

Demystifying AI

How it works — and how to protect your writing

Josh Cottrell · Tech Support, Writers Morning Out

Writers Morning Out Roundtable · July 2026



What is "Generative AI," really?

1

Software trained on huge amounts of existing text and images — not a person, not a search engine.

2

It doesn't “know” facts. It predicts the most likely next word, one piece at a time.

3

Think of it as the world's best-read autocomplete — remarkably fluent, not actually thinking.

| The quick brown fox jumps over the *lazy dog*.

YOU TYPE THE START · GENAI PREDICTS THE REST

How it's built

1

Training data

Billions of examples — books, articles, websites, code — collected before training ever begins. More text than any one person could read in thousands of lifetimes.

2

Pattern learning

The model adjusts billions of internal settings (“parameters”) until its predictions get good. This training run alone can take months and enormous computing power.

3

A frozen snapshot

Knowledge locks in at a training cutoff. No live lookup, no real-time learning — “updates” mean a new training run, not the model teaching itself.

A common mix-up this explains:

Training

Happens once, ahead of time — months of computing, before the tool ever answers anyone.


Chatting with you

Uses the model already built. Nothing retrains in real time while you talk to it – more on that later.

New versions

A model only “improves” when its maker runs a new training cycle and ships an update.

What happens when you type a prompt



When a model “shows its thinking”

Some tools run extra reasoning steps before answering — the same predict-loop from before, just longer, and sometimes shown to you as a visible scratchpad.


What you gain

- ✓ Better at multi-step tasks: plotting logic, timelines, math, structural edits
- ✓ Can catch and correct its own missteps partway through
- ✓ The visible trace shows you how it got there, not just the answer

What it costs

- ✗ Slower — more reasoning steps take real time
- ✗ More tokens generated means it can cost more (API pricing) or burn through usage limits faster
- ✗ More steps still doesn't guarantee correctness — it can be confidently wrong
- ✗ That reasoning trace is still your data, under the same retention rules


Same loop, different length



Standard

predict → predict → answer

~seconds



Thinking mode

predict → check → predict → ... → answer

~10–60 sec, sometimes shown to you

The tools you'll run into



Chatbots

ChatGPT, Claude

Conversational brainstorming, drafting, research, feedback.



Writing assistants

Grammarly, Sudowrite

Built into your draft — line edits, prose suggestions.




Image generators

Midjourney, DALL·E

Covers, marketing art, mood boards.

Using it for research



Ask it to cite sources — and verify them yourself. It can generate citations that look real but aren't.



Sensitive topics can trigger guardrails — researching an illness or treatment may get filtered or hedged.

What GenAI can't really do



It doesn't understand your book the way a person does — no grasp of your themes or intent.



It can be confidently wrong (“hallucination”) — fluent and completely incorrect at once.



It has no memory of you unless a specific tool is built to keep one.

Sounds confident



100%, every time

Actually correct



varies — always verify

The exception: when a tool remembers you


A few tools now break the “no memory” rule on purpose — a Memory feature that recalls past chats, not just the current one.

What you gain

- ✓ Skip re-explaining your project, characters, or style each time
- ✓ Surfaces old conversations you can't quite place
- ✓ Suggestions feel more tailored over time

What it costs

- ✗ More of your writing and personal details persist long-term
- ✗ Details from one project can surface inside another
- ✗ “Import memory from other providers” pools even more in one place
- ✗ One more setting worth checking periodically



Where does your data go?

■ Free / personal tiers

- Often stored and may be reviewed by people, not just machines
- Frequently used to train future models by default
- You can usually opt out in settings — but it's off by default

■ Paid business / enterprise tiers

- Usually excluded from training by contract, not just a toggle
- Often has stricter retention limits
- “Paid” alone isn't a guarantee — check the specific plan

Two numbers worth remembering:

30 days

Typical default retention before deletion, if you haven't opted into training (varies by tool).

5 years

How long Anthropic may retain opted-in Claude chats to help train future models.

No opt-out

Conversations flagged for safety review can still be kept, regardless of your settings.

Privacy risks specific to writers


- Unpublished manuscripts pasted into free tools can be retained and reviewed.
- Publisher and agent contracts increasingly include AI-use clauses — check yours.
- Policies vary a lot by tool — read the fine print rather than assume:

	Tool	Trains on your content?
S	Sudowrite	States it does not train on user writing
G	Grammarly	Does not use your content to train models; shares some usage metadata with AI providers
C	ChatGPT (Free / Plus)	Trains on chats by default — opt out in Settings → Data Controls
CI	Claude (Free / Pro / Max)	Trains on chats by default — opt out in Settings → Privacy

A practical privacy checklist



Check your settings — turn off “improve/train the model” toggles (ChatGPT: Settings → Data Controls; Claude: Settings → Privacy).



Be cautious pasting unpublished manuscripts into free consumer tools.



“Paid” isn't automatically “private” — check the specific plan's policy, don't assume.



Read a new tool's privacy policy once, focused on “training” and “retention.”



Use a strong, unique password and enable two-factor authentication on writing/AI accounts.

It runs somewhere: energy, water, and neighbors

Every prompt runs on physical hardware, in a real building, drawing real power and water — that's the cost behind the screen.

950 TWh

Projected global data-center electricity use by 2030

Roughly double 2025's level — about 3% of world electricity demand. (IEA, 2026)

100,000+ homes

What one large AI data center can draw

Some hyperscale AI facilities now exceed the power use of 100,000 households.

264B gallons

Water used by AI data centers globally in 2025

For cooling — often drawn from regions already under drought stress.

75+ projects

Faced local community pushback in Q1 2026 alone

Mostly over water use and strain on the local power grid. (Data Center Watch)

This isn't a reason not to use it — it's a footprint worth knowing is there.



Thank you

Josh Cottrell · joshcottrell@gmail.com

