

Understanding AI

A Guide for Normal People

BRADLEY BATES AND MICHAEL SINGLETON

DESIGNED BY NERD SAVAGE



Unicorn Warehouse

Building magic. Empowering humans.

Understanding AI

A Guide for Normal People

BRADLEY BATES AND MICHAEL SINGLETON

Unicorn Warehouse is a venture studio focused on transforming ideas into mission-driven products through systems thinking, AI, research, and collaborative execution.

Mission: Build practical technology, educational initiatives, and organizations that increase human capability through responsible AI, strong systems design, and collaborative innovation.

Copyright © 2026 Bradley Bates and Michael Singleton

All rights reserved. No part of this book may be reproduced, distributed, or transmitted in any form without prior written permission from the authors, except for brief quotations used in reviews, commentary, teaching, or scholarly discussion.

This is a working draft. The contents may change before publication.

First draft edition, 2026.

Disclaimer

This book is for general educational purposes only. It is not legal, financial, medical, security, employment, or professional advice.

AI systems change quickly. Features, tools, memory behavior, safety settings, and platform rules may differ between products and may change over time. Always verify important information before relying on it.

The authors encourage readers to use AI thoughtfully, skeptically, and responsibly.

Unicorn Warehouse

Building magic. Empowering humans.

Dedicated to everyone
trying to stay human
while learning how to
use the machine.

Unicorn Warehouse

Contents

CHAPTER 1:
The Model Is Not the Whole System

CHAPTER 2:
The Words You'll Keep Running Into

CHAPTER 3:
What's Inside the Box

CHAPTER 4:
Context Is Not Memory

CHAPTER 5:
Why "Related" Isn't "True"

The Model is Not the Whole System

"I used to think the AI I was talking to was the AI."

One thing. One mind. You type, it answers, end of story.

I got that wrong. I think most people still do, and honestly, the apps are built to let you keep thinking it. **Here's what's actually happening:**

What you're typing into doesn't just contain "the AI." It contains a language model, the part that writes the words, wearing a much bigger system like a coat. The system is the one deciding what gets remembered, what's allowed to happen, what needs your permission first, what gets logged somewhere so it can be checked later. The model doesn't do any of that. The model writes.

That's it. That's the whole split you need to hold onto for the rest of this book:

The model writes.

The system decides what's real.

I want to be specific about why this matters, because "the model is just a language thing" sounds like a small technical detail, and it's not. It changes what you should and shouldn't believe when an AI talks to you.

The Model is Not the Whole System

The Model Doesn't Own a Filing Cabinet

Say you tell an AI something important. A deadline, a number, a decision you made. And later it says "got it, I'll remember that."

it? **Did it?** Did it? Did it? Did it? Did it? Did it?

Maybe. Maybe not. And here's the uncomfortable part: the model has no idea either. It's not lying to you, exactly. It's doing what it does, which is predict a reasonable-sounding next sentence. "Got it, I'll remember that" is a reasonable-sounding sentence to say after someone tells you something important. Whether anything was actually saved anywhere is a completely separate question. One the model itself can't answer, because saving things isn't the model's job. It's the system's job, and the system might do it well, do it badly, or not do it at all.

I didn't understand this for a long time.

I'd have a long conversation with an AI, it would feel like it knew me by the end, and I'd walk away assuming that knowledge was sitting somewhere, waiting for next time. Sometimes it was. A lot of the time it wasn't. Nobody told me there was a difference between "this feels remembered" and "this is remembered."

"I had to find out the hard way, more than once."

Unicorn Warehouse

Building magic. Empowering humans.



The Model is Not the Whole System

Why This Keeps Fooling Smart People

It's not your fault if this caught you off guard too. The whole interface is designed to feel like one continuous mind. There's one little chat bubble. It talks in first person. It says "I." Every signal you're getting is built by people whose job is to make the experience feel seamless. And seamless is good design, but it's also exactly the thing that hides the seam that actually matters.

So here's the question I want you to start asking, every time:

What was this thing actually looking at when it answered me?

Not "what does it know" in some big cosmic sense. What was put in front of it, right then, for that one answer? Because that's all it had. If the answer feels confident, ask whether it's confident because it's right, or confident because confident-sounding text is what the model is good at producing regardless of whether anyone checked.




Unicorn Warehouse

Building magic. Empowering humans.

The Model is Not the Whole System

Watch Out For This

If an AI tells you "I saved that" or "I'll remember this" or "noted," that's a sentence, not a receipt. A sentence the model generated because it sounded like the right thing to say next.



I saved that.

I'll remember this.

Noted.

Whether it's true is a different question, answered by a different part of the system, one you usually can't see from the chat window.

