

CHEAT SHEET

AI Skills Framework

THE WWWHR FRAMEWORK

What

What tool should you use? Is it secure? Does it have the features most useful for your project? Does it have your information handy so that you don't need to spend much time getting it primed for your work?

When

When should you use AI during your project? A lot of the battle is know when to think about using AI, and when not to. At what steps in the writing or drafting process should you turn and ask AI for ideas or some work? Refer to the AI Tasks section for more insight here, and continue building your skills to get better at it.

Why

Why are you using AI? A vision of the results you're looking for will make it much easier to quickly identify useful results. If you keep your vision vague, then you can spend hours diving down AI rabbit holes. Take a moment to think about what the final version should look like.

How

How can you best leverage AI for this task? What prompting, saved prompts, context, documents, information, data can you give the AI to get the best results quickly?

Refine

Refine. Take stock of what works and what doesn't. Start saving prompts, templates, examples, and any other information that has helped you get better AI results, so that you can easily reuse it later. Take stock of your writing and other practice processes so you can better decide the WWWH in the future for each of the major tasks you work on regularly. Revisit your processes (including using AI to help) so that you can benchmark your work and continuously improve. The age of AI is the age of process.

WRITING & ARGUMENTATION SKILLS

Summarizing

Condensing legal documents, cases, or fact records into clear, useful summaries tailored to different goals and audiences.

- Tailor summaries by issue, audience, or task: internal, client-facing, or judicial.
- Use different lenses: timeline, doctrine, persuasive framing.
- *Sample prompt: "Summarize this deposition transcript focusing only on gaps in the timeline and witness credibility."*

Drafting

Generating first drafts or alternative versions of legal documents—from contracts to emails—faster and with creative variation.

- Start from bullet points, past templates, or messy notes.
- Ask for multiple takes to test different framings or tones.
- *Sample prompt: “Draft a first-pass operating agreement from these notes, using plain language for a startup client.”*

Styling

Adjusting tone, voice, and formality for different audiences and goals.

- Shift from neutral to persuasive, dense to digestible, formal to informal.
- Ask AI to mimic tone of a specific judge, partner, or past brief.
- *Sample prompt: “Rewrite this letter in a warmer, more collaborative tone while keeping the legal ask firm.”*

Brainstorming

Generating multiple ideas, framings, arguments, or structures to explore your best rhetorical approach.

- Use to kick off drafts or get unstuck mid-writing.
- Request ideas targeted to specific goals: urgency, empathy, impact.
- *Sample prompt: “Give me five different openings to this motion that emphasize urgency in different ways.”*

Reframing

Reworking facts, arguments, or organization to shift emphasis or impact.

- Try chronological, thematic, or issue-based ordering.
- Use to test alternative fact framings or counter-narratives.
- *Sample prompt: “Reframe this section to emphasize institutional failure over individual negligence.”*

Editing

Improving legal text at all levels: clarity, logic, and style.

- Layer edits: grammar only, logic only, tone only.
- Ask for flags, explanations, and rewritten alternatives.
- *Sample prompt: “Edit this memo for clarity and flow—leave comments explaining any logic gaps.”*

Comparing

Analyzing multiple legal texts side-by-side—e.g., briefs, clauses, rulings—to evaluate and improve.

- Use for strategic comparisons, e.g., your motion vs. theirs.
- Highlight tone, clarity, or logic differences.
- Sample prompt: “Compare these two versions of the clause and identify which is clearer and why.”

Synthesizing

Combining multiple legal sources into a coherent rule, theme, or analysis.

- Ask AI to extract common elements and unify them.
- Great for multi-case research or layered standards.
- Sample prompt: “Synthesize these five cases into a single articulation of the controlling test and its exceptions.”

Layering or Structuring

Helping you move between levels of abstraction: rule, application, policy, etc.

- Ask AI to bridge between layers or shift perspective.
- Use to connect granular facts to overarching doctrine.
- Sample prompt: “Turn this rule and analysis into a big-picture summary that connects back to client goals.”

Narrative Structuring

Shaping legal writing into a persuasive, emotionally engaging story.

- Ask for rising tension, character framing, or momentum building.
- Use to highlight injustice, credibility, or cause-effect arcs.
- Sample prompt: “Turn this timeline into a compelling narrative that highlights bad faith by the opposing party.”

Calibration

Tuning tone, aggressiveness, or caution based on posture and audience.

- Shift from aggressive to conciliatory, confident to cautious.
- Ask AI to calibrate tone to match the reader (judge, GC, regulator).
- Sample prompt: “Revise this brief to sound more neutral and solution-oriented, assuming a skeptical judge.”

Consistency Checking

Ensuring coherent voice, structure, and terminology across complex or multi-author documents.

- Flag inconsistencies in style, tone, argument, or formatting.
- Great for finalizing court filings or client-facing documents.
- Sample prompt: “Review this brief and highlight any inconsistent terminology or shifts in tone.”

Genre & Context Shifting

Transforming content across legal genres, formats, or contexts.

- Ask AI to convert a memo into a client alert, a checklist into a slide deck, or a contract into an argument outline.
- Use when repurposing content for different audiences or uses.
- *Sample prompt: “Turn this client FAQ into a training handout for new sales staff on legal dos and don’ts.”*

Emulation

Mimicking the writing style or structure of a person, judge, firm, or institution.

- Train AI on past writing to generate consistent tone, syntax, and formatting.
- Use for ghostwriting, reply briefs, or firm-branded content.
- *Sample prompt: “Rewrite this memo in the style of Judge X’s rulings from these two orders.”*

Personas

Tailoring tone, content, and framing to a specific audience persona—e.g., GC, regulator, skeptical judge.

- Ask AI to prioritize information and adjust tone to match the recipient’s mindset and goals
- Use for persuasive writing or strategic advice.
- *Sample prompt: “Frame this analysis as if written for a cautious GC with little tolerance for legal uncertainty.”*

Applying Rhetoric or Logic Structures

Using formal techniques to sharpen persuasion or clarity.

- Ask AI to apply structures like IRAC, Rule-Proof-Application, classical rhetorical moves, or analogical framing.
- Use to revise, test, or teach persuasive logic.
- *Sample prompt: “Rewrite this argument using a Rule-Proof-Application structure, with a metaphor to reinforce the policy theme.”*

Multilingual Support

Translating legal content and summarizing or clarifying it across languages.

- Use AI to translate and annotate foreign-language contracts, rulings, or correspondence.
- Especially useful in cross-border deals or international litigation.
- *Sample prompt: “Translate this Spanish employment agreement into English and highlight any unusual severance provisions.”*

Sanity Checks & Gut Checks

Asking AI for a second opinion on structure, logic, tone, or completeness.

- Use at the end of drafting or reviewing to check blind spots.
- Ask AI to rate strength, clarity, and risks.
- *Sample prompt: “Gut check this letter—do I sound too aggressive and have I clearly stated the client’s ask?”*

Extracting

Pulling key legal material from large or dense content sets.

- Use to isolate precedent, patterns, red flags, or useful citations.
- Great for filings, rulings, contracts, or data-heavy records.
- *Sample prompt: “Extract the governing law clauses and any termination triggers from these 30 contracts.”*

Pattern Recognition & Trends

Identifying hidden themes, recurring issues, or strategic patterns across documents or matters.

- Useful for surfacing market norms, repeated disputes, or overused arguments.
- Ask AI to detect what’s implicit, not just explicit.
- *Sample prompt: “From these NDAs, what appears to be the market standard on termination and indemnity?”*

THINKING & REASONING SKILLS

Strategizing

Mapping out legal options, outcomes, and back-planned tactics aligned with case or client goals.

- Use AI to generate pros/cons, likely responses, or settlement implications.
- Ideal for litigation, negotiation, or regulatory response planning.
- *Sample prompt: “Outline three strategy options for handling this breach claim, each with risk and cost trade-offs.”*

Pressure Testing

Challenging arguments by simulating scrutiny from judges, opposing counsel, or critical readers.

- Use AI to find logical gaps, weak premises, or missing support.
- Ask for counterarguments and rhetorical vulnerabilities.
- *Sample prompt: “Act as opposing counsel and critique this motion for any weak assumptions or unsupported conclusions.”*

Risk Analysis & Forecasting

Modeling potential future legal outcomes or regulatory reactions.

- Predict enforcement, litigation, discovery, or audit risks.
- Use to prepare clients for contingency planning.
- *Sample prompt: “Given this fact pattern and recent SEC actions, what enforcement risks are most likely?”*

Evaluation

Spot-checking logic, soundness, and sufficiency of legal reasoning.

