Trudeau.ai

AI LITERACY FOR LAW STUDENTS

Using AI while Building Legal Judgment

Developed by Professor Chris Trudeau

Version 1, Jan 2026

Disclaimer: This guide is not endorsed by, sponsored by, or affiliated with Detroit Mercy Law School. It was written and developed solely by Professor Trudeau in his individual capacity.

If you are enrolled in one of Professor Trudeau's courses, you should follow this guide. It reflects how he expects students in his classes to think about and use AI responsibly as part of their learning.

If you are reading this guide but are not a student in one of his courses, treat it as one professor's approach, not universal law school policy. AI rules, expectations, and honor code requirements vary by school and by instructor, so you should always confirm what is permitted and expected in your own classes BEFORE using AI.

1. What This Guide Is Not

Before we talk about what AI can do for you as a law student, I want to clear the air. A lot of confusion about AI comes from people expecting it to be something it is not. This guide is not a user manual for any particular tool. I am not going to teach you which buttons to click, which features are hot this week, or how to write prompts that look like they were drafted by a prompt engineer who has not slept since 2021. Those things change constantly, and the moment you learn them, the interface changes, and you are back where you started.

This guide is also not a shortcut guide. It is not about avoiding reading, skipping the hard parts, or turning law school into a scavenger hunt for answers. If you are hoping for that, you will be disappointed, and that is intentional.

It is not a replacement for thinking, and it is not a way to outsource responsibility. Anything you submit with your name on it is still yours, even if a tool helped generate it. In law school, you are learning to think and act like competent, ethical lawyers. Remember, in practice, you will be the one that clients turn to (and pay) when they need help. You will be responsible when things go wrong. AI does not change those realities.

Finally, this guide is not tied to one course. I wrote it for law students, but the habits here should travel with you across classes and into your career. The tools will change, but your professional judgment will always be there. That's the number one rule for using AI in law practice – never outsource your professional judgment to AI.

2. What This Guide Can Do for Law Students

AI can be a helpful study partner, but if you use it poorly, it can quietly interfere with learning in ways you often do not notice until sitting in the exam or after receiving feedback on that writing assignment. The point of this guide is to keep you on the helpful side of that line.

I want you to use AI in ways that support learning rather than replace it. Many students love AI because it gives explanations that feel clear and reassuring. That feeling can be useful, but it can also be misleading. One of the central risks of AI-assisted studying is confusing recognition with understanding. You recognize the rule. You recognize the structure. You recognize the vocabulary. But then, on the exam (or in real life), the fact scenario changes and suddenly your recognition is not enough. Recognition doesn't help you apply difficult concepts to varied factual scenarios. That's what mastery does.

I also want you to avoid false confidence. Modern AI systems are fluent by design. They explain things smoothly, organize ideas neatly, and rarely hesitate (or tell you they do not know something) unless you tell them to. Do not confuse fluency with accuracy. A clean, clear explanation can still be wrong, incomplete, or subtly off in ways that matter in legal analysis.

There is also a quieter risk that matters for students. When a tool makes the hard parts feel easy, it can tempt you into what some researchers call **cognitive outsourcing**. The tool does the work, you get the product, and your brain never has to build the muscle. That is what we **do not** want.

Study on Using AI to Write Essays

A research team associated with MIT's Media Lab ran an essay writing study that gets at this in a concrete way. Participants wrote under different conditions, including writing with ChatGPT, writing with a search engine, and writing with no tools. The researchers used EEG measures during the writing task and also looked at the essays themselves. Their headline finding is simple: the group writing with the AI assistant showed the weakest brain connectivity, reported the lowest sense of ownership over what they wrote, and struggled more when asked to recall or quote their own work. In other words, the writing may look fine on the screen, but the learning can be less than it would be if you wrote the essay the old-fashioned way.

If you want to see the study and the authors' materials directly, start with the paper here: https://arxiv.org/abs/2506.08872 The authors also maintain a short FAQ for readers here: https://www.brainonllm.com/faq

Most importantly, this guide is about building good professional judgment. Law school is not just about acquiring information. It is about learning how to evaluate claims,

testing your reasoning, and deciding when something is reliable enough to use. AI does not change that goal. It raises the stakes because it makes it easier to accept an answer without doing the work that makes you better.