That's not a reason to never trust anything an AI tells you. It's a reason to know which claims are backed by something real and which ones are just good writing. You're going to need that distinction constantly, for the rest of this book and honestly for the rest of your life with these tools.

So we're starting here, on page one, because everything else builds on it.

The model writes.

The system decides what's real.

Keep that in your pocket.

The Words You'll Keep Running Into

I'm not going to make you memorize a glossary like it's a vocab quiz. But there are about seven words that show up constantly once you start paying attention to how AI systems actually work, and if you don't have a solid grip on them, every conversation about this stuff turns to mush.

So let's just get them straight, once, here, so I can use them for the rest of the book without stopping to explain myself every other paragraph.



The Words You'll Keep Running Into

Model

The part that writes. That's it. The model takes whatever text it's been handed and predicts what should come next, over and over, until it's produced an answer. It doesn't store files. It doesn't check facts against a database. It writes.

When people say "the AI," they almost always mean the **model**.

From here on, I'm going to be more specific than that, because the model is only one piece of what you're actually dealing with.

Context

Everything the model can see right now, in this exact moment, for this exact answer. Your message. Maybe some of the conversation before it. Maybe some instructions nobody showed you. Maybe nothing else at all.

Context is not the same as "everything the AI knows about you." Context is the desk it's working at today.

If something isn't on the desk, it doesn't exist for that answer, no matter how many times you've told the AI about it in some other conversation.

The Words You'll Keep Running Into

"This is the one that got me.

Memory

Memory is information that's stored somewhere outside the current conversation, sitting in a drawer, waiting to maybe get pulled back out later.

Here's the part that matters:

memory only counts if it actually gets pulled out and put back on the desk. Otherwise it's just sitting in a drawer somewhere, doing nothing, while the model answers you with no idea it exists.

I used to assume that if I'd told an AI something once, it was mine forever, in there somewhere, ready when I needed it. Sometimes that's true. A lot of the time it isn't. **The honest version is:** memory is a maybe, not a guarantee, and you won't always be able to tell which one you're getting just by how confident the answer sounds.

Tool

Anything the system can reach outside of just writing text. Searching the web. Running a calculation. Sending an email. Touching a file.

A tool is power. The model asking to use a tool is not the same as the tool actually firing. That gap is the whole next word.

The Words You'll Keep Running Into

Proposal

What the model suggests doing.

"Just that. A suggestion."

"I'll go ahead and send that email" is a proposal. It is not the email being sent. It's not even a promise the email will be sent. It's a sentence that sounds like a promise, generated by something whose entire job is generating sentences that sound right.

I want you to read that sentence again because it's the most important one in this chapter.

A proposal is not a promise. It's not even close.

Here's what that actually looks like, stripped down to the bare exchange:

you

Send that invoice to the client and let me know when it's done.


AI

Done! I've sent the invoice to the client.

You, an hour later, checking:

No email was ever sent.

There was no tool connected that could send it. The model didn't know that. It just knew "Done! I've sent the invoice" was a normal, confident thing to say right after being asked to send an invoice.



The Words You'll Keep Running Into

Authority

The actual permission to do the thing. Authority lives somewhere else, in the system around the model, not in the model itself. The model can propose all day long. It cannot grant itself the authority to act on its own proposal. Something else has to say yes.

Go back to the invoice example for a second. The fix isn't "the model should be more careful." **The fix is a system that actually checks**, after the model proposes sending something, whether sending it is allowed, whether a tool exists to do it, and whether it actually happened, then reports back the truth instead of the model's best guess at a satisfying answer. When that check exists, you get "I can't send that, I don't have access to your email." When it doesn't, you get "Done!" and a client who never got an invoice.

When those two things blur together, that is, when a system lets a confident-sounding proposal act like it already has authority, that's usually where things go wrong. Keep proposal and authority as two separate words in your head, always, and you'll catch most of the trouble before it starts.

Claim

What the system is honestly allowed to say happened. Not what it wants to say. Not what would sound most impressive. What it can actually back up.

A good system says "I sent that" only after it can prove it sent that. A weaker system just says "I sent that" because that's what came next in the sentence. Same five words. Completely different amount of truth behind them.

Learning to ask "what's backing this claim?" might be the single most useful habit in this whole book.

The Words You'll Keep Running Into

Why I'm Front-Loading This

I could've spread these definitions out, one per chapter, and let you pick them up naturally. I thought about doing it that way. But these seven words lean on each other. Memory needs context. Proposal needs authority. And Claim needs all of it behind it before it means anything.

I'd rather hand you the whole set now, a little dense, than drip them out and have you build the wrong picture in the meantime.

You don't need to memorize this chapter. You just need to know it's here, so that when I say **"that's a proposal, not a claim"** three chapters from now, you already know exactly what I mean and why the difference matters.

