



A NO-NONSENSE GUIDE FOR CURIOUS HUMANS

How AI **Actually** Works



Nine big ideas that explain what's really going on when you talk to a chatbot — told through things you already understand.

Ages 11+

9 ideas · 3 parts

No tech background needed



Write me a poem about my cat, Pixel.

Pixel the cat with eyes so bright, chasing
shadows through the night...

You type. It answers.
But **what is actually
happening** in there?

PART ONE


How it learns

Before AI can answer anything, it has to be built and trained. Here's where it all comes from.



It's really good autocomplete.

The weather today is



AI doesn't "know" the answer. It guesses the **next word** , over and over, like the suggestion bar on your phone — just enormously better.

WHAT'S REALLY HAPPENING

For every word, it weighs **millions of options** and picks one that fits. String enough good guesses together and you get a sentence, a poem, an essay.

It read almost everything.



books, websites, conversations — trillions of words

To learn its guessing game, AI was trained on a **colossal pile of text** — close to everything humans have written online.

WHAT'S REALLY HAPPENING

By seeing patterns repeat billions of times, it slowly learns which words tend to follow which. **It never memorized the library** — it absorbed the style of it.

It reads in chunks, not letters.

s t r a w b e r r y

How many R's are in "strawberry"? Try the buttons.

Count like a human

See it like AI

Ask an AI how many R's are in "**strawberry**" and it can get it wrong. Not because it's dumb — because it never sees the letters.

WHAT'S REALLY HAPPENING

AI breaks words into **tokens** — chunks like *straw·berry*. It works with chunks, so spelling-and-counting puzzles are genuinely hard for it.

PART TWO

How it “thinks”

Now it's trained. Understanding how it generates answers explains its quirks — including the weird ones.



It's a writer, not a librarian.



A Librarian

Looks up the true answer and reads it back to you.

not



A Writer

Writes something that *sounds* right — confident, fluent, sometimes wrong.

AI isn't fetching facts from a shelf. It's **composing** a believable answer on the spot.

WHAT'S REALLY HAPPENING

That's why it can **"hallucinate"** — state a wrong fact with total confidence. It optimised for sounding right, not for being right. **So double-check important stuff.**

It has a desk, not a memory.



the context window — older notes slide off the edge

AI can only hold so much at once — like a desk that fits a limited stack of notes. This is its **context window**.

WHAT'S REALLY HAPPENING

In a long chat, the earliest things you said can **fall off the desk**. That's why it sometimes "forgets" what you told it ten minutes ago.

It rolls the dice.



ask twice, get two different answers

Ask the exact same question twice and you'll often get **different answers**. That's on purpose.

WHAT'S REALLY HAPPENING

Instead of always taking the single most likely word, it adds a little **controlled randomness**. That's what makes it feel creative instead of robotic.

PART THREE

How to use it well

You understand the machine now.

Here's how to actually get great results
from it — and stay in charge.



It's spiky, not smooth.

EASY
FOR AI

HARD
FOR AI

Write a story

Simple math

Explain a
topic

Count letters

AI's skill is a **jagged frontier** . It can ace a task that stumps adults, then trip over something a 6-year-old finds easy.

WHAT'S REALLY HAPPENING

"Hard for humans" and "hard for AI" are **different maps** . You can't assume that because it nailed one thing, it'll nail the next. **Test it.**

You're the director.

VAGUE

"Tell me about dogs."

DIRECTED

"Explain **why dogs wag their tails**

In 3 short bullet points for a 12-year-old

The AI can't read your mind. The clearer your instructions, the better it performs — like directing an actor.

WHAT'S REALLY HAPPENING

Tell it **who it's for, how long, and what style**. Good directions ("a prompt") are the single biggest difference between a meh answer and a great one.

Always make it prove it.



AI gives you a "fact"

sounds confident and clear



You ask: "How do you know?"

check a real source yourself



Now you can trust it

verified, not just believed

AI is a brilliant assistant, not an oracle. Treat its answers like a smart friend's guess: **useful, but worth checking.**

WHAT'S REALLY HAPPENING

For anything that matters — homework facts, health, news — **verify it elsewhere** . You stay the human in charge. That's the whole point.

The whole machine, in nine ideas.

If you remember nothing else, remember these.

01

Smart autocomplete

It predicts the next word, again and again.

02

Read everything

Trained on a huge pile of human writing.

03

Sees chunks

Reads tokens, not letters — so spelling is hard.

04

Writer, not librarian

Composes plausible answers — can be confidently wrong.

05

Desk, not memory

Limited context — old chat falls off the edge.

06

Rolls the dice

A bit of randomness makes it feel creative.

07

Spiky skills

Great at some things, terrible at others.

08

You direct it

Clear instructions → far better answers.

09

Make it prove it

Verify anything that matters. Stay in charge.



You now understand AI better than **most adults.**

It's not magic, and it's not a brain. It's a fast, fluent guessing machine — and now you know how to drive it.

Go ask it something. Then make it prove it.