If you take away only one idea from this guide, take this one – AI can aid your thinking, but it cannot do the thinking for you. The moment you stop checking, questioning, and translating an AI answer into your own understanding is the moment it stops helping you learn.

---

3. How to Read and Use This Guide

This guide is not meant to be read once and forgotten. It is a reference you will understand better after you have used AI for a while and have seen the places where it shines and the places where it quietly stumbles.

You do not need to memorize every technical term right away. The goal is not technical fluency. The goal is conceptual clarity. As you read, focus on how these systems behave, where they tend to fail, and how your own reactions to their output can shape your judgment. Here is the question I want you to keep in the background as you use any AI tool:

"If this output were wrong, how would I know?"

That single question turns AI from an answer-providing machine into a learning tool.

---

4. What AI Is Actually Doing When It "Helps" You

It is easy to think of AI as a smart person sitting next to you, patiently explaining the law. That is how it feels when the output is clear, confident, and well organized. That feeling is the first thing you should learn to treat carefully.

Most chatbots you will use are powered by a large language model (LLM). A large language model is trained on massive amounts of text and learns patterns in language. When you ask a question, it generates a response by predicting what text is likely to come next, given what you just typed. Put simply, most LLMs are just really sophisticated and well-trained text prediction tools.

That means the LLM is not starting with facts and reasoning toward a conclusion in the way you do. It is starting with patterns and generating an answer that fits those patterns. That is why it can be so helpful for certain things, like outlining, explaining, and reorganizing ideas. It is also why it can be wrong while sounding completely sure.

There are two technical ideas that you should know about because they affect your day-to-day experience with AI:

1. **Tokens.** Models do not process words the way humans do. They process tokens, which are smaller chunks of text. Tokens are one reason long prompts hit limits faster than you expect. OpenAI has a more detailed (yet still clear) explanation of tokens and how to count them for English language text here: https://help.openai.com/en/articles/4936856-what-are-tokens-and-how-to-count-them
2. **The context window.** This is the amount of text the model can consider at once. If a key detail is not in that window, the model cannot use it. If it is in the window but buried, the model might still miss it. When people say, AI ignored what I told it, a context problem is often the reason.

If you are curious about the technical architecture behind LLMs, many are built using a transformer design. You do not need the math, but it helps to know that the model uses attention to decide what parts of your input matter as it generates a response. If you want to read the foundational paper that really launched LLMs as we know them, you can find it here: https://arxiv.org/abs/1706.03762

---

5. Why AI Gets Things Wrong (And Why That's Normal)

When a model hallucinates, it produces something that sounds plausible but is not supported by reliable evidence. In law school, the most common hallucination pattern is not the wild fantasy answer. It is the answer that is mostly right and slightly wrong. That slight wrongness is where the trouble lives.

Hallucinations show up in predictable ways. The model invents a citation. It states a rule too broadly. It blends two similar doctrines. It answers a question that feels close to yours instead of the one you actually asked. It gives you a confident answer when the honest answer is, it depends.

Why does this happen? Because the model is doing what it was designed to do. It is generating likely text. It is not checking a database of truth. It is not pulling your casebook off a shelf. Unless the system is connected to sources, it is not looking anything up. If you want the clearest explanation of this from the people building the systems, start here: https://openai.com/index/why-language-models-hallucinate/

If you are thinking, this is interesting, but it still feels hypothetical, here are more cautionary tales than you'll ever need in the legal world. There is a running list of real cases where lawyers and litigants filed AI-generated citations, quotations, or other hallucinated material and got sanctioned, embarrassed, or worse. It is maintained by Professor Damien Charlotin, and it is a useful reminder that courts do not care whether the mistake came from a human associate or a chatbot. They care that it is in the filing. The constantly updated database is here:

https://www.damiencharlotin.com/hallucinations/?q=&sort_by=-date&states=USA&period_idx=0

Here is the part that should make this feel real. Imagine you are a young lawyer with a motion due tomorrow. You are tired. You use a chatbot to draft a section of the brief, and it gives you a case name that sounds perfect. It even gives you a quote. You paste it in, maybe you tweak a sentence or two, and you move on because it looks clean, and you have ten other things to do.

