

Appendix D: Platform Tips and Workflow Strategies

A few hundred conversations in, you realize that how you use GenAI platforms matters almost as much as what you ask them to do. The tips below fall into three categories: platform mechanics (how the platforms work), prompt-craft habits (how to talk to the AI in ways that produce better results), and workflow strategies (how to build GenAI into your teaching without losing track of what you built). None of these require a technical background — just some intentionality about process.

PART I: PLATFORM MECHANICS

TIP 1: Stop and Edit When You Hit ENTER Too Soon

Accidentally hit ENTER before finishing your prompt? Don't let the LLM generate a response to incomplete instructions. Click the STOP button, then click EDIT on your prompt. Finish what you meant to say, then resubmit.

Where to find STOP in ChatGPT, Claude, and Gemini:

Click the black square (blue in Gemini) that appears while the LLM is generating
Hover over your message and click the pencil icon to edit
Complete your prompt and resubmit

Stopping early saves tokens, avoids wasted output, and keeps you from having to open your next message with “ignore that, here's what I meant.”

TIP 2: Attach Documents to Individual Prompts (Without Adding Them to Your Knowledge Base)

When building CustomGPTs, you can upload files to the Knowledge section, making them permanently available for all conversations. See Chapter 10 for instructions. But sometimes you want a document available for just one prompt without cluttering the knowledge base.

How to do it: In any prompt, click the plus sign or paperclip icon and upload your file. The AI can reference it for that conversation thread, but it won't become part of the CustomGPT's permanent knowledge.

Example: You've built a CustomGPT for your Contracts course with the Restatement in "knowledge." A student asks you to review their draft answer. Upload the draft as a prompt attachment, not to "knowledge," so it doesn't pollute the permanent materials.

When to use each:

Use Knowledge (permanent)	Use Prompt Attachment (temporary)
Course syllabus	Student work samples you're grading
Key cases (Pennoyer, Shoe, etc.)	Draft assignments you're testing
Restatements, model codes	One-off research articles
Assignment templates	Previous semester's exam for comparison

If you upload every student paper to your CustomGPT's knowledge base, it becomes polluted with documents that were never meant to be there. Keep knowledge clean with permanent reference materials and attach individual documents when needed.

TIP 3: Rename Conversations to Find Them Later

By default, GenAI platforms auto-generate conversation titles based on your first prompt. These are often generic ("New chat," "Contracts help") and useless when you're searching for a specific conversation three months later.

How to rename in Claude:

- Click the dropdown next to the conversation title at the top of the prompt window
- Choose Rename
- Type a descriptive name
- Press ENTER to save

How to rename in ChatGPT:

- Click the conversation title in the left sidebar under "Your chats"
- Click the three dots (:)
- Choose Rename
- Type a new name
- Click outside the text box or press ENTER to save

How to rename in Gemini:

- Click the three dots (:) at the top right of the chat window
- Select Rename
- Type the new name

- Click Rename

Best practices for naming:

- Include the date: “Contracts hypotheticals – Sept 2025”
- Include subject and task: “Civ Pro – Jurisdiction outline”
- Include the deliverable: “Final exam draft – Torts Fall 2025”
- Be specific enough to find it later: “Model answer – Pennoyer fact pattern” beats “Civ Pro” by a mile
- Add a prefix for special categories. When creating prompts to demo for other faculty, I use “FACULTY DEMO:” followed by a description.

Trust me on this one. You’ll generate dozens or hundreds of conversations, and when you need that hypothetical from last semester, a descriptive title is the difference between finding it in ten seconds and giving up and rebuilding it from scratch.

TIP 4: Highlight Text to Focus AI’s Attention

In Claude and ChatGPT, you don’t need to retype or copy-paste portions of previous messages when you want to discuss them. Highlight any text from earlier in the conversation, then ask your question. The platform focuses its response on the highlighted selection.

How to do it in Claude:

- Click and highlight any text in a previous chat (yours or Claude’s)
- Click “Reply”
- Type your question in the prompt box
- Claude focuses its response on the highlighted portion

How to do it in ChatGPT:

- Click and highlight any text from earlier in the conversation
- Click “Ask ChatGPT”
- Type your question
- ChatGPT references the highlighted selection in its response

Why this matters for legal education: When working with long case analyses, multi-part hypotheticals, or detailed legal explanations, you often want to zoom in on one specific aspect. Instead of writing “In your previous response where you discussed the third element of negligence, can you explain the zone of danger test in more detail?”, highlight the relevant sentence and ask “Explain the zone of danger test.”