- Ask AI to flag unsupported assertions, overstatements, or misapplied rules.
- Use as a final gut-check or critique pass.
- *Sample prompt: “Evaluate this memo and point out any places where the reasoning doesn’t follow from the facts.”*

Statutory Comparison & Synthesis

Analyzing multiple statutes across jurisdictions and synthesizing differences.

1. Helpful for compliance, multi-state litigation, or transactional planning.
2. *Sample prompt: “Compare data breach notification statutes in California, Texas, and Florida for financial institutions.”*

Negotiation Mapping

Planning negotiation strategy and anticipating positions, fallbacks, and options.

- Use AI to simulate BATNAs, deal dynamics, and creative middle grounds.
- Tailor by party, leverage, and goals.
- *Sample prompt: “Map out three likely negotiation paths with fallback terms for the indemnity clause.”*

Fact & Evidence Mapping

Linking facts to legal issues and available evidence in a structured, visual or logical format.

- Build issue trees, fact matrices, or burden-of-proof maps.
- Helps with discovery planning, motion support, or trial prep.
- *Sample prompt: “Create an issue-to-fact matrix showing which facts support each element of the fraud claim.”*

Reflective Lawyering

Analyzing your own work patterns, writing style, and practice habits for improvement.

- Use AI to review past briefs, memos, or time logs and generate coaching suggestions.
- Great for identifying overused structures, tone issues, or recurring blind spots.
- *Sample prompt: “Based on my last five motions, what clarity issues or style inconsistencies appear most often?”*

Generating Hypotheticals & Analogies

Creating hypotheticals or real-world analogies to test legal logic or teach others.

- Use to prepare for oral argument, explain issues to clients, or draft more vivid writing.
- Ask for analogies matched to tone or audience.
- *Sample prompt: “Generate three hypotheticals that test the outer limits of this ADA compliance standard.”*

CLIENT-FACING & COMMUNICATION SKILLS

Creating Client-Facing Materials

Drafting summaries, guidance, and explainers tailored for client use.

- Use for plain-language versions of legal documents, alerts, or deliverables.
- Adapt tone, detail, and focus for the specific client relationship.
- *Sample prompt: “Summarize this arbitration clause for a client unfamiliar with legal processes—include benefits and risks.”*

Visual Creation

Using AI to generate timelines, issue maps, flowcharts, or tables that clarify complex facts or rules.

- Great for litigation timelines, contract logic trees, or regulatory charts.
- Ask AI to draft diagrams based on narratives or checklists.
- *Sample prompt: “Turn this set of events into a visual timeline that highlights causation and delay points.”*

Follow-Up Communications

Drafting clear post-meeting emails, to-do lists, decision summaries, or progress reports.

- Use AI to convert notes or recordings into polished follow-ups.
- Tailor format and tone to client expectations.
- *Sample prompt: “Summarize this strategy meeting into a client-ready recap with next steps.”*

Client Onboarding Packets

Creating comprehensive welcome and orientation materials tailored to the matter and client.

- Include expected timelines, documents needed, contact info, and process overview.
- Ask AI to pull from precedent packets and customize per matter type.
- *Sample prompt: “Build a new client onboarding packet for a data privacy compliance matter, modeled on this past example.”*

Compliance Monitoring

Helping clients stay ahead of deadlines, updates, and jurisdictional obligations.

- Generate compliance calendars, reminders, or policy summaries.
- Use to automate ongoing compliance alerts or notices.
- *Sample prompt: “Create a compliance calendar for a biotech startup operating in California and New York.”*

Plain-Language Explanation

Translating legal documents or concepts into client-accessible prose.

- Eliminate jargon, explain terms, and connect issues to business impact.
- Ask for tone variation by client experience level.
- *Sample prompt: “Rewrite this indemnity clause so a non-lawyer board member can understand the key risks.”*

Client Q&A Assistance

Supporting fast, accurate responses to routine or predictable client questions.

- Use AI to draft answers from internal precedent, handbooks, or firm knowledge bases.
- Can be integrated into chat-style help or live prep.
- *Sample prompt: “Draft a response to a client asking what ‘reasonable efforts’ means in this clause.”*

Educational Content Drafting

Creating CLE decks, client alerts, or internal training materials on legal topics.

- Ask AI to summarize statutes, provide examples, or build quiz questions.
- Tailor by audience sophistication and tone.
- *Sample prompt: “Create a 5-slide training outline explaining the difference between employees and contractors under California law.”*

Marketing & Thought Leadership Writing

Producing blog posts, newsletter content, or social media updates.

- Translate legal substance into conversational, high-value insights.
- Ask for tone to match your brand, audience, or publication.
- *Sample prompt: “Write a 500-word blog post summarizing the recent SEC crypto enforcement action and its implications for startups.”*

Proposal & Pitch Drafting

Creating tailored materials for prospective clients or RFPs.

- Summarize experience, customize value propositions, and draft team bios.
- Use AI to assemble pitches from firm bios, past matters, and client goals.

- Sample prompt: “Draft a pitch letter for a pharmaceutical client based on these sample bios and experience summaries.”

ADMINISTRATIVE, WORKFLOW & PRACTICE IMPROVEMENT SKILLS

Organizing

Structuring large or messy content—e.g., discovery files, contract sets, or research notes—into usable formats

- Ask AI to group, sort, or label content based on topic, issue, or chronology.
- Use for clarity in long-term matters or complex documents.
- Sample prompt: “Organize these 300 emails into five categories based on relevance to the whistleblower claim.”

Legal Tech Tool Evaluation

Using AI to compare and assess potential software tools or vendor services.

- Especially helpful for GCs, ops leads, or innovation teams.
- Sample prompt: “Compare Ironclad, SpotDraft, and LinkSquares based on features useful for a 5-person in-house team.”

Internal Policy & Manual Drafting

Drafting or updating internal firm policies, handbooks, or procedures.

- Ask AI to align with best practices and local laws.
- Sample prompt: “Create a social media policy for a midsize law firm based on recent NLRB guidance.”

Admin Automation

Turning routine administrative tasks into automated workflows.

- Generate checklists, auto-fill forms, clean up transcripts, or anonymize sensitive data.
- Especially useful for task follow-ups, scheduling, and compliance logs.
- Sample prompt: “Create a pre-deposition checklist from these case notes and prior templates.”

Process Design & Improvement

Creating or refining standard legal workflows for recurring work.

- Map out steps, roles, and templates for better speed and quality.
- Ask AI to flag inefficiencies and suggest alternatives.
- Sample prompt: “Design a streamlined contract review process for a three-person in-house team.”

Information Hygiene

Checking for metadata, citation accuracy, formatting consistency, and document integrity.

- Useful at final draft stage or when combining multi-author work.
- Ask AI to flag defined term mismatches, citation inconsistencies, or stray formatting.
- *Sample prompt: “Review this memo for Bluebook accuracy, defined term consistency, and hanging headings.”*

Learning & Tutoring

Turning legal content into teachable material or personalized lessons.

- Ask AI to explain, quiz, or reframe concepts for trainees or clients.
- Excellent for new associates, clients, or lateral onboarding.
- *Sample prompt: “Explain this patent claim construction order as if I’m a new associate seeing it for the first time.”*

Work Improvement & Feedback Loops

Extracting lessons and upgrades from completed matters or documents.

- Analyze past work product for clarity, timing, or success patterns.
- Use for institutional knowledge capture and future process refinement.
- *Sample prompt: “From these past litigation memos, identify common gaps in early factual development.”*

Staying Current

Tracking new legal developments and summarizing them by relevance.

- Ask AI to monitor case law, regulations, or enforcement trends.
- Tailor updates by industry, jurisdiction, or client base.
- *Sample prompt: “Give me a weekly update on new California privacy rulings affecting SaaS companies.”*

Billing Narrative Drafting

Creating polished, compliant, and accurate billing entries.

- Turn meetings, notes, or documents into clear time entries.
- Match formatting to client or system requirements.
- *Sample prompt: “Draft billing entries from these notes for a tech client, following their narrative style guidelines.”*

Matter Timeline Estimation

Forecasting key milestones and deadlines for legal matters.

- Base predictions on similar past matters, templates, or known bottlenecks.
- Ask AI to create visual timelines or staffing plans.
- *Sample prompt: “Estimate the timeline and team for closing this \$15M Series A round, based on this checklist and client info.”*

Knowledge Management Intake

Tagging and organizing documents for internal reference libraries.