Then the judge, or the other side, tries to find the case. It does not exist. Now the issue is no longer whether you used AI. The issue is whether you filed something you did not verify. That is when you start seeing orders to show cause, sanctions, and opinions that live online forever. The painful irony is that the original mistake often takes five minutes to avoid. Pull the case. Check the quote. If you cannot find it, do not file it.

If you want an overview of why this keeps happening and why judges are increasingly irritated, the Washington Post has a useful piece that discusses the problem and references Charlotin's database:

https://www.washingtonpost.com/nation/2025/06/03/attorneys-court-ai-hallucinations-judges/

Practical Takeaway: When you see a clean answer, don't JUST ask, does this sound right? Ask, what would happen if this were wrong? That question keeps you in control.

---

6. Key AI Terms Law Students Actually Need

There are a hundred AI terms floating around. You do not need to learn all of them. But you do need to understand the ones that, if misunderstood, will cause you to use AI badly.

A large language model is a system trained on large amounts of text that generates language by predicting what comes next. If you remember nothing else, remember this. It generates plausible text. It does not guarantee true text.

Tokens are the chunks of text the model processes. Tokens help explain why long prompts get truncated, why the tool sometimes forgets something you pasted earlier, and why giving it a full casebook chapter can backfire.

The context window is the amount of text the model can consider at once. If your key facts are not in the window, the model cannot use them. If they are in the window but mixed with clutter, it may still miss them.

Tool calling, sometimes called **function calling**, is when a model can call an outside tool like a database, a calculator, or a retrieval system. This often increases reliability of the LLM because the model can fetch information instead of guessing. OpenAI's overview of function calling is here: https://platform.openai.com/docs/guides/function-calling

Retrieval augmented generation, often called **RAG**, is a design where the system retrieves relevant documents and then generates an answer based on those materials. This is what I use to create AI tutors for my specific doctrinal classes. I upload all my teaching material, so the tutor uses RAG to pull from my own materials, which helps the tutor answer the questions the way I want it to. Although RAG can reduce hallucinations and improve traceability, but it is only as good as the documents and the AI retrieval process. If you want a deep dive on what RAG is and how it functions, here is the foundational paper: https://arxiv.org/abs/2005.11401

Grounding means forcing the AI model to work from specific materials instead of its baseline training by the company that created the LLM. If you paste your notes, a rule statement from your casebook, a statute section, or your professor's framing, you are grounding the answer in the right universe. This is one of the simplest ways to reduce hallucinations without learning anything technical.

- **Tip:** One good way to do this in ChatGPT is to create a project folder for each new class topic or assignment. Then upload the relevant information from your class, and it will ground the AI for that specific project folder.

Prompt engineering is just the practice of writing prompts in a way that reliably produces the kind of output you want. For law students, the goal is usually not prettier answers. It is better practice, clearer explanations, and fewer confident mistakes.

Prompt injection is a key security concept that all lawyers should know because it reveals a deeper truth about how these systems work. Models do not naturally separate instructions from content. That means untrusted text can sometimes manipulate the model into doing the wrong thing, especially when tools are involved. Here is a great cheat sheet on preventing prompt injection that is worth the read: https://cheatsheetseries.owasp.org/cheatsheets/LLM_Prompt_Injection_Prevention_Cheat_Sheet.html

Few-shot prompting (aka few-shot learning) means you include one or more short examples in your prompt that show what you want. Think of it like showing the model a mini answer key for the format, not for the substance. It often improves consistency when you want the output to look a certain way.

7. Prompting: Law-Student Style

If you've been on social media at all lately, you've likely seen people advertising courses about prompting. And prompting is important. But the problem with those courses is that most advice about prompting is written as if your job is to produce output. In law school, your job is usually to produce understanding. Those are not the same goals.

Here is the basic idea. A prompt is not a magic phrase. It is simply the set of instructions and information you give the model. If you want better results, you give clearer inputs. That is not an AI trick. That is clear communication, which is a vital tool for all legal communication. In fact, I've spent my entire professional life advocating for clear legal communication. Now, with AI prompting, those clear communication skills are even more important. For this guide though, I'll focus on how to best prompt to aid your learning.

