

ROUND ROCK WRITERS GUILD  
SHOP TALK  
NOV 2021

## ARTIFICIAL INTELLIGENCE



## AGENDA

- A Bit of History. What is AI?
- Concerns with Using AI.
- How to Use AI
- Intelligence?

A BIT OF HISTORY

## WHAT IS AI?



### ALAN TURING (1912 - 1954)

Groundbreaking mathematician and computer scientist

Responsible for the enigma machine in WWII, which enabled the decryption of nazi codes and was key to the allies winning the war

Proposed the "Imitation Game" (now called the Turing Test) in 1949 as a measure of artificial intelligence

Subject of the movie "The Imitation Game" (if you want to know more)



## The Turing Test:

### The Setup:

The test involves three participants: a human judge, a computer, and another human. They are isolated from each other, often with only text-based communication available.

### The Process:

The human judge asks a series of questions to both the computer and the other human, aiming to determine which is the computer. There are no "domain restrictions" on the questions - they can be about any subject at all.

### The Goal:

The computer's goal is to convincingly imitate human responses so that the judge cannot tell it apart from the human participant.

### Passing the Test:

If the judge fails to reliably identify the computer more than half the time, the computer is considered to have passed the Turing Test.

The Turing Test is not a test of intelligence itself, but rather a test of a computer's ability to mimic human-like conversation and behavior. It's a benchmark for artificial intelligence, and passing it remains a significant challenge.

DOUG LENAT (1950 - 2023)

Computer scientist and AI  
Researcher

In the 1980s, ran the “Cyc” (from  
“Encyclopedia”) project

Believed AI was just a matter of  
giving a computer enough data to  
work with

“Intelligence is 10 million rules”

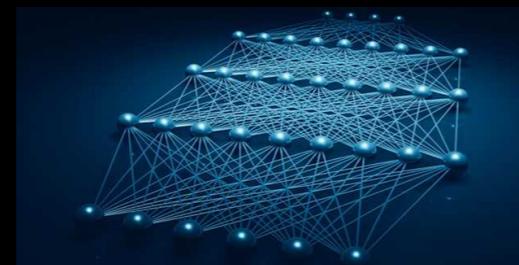


Lenat's key insight was that we humans have a lot of implicit rules in our head. For example, I know Kathi. That means that my date of death can't be before her date of birth, and vice versa. It seems silly to say that out loud, but computers, especially in the 80's, didn't have those rules. Working at MCC, Lenat had graduate students enter information from the Encyclopedia Britannica into “Cyc”, an early AI. Overnight Cyc would integrate information into its core rule base, then the following morning Lenat and his students would query Cyc to determine what new information it had learned. The results were often surprising. One morning Cyc said “Everyone is famous.” Lenat and his students realized that all of the Britannica articles about people that they had fed Cyc were about famous people - ergo Cyc had “concluded” that everyone was famous. That's an early example of AI “Hallucination”, a topic to which we'll return in a moment.

## LENAT WAS RIGHT

AT LEAST UP TO A POINT

### GPUs & Neural Networks



### The World Wide Web



Turns out, Lenat was right. Two things have happened to validate his theory.

- 1) The advent of GPUs and neural networks. GPUs were originally developed to support video game graphics, which rely on something called “vector arithmetic” to generate graphics and animation. Over time, they became cheap enough to be deployed by the thousands in neural networks. Neural networks are networks of computing nodes based on the architecture of the human brain. Those networks rely on the very same vector arithmetic needed for high-end computer games. So, in a sense, we owe the advent of AI to video games. Want to know more about how vector arithmetic works? Check out this YouTube video: <https://www.youtube.com/watch?v=s0Q3CojqRfM>
- 2) The World Wide Web, and the fact that everything is now digital provides an ideal data set for training neural networks. No more need for graduate students to enter data - now just deploy a web crawler. (A web crawler is a program that trawls the internet and indexes its content. It’s how search engines like Google are able to do what they do.)

## AI TERMINOLOGY

- **Machine Learning** - A subset of AI where algorithms improve their performance at a task through data and experience, rather than being explicitly programmed for every rule.
- **Deep Learning** - A subset of ML using multi-layered neural networks to detect patterns in data, especially effective in speech, vision, and text tasks.
- **Large Language Model** - A type of AI trained on vast amounts of text to understand and generate human-like language (e.g., GPT, Claude, Gemini).
- **Transformer** - A neural network architecture that underlies many modern LLMs; uses **attention** mechanisms to understand context in sequences like text. (GPT = "Generative Pre-Trained Transformer". Ironically, in french, phonetically it sounds like "J'ai pété", which means something totally different ...)
- **Training Data** - The dataset used to teach the AI how to perform its task.
- **Fine-Tuning** - Adjusting a pre-trained model on a smaller, specialized dataset to improve performance for a specific use case.
- **Prompt** - The input or question given to an AI to generate a response.
- **Inference** - The process of an AI making predictions or generating outputs after it's been trained.
- **Bias** - Systematic errors in AI outputs caused by imbalances or flaws in training data.
- **Alignment** - Ensuring an AI's goals and behaviors are in line with human values and intentions.
- **Explainability** - The ability to understand how and why an AI made a decision.

“J'ai pété” = “I farted”

## CONCERNS WITH USING AI





There are several issues related to intellectual property and AI.

Almost every major AI firm has been sued for the improper use of copyrighted work in training their models. The AI firms claim “fair use”. The legal situation is still being resolved, with some cases having been decided in favor of the AI firms, and some in favor of the copyright holder. One emerging idea that seems to be taking hold is the creation of a special type of license to allow the AI firms to use copyrighted work. There’s also the question of whether, even if an author was, say, paid for a copy of their book that was fed to an AI, does the manner in which the AI can constantly store and use that information create an additional type of use case? Said another way, should an author be paid more for a copy of their work if it’s used to train an AI?