- Sort by clause, issue, jurisdiction, industry, or success outcome.
- Ask AI to identify duplicates or underused models.
- *Sample prompt: "Auto-tag these agreements by jurisdiction, termination clause type, and dispute resolution method."*

Template & Clause Bank Maintenance

Improving and streamlining reusable language.

- Identify outdated, redundant, or inconsistent clauses.
- Align templates with firm or client preferences and current market norms.
- *Sample prompt: "Compare these service agreement templates and consolidate into a clean master version."*

Junior Lawyer Coaching

Simulating partner feedback, explaining edits, or offering writing support.

- Ask AI to rewrite, critique, or annotate examples to teach better habits.
- Use to model professional tone, logic, and structure.
- *Sample prompt: "Act as a partner reviewing this memo and leave comments on structure, tone, and clarity."*

Conflicts Checking Support

Summarizing and flagging potential client conflicts across past matters.

- Use AI with structured data or unstructured matter descriptions.
- Ask for high-level summaries or specific red flags.
- *Sample prompt: "Summarize any possible conflicts between this prospective biotech client and our past IP litigation work."*

Privilege & Work Product Review

Reviewing documents for privilege risk or business/legal content blending.

- Ask AI to highlight risky phrasing or recommend redactions.
- Use for decks, emails, and meeting summaries.
- *Sample prompt: "Flag any content in this internal strategy deck that could jeopardize attorney-client privilege."*

Regulatory Risk Forecasting

Modeling enforcement likelihood, reputational exposure, or audit risk.

- Ask AI to combine client facts with enforcement history and trends.
- Use to build early warning systems and response playbooks.

- Sample prompt: “Given this advertising model and FTC enforcement history, what compliance risks should we flag?”

Stress & Complexity Flagging

Having AI identify likely high-stress or complexity areas in documents or workflows.

- Helpful for workload triage or support.
- Sample prompt: “Highlight areas in this trial calendar that may be unmanageable for a two-person team.”

Coaching & Mindset Reframing

Using AI as a mindset and communication coach to reframe difficult conversations or feedback.

- Use when navigating challenging emails, team dynamics, or client pushback.
- Sample prompt: “Help me reframe this negative feedback email so it’s firm but constructive.”

Wellness-Integrated Time Management

Helping lawyers balance billable tasks with breaks, wellness, and focused deep work.

- Use AI to restructure task lists with built-in recovery time.
- Sample prompt: “Break this to-do list into 90-minute sprints with short breaks and lower-cognitive-load tasks in the afternoon.”

LITIGATION PRACTICE SUPPORT SKILLS

Deposition Preparation

Generating outlines, question sets, impeachment material, and prep packets for witnesses or attorneys.

- Ask AI to simulate opposing answers or anticipate evasive responses.
- Build themes and key facts into the question flow.
- Sample prompt: “Create a deposition outline focused on the CFO’s knowledge of internal controls and inconsistencies in their past statements.”

Motion Practice Planning

Outlining motion strategy, argument structure, and local rule requirements.

- Ask AI to prioritize motion types based on facts, timing, and judge profile.
- Use for early-stage planning or motion sequencing.
- Sample prompt: “Given this case and judge, suggest three likely-to-succeed pretrial motions and how to frame each one.”

Discovery Management

Summarizing responses, spotting inconsistencies, and generating follow-up tasks.

- Use AI to find incomplete answers, evasions, or cross-document inconsistencies.
- Ask for targeted follow-up questions or sanction motion options.
- Sample prompt: “Summarize these RFP responses and flag which ones are incomplete, evasive, or raise privilege concerns.”

Pretrial Planning

Generating checklists, exhibit lists, trial themes, and logistical plans.

- Ask AI to organize tasks across teams, prep timelines, or generate trial week workflows.
- Use for coordination across litigation support, witness teams, and filings.
- Sample prompt: “Draft a pretrial checklist for a two-week federal jury trial including deadlines, logistics, and witness prep.”

Drafting Jury Instructions

Customizing pattern instructions or proposing tailored language to match your theory of the case.

- Ask AI to synthesize model instructions and suggest edits.
- Use to align language with your key themes and burdens.
- Sample prompt: “Draft a negligence instruction based on 9th Circuit pattern language, tailored to a product defect claim.”

Simulating Cross/Direct Examination

Using AI to simulate witnesses or attorneys to test question impact and story flow.

- Ask AI to role-play evasive or hostile witnesses based on real deposition or trial testimony.
- Use to train junior litigators or rehearse key moments.
- Sample prompt: “Play the role of the plaintiff in this employment case and answer 10 direct questions. Then switch to hostile cross-mode.”

Brief Bank Mining

Extracting reusable arguments, rule language, or structure from a set of past filings.

- Ask AI to find persuasive hooks, analogies, or citations used previously.
- Use to speed drafting or maintain firm tone consistency.
- Sample prompt: “Pull the strongest causation arguments from our last five MSJs in product liability cases.”

TRANSACTIONAL PRACTICE SUPPORT SKILLS

Due Diligence Review

Summarizing, flagging, and organizing key clauses, deviations, and risks from documents.

- Ask AI to compare terms to standards, prior deals, or firm templates.
- Generate summary tables, red flag memos, or client-ready summaries.
- *Sample prompt: “Review these 20 commercial leases and identify any unusual assignment or subletting restrictions.”*

Term Sheet Drafting & Translation

Turning bullet points into structured term sheets—or full agreements from term sheets.

- Ask AI to identify missing standard terms or flag ambiguous phrasing.
- Use to align parties early and surface redline hotspots.
- *Sample prompt: “Draft a convertible note term sheet from these negotiation notes and add standard terms where missing.”*

Closing Checklist Creation

Generating comprehensive deal checklists and assigning tasks by role and timeline.

- Tailor checklists to deal type, jurisdiction, or client.
- Ask AI to group documents by priority, party, or approval stage.
- *Sample prompt: “Create a closing checklist for a private asset purchase involving IP, employment, and international compliance.”*

Negotiation Simulation

Testing how proposed changes will impact leverage, position, or deal risk.

- Ask AI to simulate counterparty responses or highlight concession paths.
- Use to explore fallback positions or escalation triggers.
- *Sample prompt: “How might the other side respond if we insist on deleting the limitation of liability cap?”*

Post-Deal Analysis

Extracting insights from what caused deal friction, delays, or concessions.

- Use for firm knowledge capture, template improvements, or client reporting.
- Ask AI to identify repeat bottlenecks or decision inflection points.
- *Sample prompt: “From these annotated agreements, summarize which clauses required the most negotiation and why.”*

Drafting & Editing (Transactional)

Generating, polishing, or explaining contracts and commercial documents.

- Ask AI to draft clauses, clean up dense language, or rewrite for plain English.
- Great for translating business needs into legal language (and vice versa).
- *Sample prompt: “Rewrite this indemnification section in plain English and explain the business impact of each subclause.”*

CHEAT SHEET

10 Highly Detailed AI Prompts for Writing a Legal Brief

These prompts guide AI through each stage of writing a legal brief. But note something important: Each assumes you are providing high-quality source materials, not asking the AI to find that source material for you. This ensures you have much better accuracy and quality work product. The AI is working with your source notes, authorities, and ideas and brainstorming and ideating.

Each prompt is designed to maximize AI's strengths while keeping the lawyer in control of strategy, legal analysis, and argumentation.

1. AI-ASSISTED CASE & STATUTE EXTRACTION (PRE-DRAFTING RESEARCH STAGE)

Goal: Summarize and structure key takeaways from uploaded legal sources (cases, statutes, prior briefs).

Prompt:

Context: You are assisting a lawyer in analyzing case law, statutes, and relevant briefs. Your task is to extract and summarize key legal principles only from the provided materials.

Task: Review the uploaded documents and summarize:

- **Governing Rule Statements:** Clearly articulate the controlling legal principles
- **Key Cases & Holdings:** Summarize major precedents, including their reasoning and jurisdictional weight
- **Supporting vs. Opposing Authority:** Identify favorable and unfavorable case law
- **Unresolved Issues or Conflicts:** Highlight any ambiguities or competing interpretations

Materials Provided: [Attach cases, statutes, prior briefs].

Output Format:

- **Rule Statement:** [Summarized rule]
- **Key Case Law:** [Summarized cases]
- **Counterarguments or Weaknesses:** [Potential issues]
- **Unresolved Questions**

Why This Works: Forces AI to only analyze verified legal materials and structure the research like a legal memo.