When you prompt for learning, you want a few things:

1. First, you want a role that supports learning, not a role that replaces your work. Ask the system to be a tutor, a quizzer, a Socratic partner, or a skeptical reviewer. At the time I am writing this guide, ChatGPT 5.2 has two modes that are useful for learning – "Study and Learn" mode and "Quizzes" mode. Play around with both of those modes so that the LLM knows to help you learn, not just give you the answers.
2. Second, you want context. Tell it what class you are in, what topic you are working on, and what you already understand. If you have a definition from your notes or your casebook, paste it. Models do better when you give them the exact material you want them to use.
3. Third, you want constraints. Tell it what you do not want. For example, do not give me the final answer first. Ask me questions. Force me to explain it. Do not cite cases unless you are confident they are real. If you are unsure, say so.

The best prompts for creating material have one more feature. They include a verification hook. Ask the model to cite where it is getting a rule from, or to label what is a direct quote versus what is paraphrase. Then you check it to verify. I always assume that any rule or quote that the AI writes is off in some slight way. It may not be, but if you assume it is wrong in some way, that forces you to verify it.

Deeper Dives to Learn about Prompting Techniques

- This post by Thompson Reuters explains prompt engineering techniques for legal work, so it is worth a read: https://legal.thomsonreuters.com/blog/prompt-engineering-best-ai-output/
- If you want a deeper dive, here's a guide from Microsoft that discusses few-shot prompting, and many more things about prompting: https://learn.microsoft.com/en-us/azure/ai-foundry/openai/concepts/prompt-engineering?view=foundry-classic
- If you want to go deeper on prompt construction, OpenAI's prompt engineering guide is helpful: https://platform.openai.com/docs/guides/prompt-engineering
- OpenAI also has a practical best practices article with prompt formats that work well: https://help.openai.com/en/articles/6654000-best-practices-for-prompt-engineering-with-the-openai-api

---

8. AI Assistance v. AI Substitution

This is where students get into trouble, often without realizing it. This guide explains how AI can assist you in ways that strengthen learning. It can help you practice, test your understanding, generate examples, and pressure test your reasoning. However, AI can also substitute for your work in a way that feels productive but actually prevents learning.

Here is a simple way to tell the difference. AI assistance increases your ability to explain the material in your own words. But substitution makes you dependent on the tool to produce the explanation for you.

If AI helped you understand a case, you should still be able to brief it on your own without AI's help. If AI helped you understand a rule, you should still be able to apply it to a new fact pattern without AI. If AI helped you draft something, you should still be able to defend every sentence as accurate and appropriate.

Remember this: The biggest trap with AI is the comfort trap. AI outputs seem straightforward and read well. But remember, law school is messy. Do not confuse clear answers with correct answers.

9. The Human Side of AI Use

If AI errors were obvious, using AI would be easy. The challenge is that AI errors often come wrapped in confident prose, and humans are not built to resist confident prose. There are two concepts you should understand to fully appreciate the impact that AI has on human judgment.

First, is the ELIZA effect. Have you ever chatted with ChatGPT and thought it understood you like a human would? If so, you've experienced the ELIZA effect, which simply means that people quickly project human-like understanding onto conversational systems. This is not because people are careless. It is because language is one of our strongest signals of intelligence. The original ELIZA paper by Joseph Weizenbaum is still a classic, and it was written in 1966:
https://cse.buffalo.edu/~rapaport/572/S02/weizenbaum.eliza.1966.pdf

> If you want more and are curious about human-AI interaction like I am, here is a modern explanation of the ELIZA effect for today's AI systems:
> https://www.nngroup.com/articles/eliza-effect-ai/

Second, is automation bias. Like the ELIZA effect, this bias has been around for decades. What automation bias means is that when a system provides answers, humans tend to rely on those answers even when the system is imperfect. Think about calculators. If you put in 1234 x 1234, you would blindly trust what the calculator tells you the answer is versus doing the math yourself, right? After all, it's a calculator. It was programmed to do this exact thing, so it has to know more than you about math, right? But what if the calculator was wrong because it had faulty programming or some other glitch? You wouldn't know it. (By the way, the answer is 1,522,756. I used my calculator and did the math on paper.)

