until it matches your words. Some of the newest ones are built more like text models, but the family is genuinely separate.

Speech models.

Two directions. One turns your voice into text, which is what happens when you talk to your phone. The other turns text into a voice, which is what happens when a helper reads its answer aloud. Both are now very

good, and both are usually bolted on beside a text model rather than being part of it.

Reasoning models.

Text models that have been trained and set up to work a problem through before they answer, rather than replying immediately. Chapter nine deals with these properly. Think of them as the same family with a different working habit.

The models you never see.

The largest group by far, and the one nobody writes books about. The thing that decides which post you see next. Whether a payment looks like fraud. Whether a scan shows something worth a second look. Whether an email is junk. None of these write anything. They score, sort and classify. They are AI models too.

Can I have that as a table?

Here it is, and it is the only table in the book.

Kind	What you give it	What you get back	Where you meet it
Text	Words	Words	Chat helpers, writing tools
Image	Words, or a picture	A picture	Picture makers, photo editing
Speech	Voice, or words	Words, or a voice	Dictation, read-aloud, voice chat
Reasoning	A hard problem	A worked answer	The thinking setting in a chat helper
Scoring	Data	A judgement or a ranking	Fraud checks, recommendations, spam filters

Which one am I using?

Usually several at once without realising, and that is the modern trick. When you photograph a letter and ask a helper to explain it, an image ability reads the page and a text model writes the explanation. When you talk to it and it talks back, a speech model sits at each end. The one doorway hides three or four machines.

Does knowing this change anything?

It changes one thing, and it saves a lot of confusion. When somebody tells you that AI cannot count, or that AI invents facts, ask which kind they mean. Those are true of text models and they are not general laws of the field. Most sweeping statements about what AI can and cannot do fall apart the moment you ask which sort of model is meant.

Try this now. Think of three places where a computer already makes a judgement about you: what you are shown, what you are offered, what gets blocked. Ask a helper which kind of model is likely doing each. It is a good way to see how much of this is already around you.

What happens next

Is this the whole book?

No. This is the opening of one of four. The full book runs to twenty-four chapters, and the other three do the same job for AI agents, for Claude, and for getting started from nothing at all.

What is in the rest?

Why it says things that are simply not true, and how to catch it. What it does not know. How to ask so it understands you. What to keep out of it. How to choose one without agonising. And a fortnight of small, harmless things to try.

Where do I get them?

At ai-is-easy.com. Each book is under three dollars, and all four together cost less than a cup of coffee and a sandwich. PDF and EPUB, instantly, no account needed.

Alex Poolman

Thoughts in Action, August 2026