Example workflow:

- AI generates a 500-word analysis of a Civil Procedure hypothetical
- You highlight just the sentence about subject matter jurisdiction
- You ask: “Is this still accurate after Grable?”

- AI focuses only on that jurisdictional point

Highlighting reduces ambiguity and prevents the AI from re-explaining material you already understand. It's particularly useful when iterating on assignments or when asking students to critique specific portions of AI-generated analysis.

TIP 5: Share a Conversation

Claude, ChatGPT, and Gemini all let you generate a shareable link to a conversation or a custom tool.

How to share in Claude:

Click the share icon (box with arrow) at the top right of the conversation
Choose "Share link"
Copy the link

How to share in ChatGPT:

Click the share icon at the top right of the conversation
Choose "Copy link"
You can control whether your name appears

How to share in Gemini:

To share a chat: Click the share icon at the top right of the chat window

To share a Gem: Click the three dots (:) next to the Gem and select Share

For faculty:

If you've built a good Socratic exchange demonstrating how to analyze a Fourth Amendment stop, share the link with students before the class discussion. They arrive having already seen the analytical process, and your class time goes further.

For students: A shareable link is the cleanest way to submit a GenAI Appendix. Instead of copy-pasting a 30-message conversation into a Word document, the student pastes a link. You see exactly what happened, in order, in the platform's own formatting.

One caution: shared links are readable by anyone who has them. Don't share conversations containing student PII, confidential materials, or exam questions that haven't gone out yet.

TIP 6: Start Fresh When Quality Starts to Slip

Context windows are large, but conversations don't age gracefully. After dozens of exchanges on a complex project, the model can lose track of early instructions, start to repeat itself, or drift from the constraints you established at the start. If you're working on something extended — drafting an exam, iterating on a major assignment,

developing a full course unit — and you notice the quality declining, don't troubleshoot the existing conversation. Start a new one.

Keep a short "context summary" for complex projects: a paragraph describing what you're building, your key requirements, and decisions already made. Paste it at the top of each new conversation. Thirty seconds of setup saves significant frustration.

A better long-term fix: Claude Projects

If you're working on something that will span multiple sessions — a full course unit, a semester-long exam bank, a research project — Claude Projects is worth setting up instead of relying on context summaries. A Project lets you write persistent instructions and upload reference documents that stay active across every conversation within that project. Your constraints, role assignments, and background materials are present automatically every time you open a new chat, without pasting anything.

How to create a Project in Claude:

- Click "New Project" in the left sidebar
- Give the project a descriptive name ("Torts Fall 2025" or "Bar Prep Module")
- Click "Set project instructions" and write your standing context — who you are, what you're building, your key requirements and constraints
- Upload any reference documents you want available throughout (syllabus, case list, assignment templates)

Every new conversation you start within the project inherits all of it

The caveat from the main tip still applies: individual conversations within a Project can still get long and stale. Projects don't fix that. What they do is ensure that when you start a fresh conversation, you're not starting from zero. The best practice is to start a new chat within the project.

ChatGPT's equivalent is a CustomGPT, which requires more setup but offers more configuration options. See Chapter 10. Gemini's equivalent is Gems. Both work on the same principle: persistent instructions and knowledge that survive across conversations.

PART II: PROMPT-CRAFT HABITS

TIP 7: Ask the AI to Ask You Questions First

Before you let the AI generate anything for a complex task, ask it to interview you. End your prompt with: "Before you start, ask me any clarifying questions you need to do this well."

The AI often knows exactly what information it needs to help you — it just defaults to guessing. Giving it permission to ask forces the model to surface those gaps rather than

paper over them.

Example: Instead of “Write me a hypothetical for my Torts class on products liability,” try: “I want a hypothetical for my 1L Torts class on products liability. Before you write it, ask me what issues I want to test, the appropriate level of complexity, and what fact patterns I’ve already used this semester.”

You’ll spend two minutes answering questions and get a hypothetical that fits your course. The alternative is a generic scenario that misses three of your learning objectives and needs twenty minutes of editing before it’s usable.

TIP 8: Ask for Multiple Versions

When a first draft isn’t quite right, resist the urge to edit it yourself immediately. Ask for two or three versions with different approaches, tones, or emphases. Version two is often exactly what you wanted, or the differences across versions clarify what you need.