Similarly, AI can be used to create new material that's closely derivative of existing work, and could be viewed as competing with that work. One could even go as far as uploading an author's work (or a movie, TV show, or whatever) and having the AI generate new material based on it. Say, for example, I'm a huge fan of Dan's "Vampire Camp" series, and I don't want to wait for him to write more. I could upload that to an AI and have the AI write stories set in that world. Clearly, that would be copyright infringement.

One can see why opinions against AI are so prevalent in the creative world.

You can, and people have done, use AI to generate work that you then “publish”, especially with self-publishing being so easy. This is rightly (IMO) viewed as unethical, and many publishers require you to disclose if and how you’ve used AI. More generally, AI-written creative work is substandard. This may change as models get more sophisticated.

To date, no copyright has been issued to an AI.



AI consumes tremendous amounts of power. Training a massive model like GPT-3 can consume ~1300 MWh. That's about what 130 average US homes would use in a year. The operational phase of AI also consumes tremendous power: ChatGPT's daily power usage is equivalent to ~33,000 US households. That's equivalent to a town 1/4 the size of Round Rock. Further, these power needs are increasing rapidly.

A single ChatGPT query uses about 2.9 Wh, while a Google query uses 0.3 Wh. That means a single ChatGPT query uses as much power as my e-Bike uses (in Eco mode) to go 2/3 of a mile (360Wh battery). Generating an image can take the same amount of power it takes to charge a cell phone.

OpenAI estimates that 10% of ChatGPT's power usage is in response to small talk - e.g. "Good Morning". That's a testament to how good these LLM's are - people are anthropomorphizing the technology.

Why does AI use so much energy? Remember those neural networks and their GPUs we talked about? Well, anyone who's familiar with gaming computers knows GPU's are SERIOUSLY power-hungry. In 2023, OpenAI was handling 10 million inquiries daily on an estimate 30,000 NVIDIA GPU's. (OpenAI doesn't publish how many GPUs they have). Today, OpenAI handles over a billion queries daily. All those GPUs doing all that vector arithmetic on all those queries takes a lot of juice.

Anyone remember the panic caused by the "DeepSeek" announcement earlier this year? Talk about unintended consequences. AI is considered sensitive technology, and there are export restrictions on those GPUs to places like China. Innovating in response to this constraint, DeepSeek, based in the city of Hangzhou in the province of Zhejiang, China, released an AI earlier this year that makes MUCH more efficient use of its GPUs by only "activating" the GPUs needed to process a particular input. This resulted in an AI that's much less energy-hungry and, therefore, more cost efficient. The day following DeepSeek's announcement, Nvidia (the foremost maker of GPUs) lost almost \$600 billion in market share. That shows us there's room for innovation in the "demand" side of AI's power needs.



(Text below contributed by ChatGPT.)

AI “hallucinations” happen when a model produces confident but false or unfounded information. This isn’t because the AI is “lying” in a human sense—it’s a byproduct of how large language models (LLMs) work.

---

### Why Hallucinations Happen

1. Prediction, not truth

LLMs like ChatGPT, Claude, DeepSeek, and Gemini are trained to predict the most likely next word based on patterns in their training data—not to verify facts. If the training or context lacks correct information, the model will still produce something that sounds plausible, even if it’s wrong.

2. Incomplete or biased training data

If a concept is underrepresented, outdated, or incorrectly represented in the training data, the model’s output will reflect those gaps or errors.

3. Overgeneralization

Models sometimes “fill in the blanks” by combining unrelated knowledge patterns—similar to a person guessing the ending of a story they’ve never read.

4. Ambiguous or leading prompts

If a question is vague or framed in a way that suggests an answer, the AI may generate speculative content that isn’t grounded in verified data.

5. Lack of real-time fact-checking

Unless connected to a current, authoritative source (like a live search), the model can’t verify whether its statements are still true.

### How to Guard Against Hallucinations

## 1. Use authoritative sources

- Where accuracy matters (e.g., medical, legal, financial advice), connect the AI to trusted external databases or have it cite sources.
- Example: Some AI systems (like Perplexity or SearchGPT) retrieve and show references.

## 2. Ask for verification steps

- Request that the AI:
- Outline its reasoning chain.
- Identify what parts are fact vs. assumption.
- Suggest ways to confirm the claim.

## 3. Cross-check with other tools

- Run the same query in multiple AIs or in traditional search engines.
- Differences in answers often point to areas worth investigating.

## 4. Phrase prompts to encourage caution

- Instead of: "What's the answer to X?"
- Try: "What are the top three credible sources on X, and what do they say?"

## 5. Limit speculative generation

- If you don't want the model to guess, explicitly instruct it:

"If you're not certain, say you don't know."

## 6. Use retrieval-augmented generation (RAG)

- Combine an AI model with a real-time search index or custom database so it grounds responses in up-to-date, verified material.

## 7. Keep humans in the loop

- For high-stakes use, have an expert review outputs before acting on them.

---

### 💡 Rule of thumb:

Treat AI as a knowledgeable assistant with no instinct for truth—great at organizing and expressing information, but requiring you to be the fact-checker. (End ChatGPT response).

AI companies have tried to duck any kind of liability for hallucinations. So far they've been mostly successful, mostly because those who have sued have had a hard time proving damage. That may be about to change. Wolf River Electric is a small solar energy company in MN. When Google's AI spread inaccurate information about the company, it lost ~\$200k in contracts. It is suing Google in federal court over the loss. Unlike other suits which have claimed libel or defamation, WRE has receipts, so it's actually much easier to prove. Should the courts rule in WRE's favor, we may see more legal accountability placed on the AI firms for inaccurate information.