Now think about automation bias in the context of writing legal briefs. You put in a great prompt and, perhaps, even told the AI what cases to use. You prompted it well, so it produces something that looks and reads like a brief. After all, when you put in good prompts, LLMs are trained to respond to the prompt. But how do you know that the LLM took the exact language in the quote or understood the nuanced way that courts interpret some rules? This is where you have to keep your human judgment and not succumb to automation bias. The real risk of using AI for legal tasks is that it is so easy to just take what is there and not verify things. That is what many lawyers are struggling with now, as you've seen if you looked at the Charlotin Database discussed earlier.

10. Verification as a Learning Habit

Students often hear verify and translate it as protect yourself. That is a key part of it, but it is not the only important part. In fact, the act of verifying things is one of the fastest ways to learn because it forces you to do the work law school is trying to teach you to do. You separate strong claims from weak claims. You track down authority. You test whether a rule actually says what you think it says. So verify to protect yourself and to help yourself master the material.

A few habits pay off when verifying AI output for writing assignments. Ask AI for the sources it used and check them. If the source shows up and is titled the way the AI says, that is a good sign. Also check that the authors of the source and the dates match up. AI can misattribute things to the wrong people, and you don't want that either. Always quote check anything in quotation marks. I always assume that any quotation that AI puts in there is not accurate in some way.

---

11. A Simple Judgment Workflow for Students

A judgment workflow is a repeatable way to make decisions when something feels helpful but might also be risky. Think of it like a guardrail you put in place before you are tired, rushed, or overly impressed by a clean answer. AI is exactly the kind of tool that needs that guardrail. It produces fluent output fast. That speed tempts you to accept it, use it, and move on. A judgment workflow slows you down on purpose and keeps you in control.

This is the workflow I want you to internalize for using AI. It is intentionally simple.

- **Generate**. Depending on why you are using AI, use it to produce an output. That could be a draft explanation, a set of practice questions, a list of issues, or a summary of what you think you know.
- **Pause**. Do not treat the output as the answer. Treat it as a proposal. Ask yourself what parts you truly understand and what parts you are tempted to accept because they sound good.
- **Verify.** Check the parts that matter, meaning the parts that would cause real damage if they were wrong. Verification might mean checking a case, checking a statute, checking a definition, or checking whether an example actually matches the rule. I'd also check out any part of AI's answer that seems new to you or sounds off in some way. That's your brain throwing up a red flag. AI might not be wrong, but

if you check it and it is right, then the act of verifying helped you learn something. And if AI is wrong, you not only save yourself, but you reinforced your learning by noticing that someone was off and then verifying what the right answer was. A win-win.

- **Decide**. Only after you verify should you decide how to use the output. Sometimes you use it or parts of it. Sometimes you'll re-prompt Ai to have it change things. You might do this if you were asking AI to create practice questions. But if you are using AI to write something, maybe you'd just delete it and start over. Starting over is not failure; it is good professional judgment.

---

12. When AI Is a Good Study Partner and When It's a Bad Idea

AI is most helpful when it helps you practice. If you have already read and prepared, AI can generate practice hypos, it can quiz you on definitions, and it can force you to explain a concept until you can do it cleanly. AI is also helpful for reorganizing material you already understand. Turning messy notes into a cleaner outline can be useful, as long as the substance is yours.

AI is risky when you use it to replace the hard part. If you use AI to summarize a case you did not read, you are skipping the skill that we are trying to train you to do. If you use AI to give you the right answer to an issue spotter, you are skipping the judgment that matters most. If you use AI to write your analysis, you are outsourcing the part of legal writing that makes you a lawyer.

There is one other vital rule for using AI in law practice or for clinics. **Do not paste confidential information into AI tools you do not control.** That includes client facts in a clinic or externship. It also includes personal information you would not want repeated, forwarded, or pulled into the open.

Here is a cautionary tale. In 2025, Sam Altman (CEO of AI) publicly warned that conversations with ChatGPT do not have lawyer-client, doctor-patient, or any other type of professional confidentiality. He actually warned that chat histories could be accessible in litigation. This means that there is no privilege protecting your chat with a chatbot the way privilege can protect a conversation with a lawyer or a licensed clinician.