2. AI-ASSISTED ISSUE FRAMING (STRUCTURING THE ARGUMENT BEFORE DRAFTING)

Goal: Ensure legal issues are framed persuasively and aligned with case strategy.

Prompt:

Context: You are structuring the core legal issues for a brief. Based on the provided materials, you will craft precise, persuasive legal issue statements that set up strong arguments.

Task: Review the provided sources and draft legal issue statements in the most effective format:

- Neutral Framing: Traditional, objective statement of the issue
- Persuasive Framing: Worded to subtly favor our client's position
- Opposition's Likely Framing: How the other side will present the issue

Materials Provided: [Attach case law, client facts, opposing arguments].

Output Format:

- Neutral Issue Statement: [Traditional statement]
- Persuasive Issue Statement: [Strategic framing]
- Opposition's Issue Statement: [How the other side will frame it]

Why This Works: Helps lawyers refine their framing strategy before drafting the brief.

3. AI-STRUCTURED FIRST DRAFT (USING ONLY PROVIDED MATERIALS)

Goal: Generate a structured first draft of a legal brief section.

Prompt:

Context: You are drafting a legal brief section based solely on the provided materials. Your task is to generate a well-organized, precise, and persuasive argument.

Task: Using the supplied case law, statutes, and factual background, draft a structured brief section with:

- Clear IRAC or CREAC formatting (Issue, Rule, Application, Conclusion)
- Direct citations to provided sources (no external or AI-generated law)
- Logical flow with transitions between arguments

Materials Provided: [Attach research, case law, factual background].

Output Format:

- Issue Statement
- Rule Explanation
- Application to Facts
- Conclusion

Why This Works: AI drafts within a structured legal framework and ensures clarity without adding unreliable content.

4. AI-POWERED ARGUMENT REFINEMENT

Goal: Strengthen the persuasiveness of legal arguments.

Prompt:

Context: You are improving the strength and persuasiveness of a legal argument.

Task: Rewrite the argument to:

- Enhance logical flow and clarity
- Bolster citations with provided case law
- Preempt and neutralize opposing arguments

Original Argument: [Insert draft text].

Materials Provided: [Attach supporting case law, statutes].

Output Format:

- Improved Argument: [Rewritten version]
- Why These Changes Matter: [Explanation of improvements]

Why This Works: AI sharpens persuasion and structure while staying within legal constraints.

5. AI-GENERATED COUNTERARGUMENTS & REBUTTALS

Goal: Strengthen responses to opposing arguments.

Prompt:

Context: You are developing counterarguments for a legal brief. Using the provided opposing arguments, draft well-supported rebuttals.

Task:

- Summarize the opposition's main points
- Develop rebuttals based on provided case law
- Ensure logical structure and persuasive framing

Opposing Argument: [Insert opposing counsel's argument].

Materials Provided: [Attach case law supporting rebuttal].

Output Format:

- Opposing Argument Summary
- Rebuttal (Legal Basis & Analysis)

Why This Works: AI systematically counters opposing arguments using only trusted sources.

6. AI-ASSISTED LEGAL STYLE & TONE ENHANCEMENT

Goal: Adjust legal writing style based on audience.

Prompt:

Context: You are revising a legal brief for tone and readability. Based on the audience (judge, arbitrator, client), adjust the tone while maintaining precision.

Task: Rewrite the provided section in different tones:

- Formal & Neutral (Judicial Audience)
- Assertive & Aggressive (Adversarial)
- Conciliatory & Persuasive (Negotiation)

Original Argument: [Insert legal text]

Output Format:

- Formal & Neutral
- Assertive & Aggressive
- Conciliatory & Persuasive

Why This Works: AI helps test different persuasive styles for maximum impact.

7. AI FOR CITATION & FORMATTING REVIEW

Goal: Check citations for consistency and accuracy.

Prompt:

Context: You are reviewing a legal document for citation accuracy and proper formatting.

Task:

- Identify missing or inconsistent citations
- Ensure compliance with [Bluebook/firm style]
- Suggest corrections

Legal Text to Review: [Insert brief section].

Why This Works: AI acts as a proofreader for citation integrity and formatting consistency.

8. AI-ASSISTED CONCLUSION OPTIMIZATION

Goal: Strengthen the impact of a legal brief's conclusion.

Prompt:

Context: You are refining the conclusion of a legal brief to make it more persuasive and impactful.

Task:

- Summarize key arguments concisely
- Use compelling final phrasing
- Ensure it aligns with the overall legal theory

Original Conclusion: [Insert draft].

Why This Works: AI ensures conclusions are strong and leave an impact.

9. AI FOR LOGICAL CONSISTENCY & FLOW REVIEW

Goal: Ensure that a legal brief is logically structured, free of contradictions, and flows smoothly.

Prompt:

Context: You are reviewing a legal brief to check for logical consistency, coherence, and clear transitions.

Task:

- Identify any inconsistencies (e.g., contradictory arguments, shifts in legal reasoning)
- Check paragraph transitions to ensure a smooth, logical flow
- Suggest restructuring if needed to make arguments clearer

Legal Document to Review: [Insert brief section].

Output Format:

- Inconsistencies Identified
- Suggested Fixes for Logical Flow
- Transition Improvement Suggestions

Why This Works: AI flags weak connections between arguments and ensures that the brief flows persuasively.

10. AI FOR FACT-TO-LAW INTEGRATION REVIEW

Goal: Ensure that the brief's factual background is seamlessly integrated into legal arguments.

Prompt:

Context: You are reviewing a legal brief to assess how well it connects the factual background to the legal arguments. Your goal is to ensure that.

- Key facts are woven into the argument effectively (not just dumped in the background section)
- The law is applied clearly to the facts (strong fact-law integration)
- Any missing factual support is flagged (if an argument lacks necessary fact-based grounding)

Legal Document to Review: [Insert brief section].

Case Facts: [Insert factual summary.

Output Format:

- Facts Not Fully Integrated
- Suggested Revisions to Improve Fact-Law Connection
- Any Unsupported Legal Assertions

Why This Works: AI ensures lawyers are applying case facts persuasively rather than keeping them siloed.

CHEAT SHEET

Advanced Prompt Engineering

There's no "perfect prompt" (at least not yet). You'll want to experiment and come at prompting challenges from different angles and compare the results. But save good prompts when they work well and reused them—prompting is personal right now.

"Let's try this from different angles. Rewrite this sentence in three different ways using different styles..."

Weights and priorities are key for these language models. So reinforcing the most critical goals for your task and being clear about the most important priorities helps. This is likely part of why emerging research suggests that using strong emotional and confident language gets better results. Powerful verbs, like "transform my writing into a very clear bulleted list" may get better results than a bland version of the same request. As will "you are doing a great job!" added onto your prompt.

"Above all, ensure that you do not change the substance of my writing at all. This is the most important task to keep in mind. You are doing a great job!"

Use GAI tools to give you feedback on how to better organize, simplify, and clarify your prompts. Experts often call this self-improvement prompting:

"You're a generative AI prompting expert. Review my prompt below and revise it to be as effective, clear, and unambiguous as possible. Take it step by step and explain your reasoning at each stage..."

The most effective prompts are often ones that break down tasks into smaller steps for the GAI and then ask it to do one thing at a time. This is often called "Skeleton of Thought" prompting. The idea is that you first have the GAI tool create an outline of its response, or only the first step of its response.

"Provide 10 alternatives for better verbs in this sentence"

[then after the response] "Now rewrite the sentence using these 10 different versions."

You can give GAI tools multiple roles. You might want to give your GAI the role of both a researcher and a judge to trigger both of those perspectives. You can also give GAI diverging roles.

“Take on the role of 10 judges with different ideological perspectives about Constitutional law. Then analyze my brief from different perspectives to explore weaknesses from different angles...”

Ask GAI to explain its reasoning as it provides your answer (or provide some high-level guidance about the reasoning it should use, without listing every step). You can also ask it to provide citations for each statement it makes. This forces GAI to come up with the steps it will take to reach a conclusion—and identify potential pitfalls that you can avoid.

“Your task is to determine if a solution is correct or not. To solve the problem, do the following.

First, work out your own solution to the problem.

Second, compare your solution to the lawyer’s solution and evaluate if the lawyer’s solution is correct or not. Don’t decide if the lawyer’s solution is correct until you have done the work yourself.

Third, provide your feedback to the lawyer.”