GPT-5, the latest model from OpenAI release in August, 2025, claims a significant reduction in hallucinations.



In some cases (professional work for hire, for example) you may be required to disclose if, and how, you've used AI. Passions around the use of AI in creative spaces are quite high - on both sides of the issue. People don't want to pay for someone to simply parrot an LLM's response back to them.

Also, many readers will take umbrage at an author who uses, or is suspected of using, AI to write their work. A simple google search will turn up many spirited arguments about this.

## PRIVACY

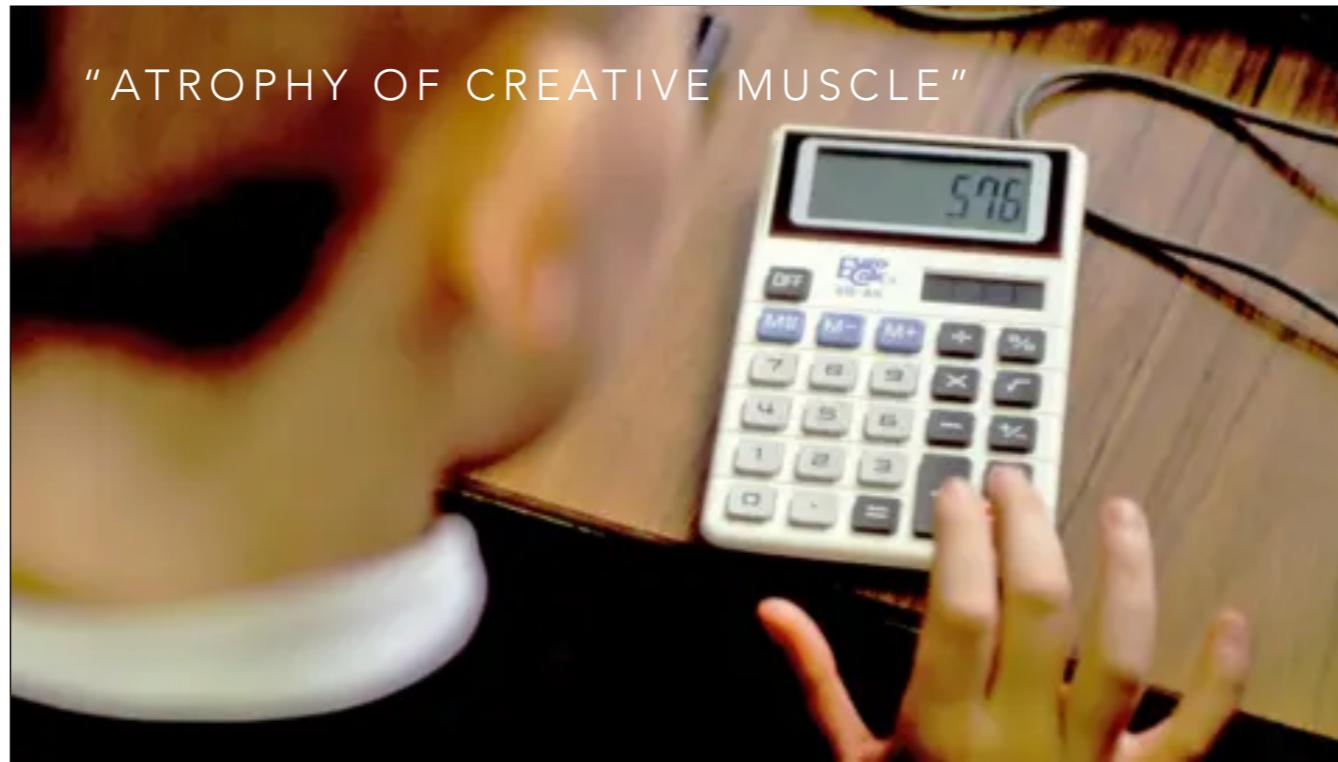


Prompts and their responses are, generally, not private. They aren't completely public, but they can be accessed. OpenAI (ChatGPT's maker) has access to the prompts and responses, and uses them to further train their models.

Anthropic (Claude's maker) has a slightly more strict policy, saying that conversations are private by default, and only a "limited number" of Anthropic employees have access to conversations for "explicit business purposes". Anthropic does not train models on user-entered data. Conversations flagged for violating the usage policy may be reviewed to improve detection of harmful activity, and of course they are also subject to Discovery.

There is no lawyer/client or doctor/patient confidentiality with AI. AI "conversations" are subject to legal discovery, so don't ask it the best way to hide a body.

Bottom line: Don't assume your chats with AI are private.



This one is my term. Sort of like how, way back in the day, our teachers would make us do math without a calculator because “you won’t always have one available”. Well, that turned out to be wrong. But what is true is that people today are much less capable of mental arithmetic than our forebears. Can you quickly answer the question “What’s 17 times 6” without pulling out your phone’s calculator or asking Google? Our forebears understood the distributive property of multiplication over addition, which makes doing this mentally easy. How many of us even understood that previous sentence? I’m not saying grandpa would have known to call it that (I’m certain he wouldn’t), but he was more likely to understand how to use it practically than people today.

BTW, here’s how you do it mentally:  $17 \times 6$  is the same as  $(10 \times 6) + (5 \times 6) + (2 \times 6)$ , or  $60 + 30 + 12 = 102$ .

It’s the same with artistic creativity. If you ask ChatGPT to write a difficult scene, you aren’t developing your ability to write such scenes. In fact, you may well be atrophying any ability you do have by your reliance on the AI. They jury’s out on that from a research perspective, but it seems reasonable. And don’t we want to become better authors, not worse ones?

