



How AI Actually Works

From room-sized computers to a helper in your pocket

A plain-English tour — no math, no jargon, no hype.

“Useful AI · minimal exposure · human verification.”

A LITTLE HISTORY

Where It All Started

1950

A famous question

Mathematician Alan Turing asks: "Can machines think?"

1956

"AI" gets its name

A summer workshop at Dartmouth College coins the term artificial intelligence.

1960s–80s

The rule-book era

Computers follow hand-written rules: "IF this, THEN that." Smart — but rigid.

1997

Machine beats champion

IBM's Deep Blue defeats the world chess champion, Garry Kasparov.

Keep in mind: back then a computer filled an entire room — and still couldn't hold a conversation.

A LITTLE HISTORY

How It Evolved

2011

Computers start talking

Siri arrives on the iPhone; IBM's Watson wins Jeopardy! on live TV.

2012

Learning by example

Instead of rules, we show computers millions of examples — and they learn patterns themselves.

2017

A better reader

A new design (the “transformer”) lets computers read whole sentences at once — the T in ChatGPT.

2022

The chatbot moment

ChatGPT arrives; 100 million people try it in two months. AI becomes a conversation.

The big shift: from following rules → to learning from examples → to holding a conversation.

The LLM: A Very Well-Read Guesser

LLM = Large Language Model. It read a giant library, and learned one skill supremely well: guessing the next word.

Try it yourself — finish this sentence:

“Peanut butter and _____”

- You said “jelly” — because you've heard it a thousand times.
- The LLM does the same, from billions of sentences.
- Every answer is built this way — one likely word at a time.



Our digital intern

It predicts — it doesn't know.
That's why it can be “confidently wrong.”

RAG: The Open-Book Test

RAG stands for “Retrieval-Augmented Generation” — *fancy words for: look it up before you answer.*



Closed-book: memory only

- Answers from what it read in training
- Can be out of date
- May fill gaps with confident guesses



Open-book: RAG

- First fetches trusted, current documents
- Answers from what it just found
- Can show its sources — so you can check

1

You ask a question

“What are this year's Medicare enrollment dates?”



2

It looks it up first

Fetches the relevant official pages and documents.



3

It answers — with sources

Now you can click and verify for yourself.

Moderated AI: The Guardrails

Think of the bumpers on a bowling lane — the big-name AI tools are wrapped in safety layers:



Checks what goes IN

Harmful requests get refused; many tools warn you before you share personal details.



Checks what comes OUT

Answers pass through safety filters that catch dangerous or inappropriate content.




People behind the curtain

Safety teams test, correct, and retrain the AI — it's not just left to run wild.

Why it matters to you: stick to well-known, moderated tools — that's where the guardrails are.

HOW IT WORKS • ALL TOGETHER

What Happens When You Press Enter



And all of this happens **in about a second.**

TODAY

AI Is Already All Around You



In your mailbox

Spam filters and bank fraud alerts have quietly used AI for years.



On the road

GPS rerouting around traffic and your weather forecast — both AI.



In your photos

Your phone grouping pictures by face and place? That's AI too.

The point: you've been using AI for years — chatbots just gave it a voice.



Today's chatbots can...

- see and describe photos
- talk and listen out loud
- draft letters and emails
- explain confusing paperwork

TOMORROW

Where It's Heading



From answering to doing

“Agents” that book the appointment or fill out the form — with your approval at every step.



More personal & private

Helpers that remember your preferences and run right on your own phone.



More rules of the road

Governments and companies are writing safety standards — like seatbelt laws, but for AI.

What will never change: it works for you — and the final decision stays human. Yours.



THE WHOLE STORY

Four Ideas to Take Home



The LLM

A very well-read guesser — it predicts words, it doesn't know facts.



RAG

The open-book test — it looks things up first and shows its sources.



Moderation

The bumpers on the lane — real guardrails, but not force fields.



You

The editor-in-chief. Every answer is a draft until a human verifies it.

“Useful AI · minimal exposure · human verification.”