Useful variations for legal education:

- “Write two versions of this hypothetical: one where the plaintiff clearly wins, and one that’s very close.”
- “Give me three versions of this exam question at different difficulty levels: straightforward, moderately difficult, and appropriately rigorous for a final.”
- “Write this feedback two ways: one for a student who needs encouragement, one for a student who needs a harder conversation.”

Multiple versions also work as a diagnostic on your own prompts. If all three versions look essentially the same, your constraints are too tight. If they look wildly different, your prompt wasn’t specific enough.

TIP 9: Ask the AI to Critique Its Own Output

After the AI produces something, ask it to evaluate what it just wrote. “What are the weaknesses in this analysis?” and “What did you leave out or oversimplify?” are underused follow-ups. The model often “knows” where its output is thin or hedged — it just doesn’t volunteer that information unless asked. You could also ask it, “What information can I provide that will help you produce stronger, more reliable output?”

Two follow-up prompts earn their keep above all others for legal education. First, whenever the AI produces legal analysis, ask: “What citations would typically support this analysis, and how confident are you that they’re real?” That question alone will catch hallucinated case names before they travel downstream into student work. Second, whenever you’ve asked the AI to write an exam question, ask: “Can you answer this question correctly without using the knowledge provided? If so, redraft using only the knowledge given to you.”

PART III: WORKFLOW TIPS FOR LEGAL EDUCATION

TIP 10: Build and Maintain a Prompt Library

Your best prompts are worth keeping. The prompt you spent forty-five minutes perfecting for your Contracts hypothetical generator shouldn't need to be recreated from scratch every semester. Keep a running document — Word, Notion, Google Doc, whatever you'll open — organized by task type.

For each prompt, save at minimum:

- The full text of the prompt
- Which platform you used (prompts don't always travel well across platforms)
- A note on what it produced and what worked
- Any variations you tested

The library grows quickly once you start. By your third semester, it's one of the most useful documents you have. As of May 2, 2026, my Prompt Library is 63 pages long separated by audience (for faculty training, staff training, class, demos, etc.) Honestly, it could be pruned a bit, but it still very useful and I rely on it often.

TIP 11: Stress-Test Your Exam Questions Before You Finalize Them

Before giving an exam, run every question through an LLM and ask it to answer as a well-prepared law student would. If the AI answers correctly from the text of the prompt alone — without applying specific case knowledge, statutory analysis, or a close read of the facts — the question may be more gameable than you intended.

The real test isn't whether the AI gets it right. Ask a follow-up: "How did you know that? What specific knowledge did you rely on?" If the answer is "the hypothetical essentially told me" rather than "I applied the Hand Formula to these specific facts," you should probably revise the question.

Running this check on a full exam takes about twenty minutes and catches problems that otherwise become noticeable (if at all) only after students have already taken it.

TIP 12: Generate a Bank of Hypothetical Variations

Once you have a hypothetical you like, ask the AI to generate variations. A single well-designed hypo can become five or six versions serving different purposes: classroom discussion, practice problems, makeup assignments, and final exam.

Sample prompt: "Here is a hypothetical I use in my Civ Pro class. [paste hypothetical] Generate four variations: one that makes subject matter jurisdiction clearer, one that makes it harder, one that adds a choice-of-law complication, and one appropriate for a final exam."

The LLM might not nail all four on the first pass, but you'll have working drafts in minutes.

TIP 13: Build Complex Assignments in Stages, Not All at Once

Don't ask for a complete, polished fifteen-part assignment in a single prompt. The LLM will make dozens of micro-decisions you never weighed in on, and the result will be generic enough to need substantial rebuilding. Construct complex assignments in stages: outline first, then section by section, reviewing at each stage before moving forward.

A typical chain for building an assignment from scratch:

Start with context and constraints: "I'm building a comparative analysis assignment for 2L students. Here's what I want them to practice: [describe learning objectives]. Here are my constraints: [length, research requirements, no AI tools]. Before you proceed, ask me any clarifying questions."

Approve an outline before drafting begins: "Here's the outline. Revise section 3 to require more independent legal research. Then draft section 1 only."

Draft section by section: Review each section before requesting the next. Don't let the LLM run ahead.

Run a final integration pass: "Here are all the sections. Make the tone and difficulty level consistent throughout, and flag anything that seems inconsistent with the learning objectives I described."

The extra time up front (maybe fifteen minutes) prevents you from receiving a 2,000-word draft that misses the point and needs to be rebuilt from scratch.