That said, it might be valid to ask the AI “What makes a good <x> scene?” and consider including the elements it provides as a starting point for your own scene.

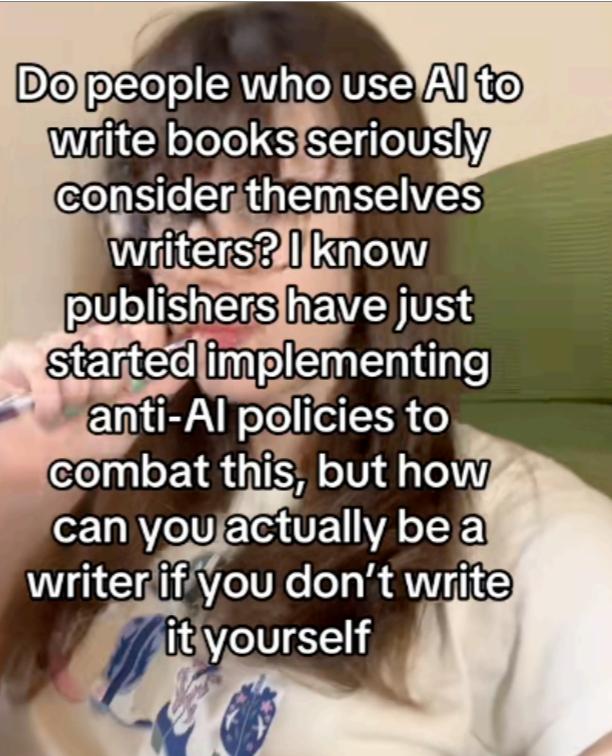
## SHOULD YOU AS AN AUTHOR USE AI?

It's a decision only you can make.

There are risks and ethical concerns, as we've discussed (hopefully respectfully and objectively).

Assuming you think there's a place for AI in your creative process, or if you just want to play around with it, the next section will talk about how to get started.

Screen grab from TikTok summing up a pervasive opinion —>



# HOW TO USE AI



## THE AI LANDSCAPE

There are way, way too many AI firms and products to list here. These are a few of the "big ones".

Model	Core Strengths
ChatGPT	Huge user base, personalized memory, strong in creative and political domains
Claude	Superior reasoning, long context, coding & written communication consistency
DeepSeek	Top structured reasoning, cost-efficiency, transparency, coding correctness
Gemini	Multimodal capabilities, nuanced responses, massive context window

**Observations & Use-Case Guide**

- Choose ChatGPT for user-focused, conversational workflows or creative ideation.
- Choose Claude when you need rigorous analysis, long-form comprehension, or reliable coding output.
- Choose DeepSeek fitting cost-sensitive projects requiring high correctness in logical or scientific domains--especially if open-source access matters.
- Choose Gemini for applications that require multimodal interaction, integration with Google services, or human-affective outputs.

The text at the left was generated by ChatGPT in response to the prompt "Compare the relative strengths of ChatGPT, Claude, DeepSeek, and Gemini."

In this presentation, black text on a white background is reserved to denote output from an AI.

## TIPS FOR AI PROMPTS

- Be specific.
- “Act as if...”
- Tell it how you want the results presented.
- Use “do” and “don’t”.
- Use examples - but be mindful of copyright.
- Consider tone and audience.
- Build on previous prompts.
- Correct mistakes and give feedback.
- Ask it to create your prompts, or what else it needs from you.

Tips from Harvard University Information Technology. <https://www.huit.harvard.edu/news/ai-prompts>.

### **Be specific**

Generic prompts like “Write a story” will produce generic results. What kind of story do you want? What genre? Is it for adults or children? How long should it be? Is it funny or serious? Adding this type of context and being specific, clear, and concise will help generate more useful outputs.

More specificity can also limit the chances of inaccurate responses. If you’re looking for a good meal near Harvard Square, asking “What is the best restaurant in Cambridge?” could return results from Cambridge, UK. Asking “What is the best restaurant in Cambridge, Massachusetts, within walking distance of Harvard Yard” gives you a better chance of finding what you’re looking for.

### **“Act as if...”**

Asking the AI to behave as if were a type of person, process, or object can be an easy way to start generating better prompts. The AI will attempt to emulate that role and tailor its answers accordingly.

For example, imagine you wanted help creating a recipe based on ingredients you have in your kitchen. If you ask, “create a recipe using these ingredients,” followed by a list, you’ll get some general recipe suggestions. But if you added “act as if you are my personal trainer” first, the AI will consider this context in its response, perhaps suggesting a healthier recipe or a meal designed to refuel after a workout.

### **Tell it how you want your output to be presented**

Generative AI can produce many different types of outputs, including code, stories, reports, summaries and synopses, dialogue, business communications, audio, images, and much more. Being specific about the type of output you want will produce better results. After describing what you want, try adding “Present this in the form of...” and your preferred output. Or start your prompt with, “Create a [format of output] about / that contains...” etc.

### **Use “do” and “don’t”**

Telling AI what you do and don’t want in your response can save time and improve your result. To use the recipe example again, you can specify what foods you do want to include in your meal, and any that you specifically do not want, e.g. “Act as if you are my personal trainer. Create a recipe that will help me refuel after my workout. Do include tomatoes, chicken, and some type of carb. Don’t include chili peppers or any ingredients containing wheat.”

### **Use examples**

Using an example of a sentence or paragraph that you’d like your output to resemble can give the AI a reference point to help improve results. However, it’s very important to not infringe on copyright. Don’t upload copyrighted works and ask the AI to rewrite / reproduce it or create a substantially similar work.