Another form of COT prompting is to ask the GAI tool to evaluate its responses using different criteria, then to redo its entire reasoning based on any insights it comes up with:

“Once you’ve completed your task, thoroughly review your answer and ensure it completely and fully answers the questions I provided you. Use whatever insights you find during this final review to produce a new response using what you learned.”

Some final GenAI tips:

- ✓ **Generated Knowledge Prompting.** Ask the GAI tool to generate potentially useful information about a given question before generating a final response. This is a kind of follow-up. The idea is that you can often get better results by first getting some starting information or data from GAI. Then use that response to craft a more detailed second prompt.
- ✓ **Affirmative instructions work better than negative ones.** Avoid saying not do something, instead specify what to do, when possible. For example, it's better to get more specific with "provide only a single line of text in your answer" as opposed to "do not provide lengthy text."
- ✓ **Repeat instructions at the end.** The research suggests that GAI tools can be susceptible to recency bias, which means that information at the end of the prompt might have more significant influence over the output than information at the beginning. So try including a reminder of the key instructions at the end of your prompt.
- ✓ **Consider giving the GAI a path if it can't find a good answer.** Say you are asking GAI to edit your writing. What if your writing is already perfect? Don't make GAI make stuff up: Tell it that if it has no improvements to simply respond with "this is already great!" Tell the GAI if it can't find a source, to respond with "I cannot find an answer." These sorts of parameters will help encourage the tool to be accurate.
- ✓ **Diverse Prompting.** The idea here is that you create multiple versions of an important prompt you want to perfect. Five is the best target number from the research. To make this easy, you can ask the GAI to create the versions for you "Please create five different versions of the following prompt. Make each version unique and from a different perspective." Then run each prompt and compare results.

Some final GenAI tips:

- ✔ Provide the GAI with any background information that could be helpful for your legal task, like your jurisdiction, the most important cases, key legal concepts, and so on.
- ✔ Provide the GAI with analogies to help it understand complex concepts or relationships more easily.
- ✔ Provide bullet points or content from your writing project to get GAI started. The more context you provide, the better.
- ✔ Keep in mind different GAI tools have different “context windows”—meaning they can only handle so much text in your prompts.
- ✔ Be clear and specific with your feedback. When providing feedback to a GAI, be as clear and specific as possible. This helps the GAI understand your expectations and refine its response accordingly.
- ✔ Provide the GAI with examples based on its response—what did it do well and what didn't it? Quotes from a GAI's response work great.
- ✔ When following up, address the GAI directly, and explicitly state that you are asking a follow-up question. This can help GAI understand the context and provide an appropriate response.
- ✔ If the GAI seems unsure about your prompt or the desired outcome, encourage it to ask clarifying questions. This can help improve its understanding of your requirements and lead to better results.
- ✔ If you're not satisfied with the GAI's response, ask it to provide alternative answers or approaches. This can help you explore different perspectives and solutions.

Some final GenAI tips:

- ✔ While providing numerous examples is helpful, ensure that each example is high-quality and relevant to your desired outcome. Poor examples can inject confusion.
- ✔ When crafting examples, try to use situations from actual legal cases or documents to provide GAI with a more accurate understanding of real-world scenarios.
- ✔ Structure your examples as prompt-answer pairs, presenting a problem and its solution, to help GAI identify patterns and learn to generate good responses.
- ✔ Mix both long and short examples to help GAI learn how to handle different levels of complexity and verbosity in responses.
- ✔ Show GAI examples of both correct and incorrect answers for a given problem, so it can learn to differentiate between the two.
- ✔ If you're not satisfied with the GAI's response, ask it to provide alternative answers or approaches. This can help you explore different perspectives and solutions.
- ✔ Start with simpler prompts and examples to familiarize the GAI with the task, then progressively introduce more complex scenarios.
- ✔ Diversify your examples: Ensure that your examples cover a wide range of scenarios to achieve diverse results.
- ✔ Randomize the order: Arrange your examples randomly unless you are specifically looking for a pattern.

Some final GenAI tips:

- ✓ Instruct GAI to generate multiple, independent answers or ideas to increase the chances of obtaining a great response.
- ✓ Use key words like "summary" or "detailed" to help guide the preferred response depth.
- ✓ Give GAI a reading grade or specific audience to write to, which will help refine its response.
- ✓ There are nearly endless ways to set parameters for detail or depth GAI will provide. For example, some experts suggest the 5 w's: asking GAI for the who, what, when, where, and why of a topic.
- ✓ You can even ask many GAI tools for a format that will optimize your ability to paste its results into a spreadsheet, or instead, a format optimized to paste into a Word document.
- ✓ Incorporate "if-then" statements to guide GAI's reasoning and focus on specific conditions. For example, "if you cannot find a source to support my statement, then respond 'no source found.'"
- ✓ Instruct GAI to compare and contrast ideas or concepts to provide a more thorough understanding.
- ✓ Instruct GAI to provide citations or references for the information it uses in its response, ensuring accuracy and credibility. For example "Provide a citation to a support source for every statement you make."
- ✓ Use explicit reasoning constraints. For example, you could instruct GAI to "consider only the relevant case law and statutes from the jurisdiction of X."

Some final GenAI tips:

- ✔ Come at difficult tasks from different angles. GAI is unlike any other tech tool you've worked with before. But it may take some creativity and trial and error. For example, researchers have confirmed better results on certain tasks by having the GAI explain its thinking process (for example, when asking the GAI to choose between options, have the GAI explain the reasoning or strengths and weaknesses of each option).
- ✔ Stack prompts of different types. Provide a role, instructions, parameters, context/hints, and examples all together: "Act as a federal judge. Explain how defamation works. Provide two paragraphs of less. Hints: intent, defamatory, falsity." Prompts with more structured guidance like this will get you the best results.
- ✔ Don't get discouraged with poor results at first. Working with GAI is a skill set, albeit one you need no technical background for. Give yourself time to experiment. There is no one right way to work with these AI tools (at least, not yet).
- ✔ Ask GAI to summarize your instructions to ensure it has a clear understanding of what you want.
- ✔ Consider using "certainty" prompting. This tactic instructs the GAI to provide a statement about how certain it is in the accuracy of its statements. This can help reduce accuracy issues and flag responses that you may want to triple check (rather than the double-check you do on everything GAI provides you).
- ✔ And finally, invest in your work. When you find a good prompt (or a good structure for certain types of prompts) save them so you can reuse them later. You can also use GAI to help you revise your prompts, improve them, organize them, compress them into fewer words, and more.

CHEAT SHEET

GAI Pitfalls

Professional Conduct and Competence

Duty of Competence

- ▶ **Review AI-Generated Content:** Ensure accuracy, relevancy, and compliance with laws.
- ▶ **Verify AI Drafts:** Scrutinize AI-drafted documents, like contracts, for client representation and legal accuracy.

Diligence

- ▶ **Active Monitoring:** Regularly check and update AI-generated content.
- ▶ **Ongoing Learning:** Stay informed about AI advancements and their implications in legal contexts.

Client Confidentiality and Data Privacy

Confidentiality Risks

- ▶ **Data Storage:** Be cautious of AI tools storing/sharing client information.
- ▶ **Secure Practices:** Implement measures to prevent unauthorized data access.

Privacy Compliance

- ▶ **Data Handling:** Understand and comply with data privacy laws relevant to AI usage.
- ▶ **Client Consent:** Inform clients about AI usage and obtain necessary consent.

Bias and Fairness

Inherent Biases in AI

- ▶ **Bias Awareness:** Recognize potential biases in AI's training data.
- ▶ **Diverse Data:** Encourage the use of varied and unbiased data sets in AI training.

Mitigation Strategies

- ▶ **Bias Detection Tools:** Utilize tools to identify and correct biases in AI outputs.
- ▶ **Critical Evaluation:** Continuously assess AI-generated content for subtle biases.

AI Limitations in Legal Contexts

Accuracy and Nuance

- ▶ **Complex Concepts:** Acknowledge AI's limitations in understanding complex legal principles.
- ▶ **Supplemental Judgment:** Use professional judgment to fill gaps in AI-generated content.

Jurisdiction-Specific Issues

- ▶ **Formatting and Regulations:** Adapt AI outputs to meet local legal standards and formats.
- ▶ **Case-Specific Tailoring:** Customize AI suggestions to suit individual case nuances.