If you are in a dispute, the other side may try to get that information, and a court can order records preserved or produced depending on the circumstances. If you want a readable summary of Altman's warning, read this article:
https://mashable.com/article/sam-altman-theo-von-apperance-chatgpt-therapy-still-bad-idea

If that feels abstract, here is the practical student takeaway. Assume your chat is not private in the legal sense. Assume it can be logged. Assume it can be retained. Assume it can be demanded through legal process. That does not mean every chat will end up in court. But it does mean that you should not treat the tool like a confidential diary.

> **For more on this:** OpenAI itself publishes information about government requests for user data and how it evaluates them, which is a useful reminder that these records can be requested: https://openai.com/trust-and-transparency/
>
> OpenAI also explains that some features, like temporary chats, still involve retention for a period of time for safety purposes. That is not a critique. It is a reality you should understand before you share sensitive information: https://openai.com/policies/row-privacy-policy/

---

13. Responsibility, Attribution, and Professional Identity

This is the part where everything comes together, because it is the part that follows you out of law school. In practice, clients rely on your work. Judges rely on your representations. Other lawyers rely on what you say. AI can help you work faster, but it cannot carry responsibility for what you file, what you advise, or what you claim is supported by authority. This section is not about fear. It is about your professional identity.

If your name is on it, you own it. This also includes using your student number or exam id numbers, of course. That rule holds true for a memo, a brief, a class assignment, and eventually client work. AI does not change responsibility. It changes workflow.

When it comes to attributing how you used AI, I want you to think less about "confessing" and more about "being able to explain yourself." If a professor, a supervisor, or someday a judge asks, you should be able to say what you used AI for and how you verified what mattered. The ABA's ethics guidance says lawyers have to consider when they need to disclose AI use to clients, and it emphasizes that AI output has to be independently reviewed and verified before it shows up in anything that matters. In other words, "I used a tool" is never the end of the story. The end of the story is always, "here is what I did to make sure it was right." Read more about the ABA's guidance here: https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf

There is also a very practical reason to build this habit now. So far, when lawyers have gotten in trouble for AI hallucinations, the cases that end better are the ones where the lawyer can show a responsible process and credible remediation. In November 2025, a judge declined to impose formal sanctions after false AI-generated citations appeared in a filing, in part, because the lawyers laid out remedial actions and addressed what went

wrong. That is why being able to explain your workflow is important. Read more about that case here:
https://www.reuters.com/legal/government/law-firm-escapes-sanctions-over-ai-generated-case-citations-2025-11-13/

For law school (or law practice), my advice is simple. Do not treat attribution like a footnote ritual. Treat it like keeping a small receipt. If you used AI in a way that mattered, be ready to describe it and be ready to point to what you checked. That is how you stay credible when someone asks hard questions later.

---

14. Closing: AI Changes Workflows, Not Responsibility

You are going to be using AI for the rest of your life. That is not a prediction. AI is the new normal.

The danger is not that AI will ruin you. The danger is that it will make you comfortable. It will make the work feel easier than it is. It will hand you answers that sound finished. If you build the habit of accepting that output in law school, you will not become a very good lawyer. Not because you lack intelligence, but because you never fully build the muscles that matter. Those muscles are knowledge and judgment. They show up when you decide what is reliable, what is missing, what needs to be checked, and what you cannot honestly say with confidence.

Professional judgment is the discipline of doing that anyway, even when the screen is offering you a neat paragraph. It is the choice to pause, verify, compare, read the actual authority, and then own the conclusion. AI can draft and suggest. It cannot take responsibility for the decision. It cannot explain itself when the judge asks, when the client asks, or when your name is on the page.

So build a different habit. Use AI to practice, not to perform. Use it to test what you know, not to supply what you do not. Use it to expose gaps in your thinking, then fill those gaps with real work. Verify what matters. Make the decision yourself. The goal is not to avoid AI. The goal is to stay responsible for your actions.

AI-Use Note: AI helped me create this guide. I used ChatGPT to brainstorm structure, tighten phrasing, and generate examples and some initial text. AI helped the most in the parts with definitions of terms and concepts. But even though AI generated some content, I edited every section, checked sources, and made the final calls. Any mistakes are mine.