### **Consider tone and audience**

Give the AI specifics about who your audience is and what sort of tone you’d like to set. For example, “Give me ideas for a best man’s speech that is funny and heartwarming but appropriate for a family audience” will generate better results than just “Write a best man’s speech”. Adding more specificity and context will further improve the output.

### **Build on previous prompts**

You don’t have to get everything into your first prompt. Try starting with a basic question and adding to it over time. Change the wording or tone or add more context and specificity to guide the AI toward the output you’re looking for.

### **Correct mistakes and give feedback**

Chat with the AI as if it’s a colleague or teammate and you’re working on a project together. Give feedback – tell it which parts of the output were useful and which parts could be improved. If you notice it got something wrong, tell it so it can correct its mistake.

### **Ask it to create your prompts or what else it needs from you**

Stuck on how to create a good prompt? Ask AI to help by creating one for you! Start with a basic idea of what you want and ask the AI to expand on it for you, like “What should I ask you to help me write a blog post about AI?”. And simply adding “Tell me what else you need to do this” at the end of any prompt can fill in any gaps that will help the AI produce better outputs.

## SAMPLE PROMPTS

### THE INSTANT EXPERT

I need to understand [Complex Topic]. Explain it to me as if I'm an intelligent beginner. Start with a one-sentence definition. Then, provide a 3-bullet-point summary of why it's important. Finally, use a simple analogy to explain the core mechanism.

### THE CREATIVE MUSE

I'm feeling blocked on my [Type of Project]. The central theme is [Your Theme]. I'm stuck on [Specific Problem]. Generate 5 unconventional ideas or starting points to get me moving again. Frame them as questions to provoke my own creativity.

## QUANTUM MECHANICS EXAMPLE

I need to understand quantum mechanics. Explain it to me as if I'm an intelligent beginner. Start with a one-sentence definition. Then, provide a 3-bullet-point summary of why it's important. Finally, use a simple analogy to explain the core mechanism.

### One-sentence definition:

Quantum mechanics is the branch of physics that describes how particles and energy behave at the smallest scales—where the rules of everyday experience break down and probabilities replace certainties.

### Why it's important:

- **Foundation of modern physics** — It underpins our understanding of atoms, molecules, and subatomic particles.
- **Technological impact** — It enables lasers, semiconductors, quantum computers, and medical imaging.
- **Deeper reality** — It reveals that nature is fundamentally probabilistic, challenging our intuitive sense of cause and effect.

### Simple analogy:

Imagine the world as a giant “cosmic casino.” In everyday life, the dice seem to land the same way every time you throw them—but at the quantum level, particles are like dice that exist in *all* possible rolls at once until you peek at them. Looking forces the “dice” to settle on a single result, but until then, reality itself is a blend of possibilities rather than a single, fixed outcome.

If you'd like, I can now walk you step-by-step from this analogy into the real math and experiments. That's where quantum mechanics really clicks.

ChatGPT's response to  
“The Instant Expert”  
on quantum  
mechanics.

Coincidentally, this response contains a hallucination! Can you spot it?

It's in the second line of the analogy. In everyday life, dice do NOT seem to land the same way every time you roll them! A better sentence would be “dice seem to land in one of a few possible states when you roll them, with each roll being un-influenced by the previous (or successor) rolls”.

# CLAUDE'S (SONNET-4) ANSWER TO "HOW TO USE"

What are the best ways for a creative writer (think speculative fiction) to best legally and ethically use you?

Here are some effective and ethical ways speculative fiction writers can leverage AI assistance:

## Brainstorming and Ideation

- Generate "what if" scenarios and explore consequences
- Develop alien cultures, future technologies, or alternate histories
- Create naming conventions for fictional worlds, species, or languages
- Explore scientific concepts that could drive your plot

## Worldbuilding Support

- Develop consistent internal logic for magic systems or technologies
- Create detailed timelines for complex narratives
- Design political systems, economies, or social structures for fictional societies
- Generate lists of details (flora/fauna, cultural practices, technological limitations)

## Character Development

- Explore character motivations and psychological profiles
- Generate dialogue samples to test different character voices
- Develop character backstories and relationships
- Create personality quirks or cultural mannerisms

## Research and Technical Details

- Explain real scientific concepts you want to incorporate or extrapolate from
- Research historical periods that inspire your alternate histories
- Understand technical processes relevant to your story
- Verify logical consistency in your speculative elements

## Editing and Revision

- Get feedback on plot holes or logical inconsistencies
- Analyze pacing and story structure
- Identify areas where worldbuilding needs more development
- Brainstorm solutions to narrative problems

## Ethical Best Practices

- Always maintain your unique voice and creative vision as the author
- Use AI as a collaborator, not a ghostwriter
- Be transparent about AI assistance if publishers or readers ask
- Don't rely on AI for entire scenes or chapters - use it for enhancement and support
- Remember that your creativity, judgment, and artistic choices are what make the work yours

The key is using AI as a sophisticated brainstorming partner and research assistant while keeping your creativity and authorship at the center of the process.

# CHATGPT'S (GPT-5) ANSWER TO "HOW TO USE"

What are the best ways for a creative writer (think speculative fiction) to best legally and ethically use you?

If you're a speculative fiction writer, the best legal and ethical way to use me is to treat me as a brainstorming partner, research assistant, and drafting tool — not as a ghostwriter whose work you pass off wholesale without adaptation.

Here's a breakdown of how to do it right:

## 1. Use me for Idea Generation, Not Direct Plagiarism

- **Good:** Ask me to generate unique plot twists, alternate history premises, alien cultures, or magic systems — and then adapt them to your own voice and world.
- **Avoid:** Asking me to "write a novel in the style of [living author]" or directly copying from copyrighted works. That risks IP infringement.