Over-Reliance and Skill Development

Dependency Concerns

- ▶ **Critical Thinking:** Avoid over-dependence on AI for legal drafting and reasoning.
- ▶ **Skill Maintenance:** Ensure continual development of legal argumentation and writing skills.

Impact on Education

- ▶ **Teaching Balance:** Use AI as a supplementary tool in legal education without compromising learning of fundamental skills.

Question Framing and AI Guidance

Effective Prompting

- ▶ **Specific Queries:** Frame questions narrowly and guide AI to relevant contexts.
- ▶ **Logical Steps:** Instruct AI to follow logical sequences in problem-solving.

Understanding AI Logic

- ▶ **Model Limitations:** Be aware of AI's reasoning differences compared to human thought processes.
- ▶ **Prompt Refinement:** Continually refine prompts for better-aligned AI responses.

Security and Data Usage by AI Providers

Information Security

- ▶ **Data Usage Policies:** Understand and adhere to AI provider's data usage and security policies.
- ▶ **Confidentiality Measures:** Avoid inputting sensitive or identifiable client information.

Transparency with Clients

- ▶ **Disclosure:** Inform clients about the use of AI tools and associated risks.
- ▶ **Consent Protocols:** Develop protocols for obtaining client consent when using AI in their matters.

Keeping Pace with Technological Advances

Continuous Update

- ▶ **AI Evolution:** Stay informed about new versions and capabilities of AI models.
- ▶ **Adaptation:** Be prepared to adapt legal practices in line with AI advancements.

Future Implications

- ▶ **Long-Term Strategies:** Plan for the integration of advanced AI tools in legal processes.
- ▶ **Ethical Considerations:** Continually reassess ethical implications as AI technology evolves.

CHEAT SHEET

Common AI-isms

DON'T GET CAUGHT WITH THESE IN YOUR WRITING!

The following writing structures, words, and formats can signal that the writer may have relied a bit too much on robots. These are not all wrong—and sometimes these structures can be useful. But they are overused in modern AI models. So watch for them.

“No X. No Y. Just Z.”

“No fluff. No theory. Just actionable insights.”

AI loves this structure. You see it everywhere. Real humans don't talk in constant negation patterns.

“It's not about X, it's about Y”

“Success isn't about working hard, it's about working smart.”

This binary contrast appears in every AI-written post. Sometimes useful, but AI overuses it to death.

Triple Threat Syndrome

“Fast, efficient, reliable.” “Boost engagement, increase conversions, maximize ROI.”

AI learned that grouping things in threes makes “good writing.” So it does it constantly. Vary your rhythm. Two things. Four things. One thing. Mix it up.

“Game-changer” and “supercharge”

If you had a euro for every time AI wrote “game-changer,” you'd buy OpenAI. Describe actual change with specific metrics instead.

Overusing “real”

“Just real strategy from real experts getting real results.”

Repeating “real” doesn't make you authentic. Specific examples do.

“To your success”

Email sign-offs like this instantly reveal AI wrote it. Even successful founders slip up and use this without reading their emails.

Sign off like a human.

Arrow overload

→ 10x your results → scale faster → join today

One arrow per post maximum. Maybe.

Emoji explosion

🚀 Boost productivity now!

💡 Revolutionary approach!

✨ Transform your business!

ChatGPT treats emojis like confetti. I don't use emojis AT ALL.

Corporate verb disease

“...highlighting the benefits...” “...facilitating better outcomes...”

“...leveraging synergies...”

Use simple verbs. “Show” not “highlighting.” “Help” not “facilitating.”

“Use” not “leverage.”

Thesaurus abuse

Unless your audience wears suits to Zoom calls, stop using:

“utilize” (use)

“implement” (start)

“optimize” (improve)

“execute” (do)

Write like you talk.

Hedging language

“It’s worth considering...” “You might want to think about...” “It’s important to note that...”

Just say the thing. Skip the diplomatic warm-up.

Short hook questions

“The best part?” “Want access?” “Ready to level up?”

Very 2023 ChatGPT energy. Ask real questions that need thought to answer.

“Enter: [thing]”

“Enter: my revolutionary framework”

Just introduce the thing normally.

“And here’s the kicker”

AI’s favorite transition phrase. Sounds like a 3am infomercial.

Em dash overload

“This approach—which many experts recommend—can transform your business—if you’re willing to put in the work—and see real results.”

Yes, writers use em dashes. No, not 6 times per paragraph. Mix your punctuation.

“X changed everything”

“This strategy changed everything.”

Really? Every single thing? Use specific changes with actual numbers.

Symbol obsession

“This symbolizes...” “Which reflects...” “Emphasizing the importance of...”

Say what happened. Say what you learned. Skip the literary analysis.

Over-structured thinking

AI loves tidy parallelism: every paragraph the same length, every sentence the same cadence.

Humans meander. They interrupt themselves. They change pace.

Perfectly symmetrical paragraphs

Three sentences. Same rhythm. Same length. No variation. Looks polished. Reads robotic.

Excessive “frictionless productivity” clichés

“Unleash your potential.” “Unlock new efficiencies.” “Streamline your workflow.”

If it sounds like a corporate wellness poster, it’s AI.

Fake authority voice

“Seasoned professionals know...” “Savvy leaders understand...”

Real experts show receipts. AI invents vague group identities.

Overusing “simply” or “easily”

“Simply follow these steps.” “Easily improve your results.”

Humans know things are rarely simple.

Teaching-tone everything

AI narrates like a YouTube tutorial:

“First, you’ll notice...” “Next, you’ll want to...” “Finally...”

Even when it’s not actually teaching.

The Inexplicable Keanu Reeves Style Pause

“The truth? You already know the answer.”

AI picked up this melodramatic beat from influencer copy. Humans rarely stage-whisper their wisdom.

The pretend mic drop

“At the end of the day, it comes down to one thing.”

Spoiler: the thing is never one thing.

Dramatic absolutes

“This is the ONLY strategy you’ll ever need.”

Real writers avoid absolutes unless they actually mean them.

Emotional inflation

“Mind-blowing.” “Jaw-dropping.” “Insanely effective.”

Instead of describing something, it cheerleads.

Abstract nouns piling up

“Transformation. Innovation. Elevation.”

AI stacks nouns with no verbs, hoping you won’t notice the sentence forgot to happen.

Anaphora addiction

“Because you care. Because you’re committed. Because you want more.”

Humans do this sometimes. AI does it like it’s getting paid per repetition.

The moral-of-the-story wrap

“So what does this teach us?”

Feels like a middle-school book report. Humans usually let the point land without announcing the lesson.

Unnecessary reassurance

“You’ve got this.” “You’re capable of achieving great things.”

Feels like a self-help bot trying too hard.

Ending with a vague call to action

“Now go make a difference.” “Start your journey today.”

Sounds like it’s closing a Peloton class.

The “humble meta wink”

“I know you’ve heard this before, but...” “This may sound cliché, yet...”

AI tries to preempt your skepticism instead of just being authentic.

Too-clean analogies

AI writes analogies like algebra: A is to B as C is to D. Real analogies wander a bit. They show personality.

The faux casual opener

“Look.” “Here’s the deal.” “Real talk.”

These add swagger without substance.

Relentless clarity clichés

“Let me be clear.” “Here’s the bottom line.” “At its core...”

They’re filler plugs AI drops before every point.

CHEAT SHEET

Building Your Own AI Harness

If you find yourself re-explaining your style, your process, and your standards every time you open AI tools like Claude or ChatGPT, you do not have a prompt problem. You have a harness problem. A harness is the system around the model: the files that tell it how to work, the scripts that tell it what routine to run, the examples that show it what good looks like, and the matter materials that give it the facts. You might not have set up a harness yet, but it's simpler than you think. And this is how you take your AI to the next level.

The point is not to build some grand "AI operating system." The point is simpler. Stop rebuilding the same workflow in every chat. Put the stable parts in files. Put the repeatable tasks in scripts. Put today's assignment in the prompt. Then the model starts each session inside your process instead of outside it.

The Taxonomy

The cleanest way to think about this is to separate seven things.

◆ The Harness File

This is the traffic cop. It tells the model how to classify today's task, which scripts to choose from, which references to consult, and what to do when files conflict.

◆ Reference Files

These are your stable knowledge sources: your style guide, research playbook, drafting rules, quality-control checklist, and practice-area notes.

◆ Example Files

These are model paragraphs, sample analyses, and excerpts that show shape and technique. They are not rules. They are illustrations.