## 2. Own the Final Voice

- **Good:** Let me give you a rough scene, then rewrite it so it fits your style and tone. Your creative fingerprint should be on the final draft.
- **Avoid:** Publishing my raw text as-is — especially for major chunks — without edits, since that could lead to originality concerns and a voice mismatch.

## 3. Use Me for Research & Worldbuilding Support

- **Good:** Ask me for historical details, scientific plausibility checks, or "what if" scenario implications so your speculative worlds feel grounded.
- **Good:** Have me create maps of relationships, magic system rules, or political hierarchies for your fictional world.
- **Avoid:** Using me for factual claims in a real-world context without double-checking — especially for sensitive or legally risky content.

## 4. Keep IP Rights in Mind

- Anything I generate for you is yours to use — but if it's based heavily on someone else's copyrighted work, the original IP holder may still have rights.
- **Good:** "Invent me an entirely new alien species with its own language."
- **Avoid:** "Write a sequel chapter to Dune" — that would be derivative of Herbert's work.

## 5. Be Transparent When Needed

- In commercial or collaborative settings (especially contests, freelance gigs, or submissions), check the rules about AI use. Some require disclosure.
- **Good practice:** Be upfront if your work had AI assistance when the venue asks. Keeps you safe from ethical or legal trouble.

## 6. Treat Me as a Thinking Partner

Here are some ethically solid creative uses:

- **Brainstorm prompts** when you're stuck.
- Play out "what-if" chains to test world consistency.
- Generate lists of possible character names, alien flora/fauna, or magic item concepts.
- Roleplay a character so you can test dialogue.
- Run sensitivity passes to help you spot potentially harmful tropes.

💡 **Bottom line:** Use me as a *springboard*, not a crutch. I work best when I give you raw material, structure, or insight — and you transform it into something unmistakably yours.

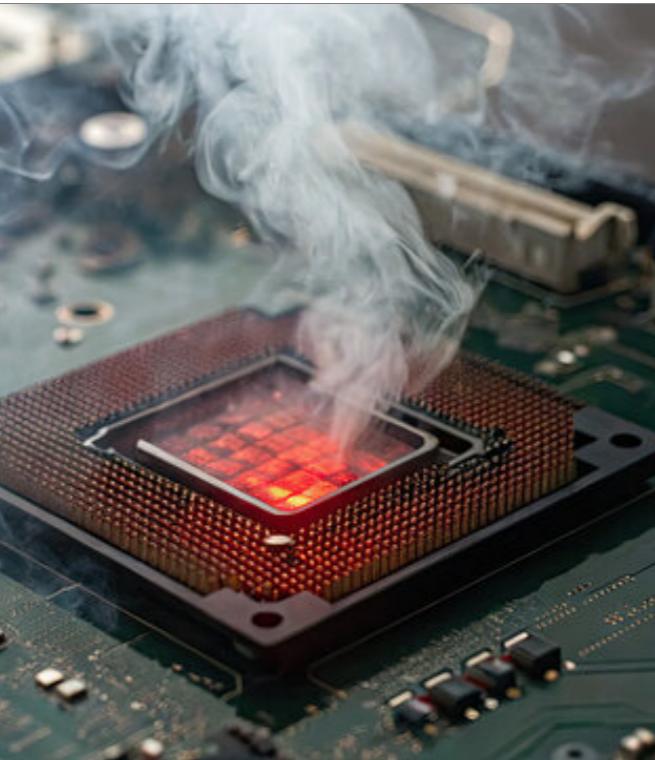
## AI CONTINUITY TRACKER

- One question I got when preparing this was “Can I load my first novel into AI and use it as a continuity tracker for the second”?
- Yes, you can. But ...
  - You probably don’t want to upload your entire work into a public AI that uses your data for training future models.
  - Use a local AI, or
  - Use a private cloud workspace (\$\$\$)
- You’ll need to periodically refresh the model with new scenes you’ve written.



## LOCAL AI

- Ollama, LM Studio, etc.
- Can be complicated to install
- Requires capable machine
- Mac is best b/c of memory architecture & low power consumption
- Performance concerns



Yes, you can run a LLM locally. There are several options - Ollama, LM Studio, and others.

Some can be complicated to install, others are straightforward. Ollama seemed pretty easy to me.

This is about to get a little technical - apologies in advance.

You'll want a capable machine, and the more GPUs you have the better. The Mac is considered the "best" for local AI, since the Apple Silicon chips have a unified memory architecture. PC's don't. (That's where the memory in the system is accessed directly by both the CPU and GPU cores, as opposed to the CPU and GPUs have separate memory.) The best Mac for local AI is the Mac Studio the the "Max" chip. It's capable of running the massive 671 million parameter DeepSeek R1 model locally. You'll need massive specs, though ... 512GB of memory and you'll have to manually tweak a few things. That machine would cost you \$9500 or more, depending on storage. By contrast, a MacBook Pro (M4 Pro) comparable to mine would set you back around \$2300. For basic tasks 16GB might do, but for anything substantial you'll want 32GB. You'll also want 512GB - 1TB of SSD storage available for your data sets (note: that's "available", meaning after you've got the OS and everything else loaded. Practically, that means a minimum of 1TB storage). (Prices as of August 2025.)

If you prefer to stick with a Windows PC, look for AI-focused PCs, and specifically look for "unified memory" in the specifications. For example, Nvidia is developing the DGX Station and DGX Spark, both of which are expected to be available later this year from "manufacturing partners" including Acer, ASUS, Dell, HP, and Lenovo. These won't be cheap though, given the specifications (128GB unified memory, 4TB SSD storage to start). Dell, HP, Lenovo, etc. all offer "AI" PC's, but they usually don't feature unified memory, putting them at a disadvantage compared to the Mac architecture.

You won't get nearly the same performance as you will with a cloud-based setup. My 8 CPU core / 16 GPU core MacBook Pro (M1 Pro chip) gets about 13 tokens/

second on Ollama phi4:14b and DeepSeekR1-14b, and 19 tokens/s on gpt-oss:20b, while OpenAI's public GPT-5 (considered a slower LLM) has been clocked at 79 tokens/s. The Mac Studio mentioned above running the huge DeepSeek model clocked in at 17-18 tokens/s running a 671b version of DeepSeek.

## WHERE TO FIND THE AIS

CHATGPT	<a href="https://chat.chatbot.app/gpt5">HTTPS://CHAT.CHATBOT.APP/GPT5</a>
CLAUDE	<a href="https://claude.ai/new">HTTPS://CLAUDE.AI/NEW</a>
DEEPEEK	<a href="https://www.deepeek.com/">HTTPS://WWW.DEEPEEK.COM/</a>
GEMINI	<a href="https://gemini.google.com/app">HTTPS://GEMINI.GOOGLE.COM/APP</a>
COPilot	<a href="https://copilot.microsoft.com/">HTTPS://COPilot.MICROSOFT.COM/</a>

## HOW TO MOVE BETWEEN AIS

### IT'S MESSY

There's no magic "import/export" feature that works between AIs. When asked how one might move from one AI to another, ChatGPT-5 suggested the following, rather manual process:

1. Gather All Material – Export AI chats, drafts, notes, images.
2. Distill into Project Dossier – Summaries, timelines, profiles, rules.
3. Create Machine-Readable Reference – JSON/text chunks for AI ingestion.
4. Choose New AI Environment – Cloud with custom knowledge base or local/private setup.
5. Load Data – Upload or paste reference docs into new AI.
6. Sanity Check – Ask test questions to verify accuracy.
7. Maintain Bible Outside AI – Keep a master copy updated for future moves.

(From ChatGPT, with some parts edited for clarity):

AI Project Portability Plan

Step 1 – Gather Everything

You want all project-relevant info in one place.

- Export AI chat logs:
- In ChatGPT: Settings → Data Controls → Export → Request Export
- Download and unzip the archive.
- Search for “<Your Project>” to pull all relevant threads.
- Gather separate documents: Drafts, outlines, character bios, maps, timelines, and any separate notes you’ve stored outside the AI.
- Collect visual references: If you’ve described or generated character/location images, grab those too.

---

Step 2 – Distill Into a “Project Dossier”

This is the human-readable master document that will survive any platform change.

Suggested structure:

1. One-Page Summary – high-level plot, major characters, world premise.

2. Timeline – main events in chronological order (including pre-story backstory).
3. Character Profiles – using a template like the one we built for your series bible.
4. Location Profiles – setting descriptions, maps, and key history.
5. Technology/Magic Systems – mechanics, rules, limitations.
6. Themes & Tone Notes – to guide future scenes.
7. Continuity Notes – details that must remain consistent.

Format: Markdown, Word, or Google Docs (Markdown is the most portable).

---

### Step 3 – Create a Machine-Readable Reference

- Take your Project Dossier and break it into small, labeled chunks (e.g., one character per file, one location per file).
- Save them in plain text or JSON with clear tags:

```
{  
  "type": "character",  
  "name": "Lady Demerzel",  
  "description": "18,000 year-old human-appearing robot ...",  
  "continuity_notes": "Hair color: blonde, Eyes: blue..."  
}
```

- These can be fed into a new AI's custom knowledge base or vector database.

---

### Step 4 – Choose Your New AI Environment

Depending on needs:

- Cloud-based with custom knowledge → ChatGPT Enterprise, Claude Projects, Gemini Workspace.
- Local/private → Ollama + ChromaDB/Weaviate for searchable embeddings.

---

### Step 5 – Load Your Data

- If the new AI supports document uploads: upload the entire Project Dossier + JSON files.
- If it doesn't: paste summaries in at the start of each new conversation or use a prompt like:

"You are assisting me with the <Project Title> project. Here is the reference material you should keep in mind for this session: [paste summary]."

---

### Step 6 – Sanity Check

- Ask the new AI sample questions to test recall and accuracy:
- “What is Lady Demerzel’s relationship to Brother Dusk (Clean XVII)?”
- “Describe Kalgan and why it matters to Gaal Dornick.”
- Compare answers to your original notes.

---

#### Step 7 – Maintain the Bible Across Platforms

- Keep the Project Dossier updated outside any one AI.
- Every time you add a chapter, re-run your extraction prompt and append results.
- This way, if you ever switch again, you’re moving a clean, complete package — not hoping your old AI “remembers” correctly.

The portability plan is a perfect way to talk about AI as a tool, not a vault — it works best when you hold the “single source of truth,” and AIs are assistants that you can swap out without losing the heart of your project.

## AUTHOR-FOCUSED AI TOOLS

Here's a side-by-side comparison of the four tools you mentioned, geared toward what your writers' group would care about most.

(Yes, I asked ChatGPT to help with this slide.)