◆ Script Files

These are stored workflows for recurring tasks. Think of them as static prompts with a job description: "authority stress test," "statute explainer," "brief-section drafter," "counterargument builder."

◆ Matter Files

These are the live case materials: facts, timeline, issues, authority notes, record cites, draft goals, and client constraints.

◆ The Session Prompt

This is the short instruction you type today. Its job is not to restate your whole process. Its job is to assign today's task.

◆ The Improvement Log

This is where the harness gets better. Every time you discover a better question, a missing quality check, or a stronger script, you record it and update the system.

If it helps, think of it like a law office. The references are your office manuals. The scripts are your checklists. The matter files are the case file. The session prompt is the assignment email. The harness file is the partner telling the team which manual and which checklist apply today.

What a Markdown File Is, and Why to Use One

A markdown file is just a plain-text file with an .md or .markdown extension. You write normal text, add simple headings and lists, save the file, and the AI tool reads it easily.

You do not need special software. Notepad works. TextEdit works if you switch to plain-text mode. If your firm blocks new software, use whatever plain-text editor you already have. Save the file as STYLE_GUIDE.md instead of STYLE_GUIDE.docx, and you are done. And if you don't want to deal with this, you can totally use a .docx or pdf instead. But markdown just makes it easy on the AI.

And to be clear: Any major AI tool like ChatGPT or Claude can create these files for you, and that is what we suggest. Simple start a session and start building these documents with the AI.

A markdown file can be as simple as this:

Style Guide

Openings

- Start with the point.
- Explain what the statute does before citing the section number.

Case Discussion

- Do not summarize a case unless the summary proves something.
- Separate the holding from dicta.

Banned Habits

- Do not say "clearly" unless the point is actually uncontested.
- Do not hide the real issue behind a doctrinal label.

That is enough. The virtue is not fancy formatting. The virtue is that the file is clean, portable, easy to update, and easy for different AI products to read.

The Folder Structure

You do not need many files to get real value. Start with a small folder.

LAWYER_AI_HARNESS/
00_HARNESS.md

REFERENCES/
01_SYSTEM.md
02_STYLE_GUIDE.md
03_RESEARCH_PLAYBOOK.md
04_DRAFTING_PLAYBOOK.md
05_QC_CHECKLIST.md
06_PRACTICE_AREA_NOTES.md

EXAMPLES/
01_GOOD_OPENINGS.md
02_GOOD_STATUTE_EXPLANATIONS.md
03_GOOD_CASE_DISCUSSIONS.md
04_GOOD_COUNTERARGUMENTS.md

SCRIPTS/
authority_stress_test.md
statute_explainer.md
brief_section_drafter.md
counterargument_builder.md
cite_risk_audit.md
rewrite_to_real_point.md

MATTERS/
Smith_v_Jones/
00_MATTER_OVERVIEW.md
01_FACTS_TIMELINE.md
02_ISSUES.md
03_AUTHORITY_NOTES.md
04_RECORD_SUPPORT.md
05_DRAFT_GOALS.md

LOGS/
LESSONS_LEARNED.md
CHANGELOG.md

That structure does two important things. It separates the stable from the live, and it gives the model a source hierarchy. You do not want standing style rules buried inside a matter timeline, and you do not want a one-off matter note quietly rewriting your drafting playbook.

What Goes in the Reference Files

The reference files are the brain. Each one should have one job.

01_SYSTEM.md should hold your standing rules. This is where you tell the model what role it is playing, how cautious it should be, how it should handle uncertainty, what sources it should prioritize, whether it should ask questions or make labeled assumptions, and what it must never do. For a legal harness, this usually means rules like: do not invent authority, do not overstate support, separate binding from persuasive authority, distinguish between allegation-stage and merits-stage reasoning, and flag missing record support.

02_STYLE_GUIDE.md should hold the things you are always fixing anyway. How do you want section openings to work? How long should paragraphs run? Do you like headings that make an argument or headings that describe a topic? What phrases do you want to avoid? How should the model write about statutes, cases, standards, and counterarguments? If you have a recognizable house style, it belongs here.

03_RESEARCH_PLAYBOOK.md is where the real leverage starts. This file should not say “find cases.” It should contain your better questions. For statutory work, should the model check neighboring provisions, omitted language, parallel statutes, amendment history, practical consequences, and hypotheticals? For case law, should it identify limiting language, emphasized facts, silent facts, weak links in a line of cases, and what the opinion assumed without saying? This is the file that keeps the tool from doing thin, first-layer research.

04_DRAFTING_PLAYBOOK.md should explain how research becomes prose. How do you want the model to move from authority to argument? How should it explain a rule without hiding behind the doctrinal label? How should it write the real point rather than the section number? How should it make a concession? How should it turn a line of cases into a principle rather than a string cite?

05_QC_CHECKLIST.md should catch the failures you see over and over. Did the draft overclaim? Did it rely on dicta as if it were a holding? Did it cite a statute number without explaining what the provision actually does? Did it quietly assume facts not in the record? Did it talk about the law in abstractions without translating the law into operation?

06_PRACTICE_AREA_NOTES.md is optional but useful if you work in a repeat area. Put there the recurring traps, the common arguments, the annoying local doctrines, the court-specific quirks, or the client-specific constraints that show up again and again.

The example files are different. A good example file is not a dump of nice paragraphs. It labels the lesson. “This example is useful because it explains a statute by operation rather than by section number.” “This example is useful because it distinguishes a case without spending a page summarizing it.” Examples work best when the model knows what it is supposed to learn from them.

What Goes in the Script Files

A script is a stored workflow for a recurring task. It is more specific than a system rule and more stable than a one-off prompt. If you keep typing the same kind of instruction, it belongs in a script.

Every good script should answer six questions:

1. What is this script for?
2. When should it be used?
3. What inputs does it need?
4. What files should it consult first?
5. What steps should it run?
6. What should the output look like?

That is enough structure to make a script reusable.

Here is a simple example:

```
# Authority Stress Test
```

```
## Use when
```

You have a draft or outline that relies on authority and you want to test whether the authority is being used well.

```
## Inputs
```

- Current draft or outline
- Matter authority notes
- Research playbook
- QC checklist

```
## Read first
```

1. MATTERS/.../03_AUTHORITY_NOTES.md
2. REFERENCES/03_RESEARCH_PLAYBOOK.md
3. REFERENCES/05_QC_CHECKLIST.md
4. Current draft

Steps

1. Identify the three authorities doing the most work.
2. For each one, state the actual holding, the key facts, and any limiting language.
3. Identify one place where the draft may be overreading or underusing the authority.
4. Suggest one hypothetical that pressure-tests the opponent's reading.
5. Rewrite the weakest authority paragraph.

Output

- Biggest authority risk
- Authority-by-authority notes
- Revised paragraph
- Questions for lawyer review

That is a script. It is a stored routine, not a giant general prompt.

A good lawyer harness usually ends up with scripts in four buckets.

Research scripts: issue mapper, statutory interpretation sweep, analogous-law finder, line-of-cases audit, authority stress test, bad-case rescue.

Drafting scripts: section drafter, case explainer, statute-to-operation rewrite, counterargument builder, transition improver, conclusion writer.

Revision scripts: rewrite to real point, cut throat-clearing, tighten headings, remove overstatement, improve concessions, reduce summary and increase proof.

Quality scripts: holding-vs-dicta check, record-support audit, cite-risk audit, stage-of-case audit, unsupported-assumption audit.

The key is to keep scripts narrow enough that they do one thing well. If a script sounds like it could govern every task, it is probably not a script. It is probably part of your SYSTEM.md.

The Master Harness File

Yes, you probably do want one master file that routes the work. This is the file that tells the model how to choose among scripts and references each time.

I would call it 00_HARNESS.md, and I would make it short and procedural.

Master Harness

You are working inside a legal AI harness.

When the user gives a task:

1. Classify the task as one of these:

- research
- drafting
- revision
- quality control
- strategy analysis
- client communication
- training/example generation

2. Choose:

- one primary script from /SCRIPTS
- one optional secondary script if needed

3. Choose the most relevant files from:

- /MATTERS for case-specific facts and issues
- /REFERENCES for standing guidance
- /EXAMPLES for illustration

4. State which files you selected and why.

5. Use this priority order if files conflict:

1. matter files
2. QC checklist and system rules
3. specific script
4. style guide
5. examples

6. Do not invent authority, facts, or record support.

7. Label assumptions.

8. Separate confirmed points from inferences.

9. Give the requested output in the requested format.

10. At the end, propose any permanent improvements for LESSONS_LEARNED.md.

If the user's prompt is vague, do not restate the whole harness. Infer the likely task type and proceed, while flagging major uncertainties.

That file does not do the work itself. It assigns the work. In practice, it keeps the model from rummaging through files randomly or treating a sample paragraph like a governing rule.

The Session Prompt

Once the harness exists, your session prompt should get shorter. It is no longer responsible for carrying the whole system.

A good session prompt usually needs five things: the task, the matter, the audience, the deliverable, and any hard constraints. That is it.

Something like this:

Use 00_HARNESS.md and the Smith_v_Jones matter files.

Task: draft the opening analysis section on whether the notice statute requires the agency to include the hearing date in the initial notice.

Audience: federal district judge.

Deliverables:

1. brief issue map
2. 700-word draft
3. list of the two biggest authority risks

Constraints:

- write in firm style
- explain what the statute does before citing the section number
- do not overclaim if the case law is mixed

Notice what the prompt is not doing. It is not restating your entire style guide. It is not reciting your QC checklist. It is not re-teaching the model how to analyze authority. Those instructions already live in the harness.

How to Set This up in Claude

Claude is well-suited to this because Projects lets you upload documents, text files, or code snippets into a project knowledge base and add project instructions that apply across chats in that project. Anthropic also says that context is not shared across chats within a project unless the information is in the project knowledge base, which is exactly why a file-based harness helps. Paid Claude plans also get caching benefits from project knowledge bases, so reused project material counts more efficiently over time.

The clean setup in Claude is:

1. Create a Project.
2. Paste the content of 00_HARNESS.md into Project Instructions.
3. Upload your reference files, example files, script files, and matter files into Project Knowledge.
4. Start each chat with a short session prompt that points Claude to the matter and the task.

That is usually enough. If the project gets crowded, upload only the standing files plus the current matter folder. You do not need every old matter living there forever. You could also use Claude CoWork now and simply put all these files into a folder and then give Claude access to it!

How to Set This up in ChatGPT

ChatGPT gives you two useful paths. Projects are the closest match to Claude Projects: files, chats, and project instructions live together, and OpenAI explicitly describes Projects as repeatable workspaces for recurring tasks. Project instructions apply only inside the project and override global custom instructions. If you build a custom GPT instead, OpenAI's guidance is to put behavior, tone, and workflow in the Instructions field, while uploaded Knowledge files should serve as reference material. Clear, text-forward files work better than complex layouts.

So the practical setup in ChatGPT is:

- **If you use Projects:** put 00_HARNESS.md into Project Instructions, upload the reference, example, script, and matter files, and run short session prompts inside that project.
- **If you use a custom GPT:** put the harness rules, style rules, and workflow rules in Instructions; upload the reference and example files as Knowledge.
- **If you only have a plain chat:** upload the relevant files for that session and tell ChatGPT the file hierarchy in the opening prompt.

One more small point matters. OpenAI says Custom Instructions apply across all chats. That makes them the wrong place for matter-specific or workflow-heavy instructions. Use them for broad, personal defaults only. Put the serious harness inside a Project or GPT.

If the Tool Is Limited, the Harness Still Works

Some firms will not allow Projects. Some will not allow custom assistants. Some will only allow a locked-down chat with file upload.

The harness still works.

Keep the same folder on your machine or in whatever approved repository your firm allows. Then, for each session, upload:

- the harness file,
- the relevant reference files,
- one or two scripts,
- the current matter files.

Then start with a short session prompt that names the task and the file order. The process is a little more manual, but the structure still saves you from rebuilding everything every time.

How to Make the Harness Improve

The right version of “self-improving” is not silent auto-editing. It is proposal, review, then update.

At the end of a useful session, run a short update prompt:

Review this session and propose updates to:

- REFERENCES/03_RESEARCH_PLAYBOOK.md
- REFERENCES/04_DRAFTING_PLAYBOOK.md
- REFERENCES/05_QC_CHECKLIST.md
- SCRIPTS/

For each proposed change:

1. give the exact text to add or revise
2. explain why it helps
3. say whether it is permanent guidance or matter-specific
4. identify the file where it belongs

Do not rewrite the files directly.

You could, in theory, have your AI autoimprove itself. But to give you more control and ensure your work isn't messed up, approving changes is probably the best bet. Over time, this is how the harness gets sharper. The style guide gets clearer. The research questions get better. The QC file catches more of the model's recurring mistakes. The script library grows only where it is actually useful.

Before you upload client material, check the product tier and your organization's rules. Also, do not assume every uploaded document is being read visually. OpenAI's file-upload guidance says that, outside Enterprise PDF visual retrieval, ChatGPT generally extracts digital text from document uploads and discards embedded images. If your PDF contains a critical chart, handwritten markup, or scanned visual that matters, summarize it in text or upload the relevant image separately.

CHEAT SHEET

Key GenAI Terms to Know

A

API (Application Programming Interface)

A set of rules enabling one software application to interact with another. OpenAI provides an API to use GenAI in various applications, such as chatbots, translation services, and content generation.

AGI (Artificial General Intelligence)

A hypothetical AI system capable of performing any intellectual task a human can do, possessing understanding and learning capacities similar to human intelligence.

Attention Mechanism

A technique in deep learning models, like the Transformer architecture, enabling the model to focus on different input parts when making predictions, thus capturing long-range dependencies and context more effectively.

Autoregression

A property of language models where the next token's prediction depends on the model-generated previous tokens, rather than just a single question input.

B

Beam Search

An NLP search algorithm that finds the most likely token sequence generated by a language model. It tracks a fixed number of the most likely sequences (called "beams") at each time step and explores multiple possibilities, unlike Greedy Search, which chooses the single most likely token.

C

Chatbot

A software application simulating human conversation through text or voice, using NLP techniques and models like GenAI.

Context

Providing background information for GenAI prompts.

Corpus

A large, structured set of texts used for training language models or conducting linguistic analysis.

COT Prompting

Chain-of-thought guidance, enabling GenAI to carry out reasoned steps.

D

Dataset

A data collection used for training, validating, and testing an AI model. In GenAI, it consists of a diverse range of texts from books, articles, and websites.

F

Fine-tuning

Adapting a pre-trained model for a specific task by training it on a set of your data.

Few-shot

Providing GenAI with examples to help it understand what you want, how you want responses formatted, and more.

G

Greedy Search

An NLP search algorithm finding the most likely token sequence generated by a language model, selecting the most likely token at each time step. This can lead to poor responses, especially when several reasoning steps are needed to reach the right answer.

GenAI

A series of AI language models developed by OpenAI.

H

Hallucination

Refers to an AI incorrectly believing something is true.

I

Instructions

Specific tasks given to GenAI, such as "write" or "order alphabetically."

L

Language Model

An AI model trained to generate text based on patterns learned from a large corpus of text data.

M

Machine Learning

An AI subfield focusing on developing algorithms enabling computers to learn from and make predictions or decisions based on data.

N

Neural Network

A computing model inspired by biological neurons' structure and function, used in deep learning to process and analyze data.

NLP (Natural Language Processing)

An AI subfield focused on computers and humans interacting using natural language.

P

POS (Part-of-Speech Tagging)

An NLP task involving labeling each word in a sentence with its grammatical role, such as noun, verb, or adjective. This is part of why GenAI is adept at analyzing text for style and grammar improvements.

Prompts/Prompt Engineering

Constructing GenAI guidance to achieve the best results for your goals.

S

Stopwords

Common words in a language (e.g., "the", "and", "is") often removed from text during preprocessing to reduce noise and improve computational efficiency.

T

Temperature

How determinative GenAI is. For a single answer, set the temperature low; for GenAI to explore multiple possible answers, turn the temperature up.

Token

An NLP basic unit of meaning, representing a single word, punctuation mark, or another text piece. OpenAI's paid plans charge per token used in your prompt and GenAI's response.

Transformer Architecture

The neural network architecture underlying GenAI and other state-of-the-art language models, characterized by its use of self-attention mechanisms, which allow the model to weigh the importance of different words in a sentence when making predictions.

Training

The process of adjusting a model's parameters to learn patterns in data. In GenAI, this involves training the model on a large dataset of text, teaching it to generate coherent and contextually relevant text based on input prompts.

Z

Zero-shot Learning

The ability of a language model to perform a task it has not seen before, based on its understanding of the language and the knowledge it has learned from pre-training. GenAI can do quite a lot of zero-shot work, which is part of what makes it so powerful.