Tool	Core Strengths	Weaknesses / Watchouts	Best Fit For	Cost
Sudowrite <a href="http://www.sudowrite.com">www.sudowrite.com</a>	Brainstorm plot ideas, scene expansion, sensory detail, style mimicry	Verbose/flowery prose if unchecked, limited continuity memory	Creativity boosts mid-draft, overcoming writer's block	\$10–\$30/mo
Raptor Write <a href="http://www.raptorwrite.com">www.raptorwrite.com</a>	Outlining + drafting in one, plot/character info visible while writing	Younger product, some features less polished, learning curve	All-in-one drafting + planning environment	\$20–\$40/mo
Plotdrive <a href="http://www.plotdrive.com">www.plotdrive.com</a>	Beat sheet templates, plot problem-solving, untangle story middles	Not strong for scene prose, best pre-draft or outline stage	Plotters and structure-minded writers	\$15–\$25/mo
Your First Draft <a href="http://yourfirstdraft.ai">yourfirstdraft.ai</a>	Rapid drafting to full manuscript, momentum over polish	Generic first drafts, heavy revision needed, weak for deep arcs	Writers needing a complete draft fast	\$10–\$20/mo

(From ChatGPT-5):

If your writers' group is thinking about Sudowrite, Raptor Write, Plotdrive, Your First Draft, and other author-focused AI tools, I'd break it down into five areas they should understand before committing time (and money).

---

### 1. Purpose & Strengths

These tools are all built around narrower, author-specific workflows than a general AI like ChatGPT. Examples:

- Sudowrite – Idea expansion, sensory detail, rewriting, brainstorming “what if” scenarios.
- Raptor Write – Long-form drafting with outlining, character sheets, and scene management baked in.
- Plotdrive – Story structure guidance, beat sheet generation, and plot problem-solving.
- Your First Draft – Speed drafting aimed at producing a rough, complete manuscript quickly.

Key takeaway: Each has a slightly different “center of gravity.” Some excel at ideation, others at pushing you through a draft, others at structural planning.

---

### 2. Integration with Your Process

Questions to ask:

- Does the tool fit where you get stuck (plotting vs. drafting vs. editing)?
- Can you export your work easily to your main writing environment (Scrivener, Word, Google Docs)?
- Does it allow you to work modularly — e.g., use it just for character bios without committing your whole draft?

Some writers treat these tools as “bolt-ons” for specific stages rather than all-in writing platforms.

---

### 3. Privacy & Rights

Critical for authors:

- Data ownership – Do you retain full rights to the text you input and what the AI generates?
- Training policy – Will your manuscript be used to train their model? (Some opt you in by default.)
- Export control – Can you delete your data from their servers?
- Terms for publication – No major publisher has yet rejected a book for using AI assistance, but they may ask whether AI content is original and rights-clear.

If privacy is a concern, look for tools that promise no training on your work, or use local/offline AI.

---

### 4. Model Limitations

Even author-specific AIs:

- Can hallucinate facts or contradict earlier text.
- May generate prose that feels generic if over-relied on.
- Struggle with long-range story consistency unless they have built-in memory features or allow you to upload a project bible.

Best practice: Treat them as collaborators, not ghostwriters — you still review, adapt, and rewrite.

---

### 5. Cost vs. Benefit

- Most of these tools run on subscription models (\$10–\$40/month, some much higher).
- Consider whether you'll use them often enough to justify the expense — and whether a general-purpose AI with a good prompt library could give you similar value.
- A few offer free tiers or trials — good for test-driving before committing.

### Final Advice for Your Group

- Match the tool to the problem: Don't pick Sudowrite because it's popular; pick it if its scene-expansion features solve your pain point.

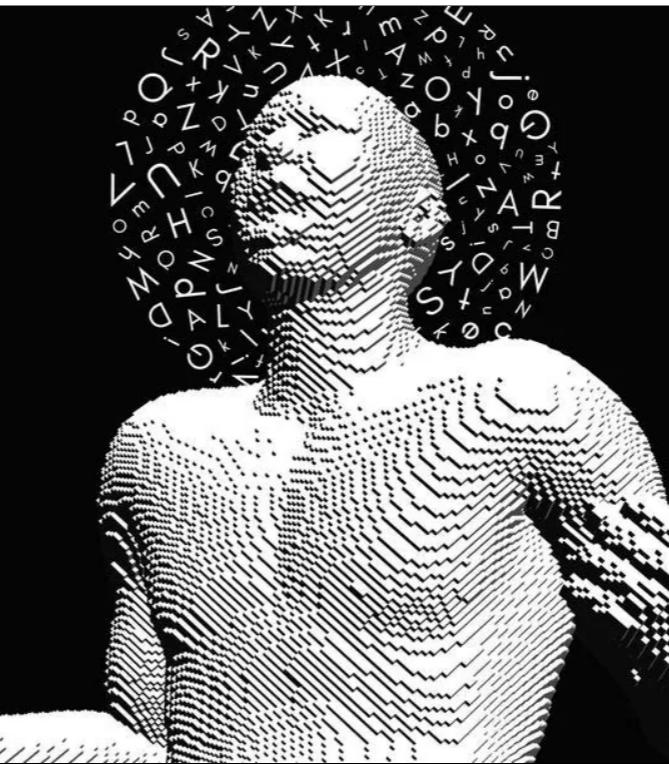
- Start with one, not five: Avoid spreading your creative process across too many platforms at once — fragmentation can kill momentum.
- Own your core documents: Always keep a local copy of your draft, notes, and bible so you can switch tools without losing your work.
- Keep your voice: The more you rewrite AI-generated text, the more it will sound like you and not the machine.

INTELLIGENCE?



## YOU CALL THIS INTELLIGENCE?

- It should be clear by now that AI is certainly artificial, but with today's state of the art, not intelligent (at least in the way humans are).
- "The singularity" is the name given to the point at which AI equals human-level intelligence.
- Futurists believe it is on the horizon. An analysis of 8600 predictions gives a 50% chance of it happening between 2041 and 2060.
- Some predict it by 2030, others say never.





In case you're wondering, that's actress Laura Birn as "Lady Demerzel", a 18,000-year old robot in Apple TV's "Foundation" series. It's worth a watch, if you like that kind of thing.