

Legal Writing With AI

Practical Benefits, Ethical Use, and Best Practices

Ross Guberman

Founder of BriefCatch

Author of Point Made: How to Write Like the Nation's Top Advocates

LLM Comparison for Legal Use Cases

The following table compares major AI models across key legal functions. Ratings reflect capabilities as of early 2026.

Use Case	ChatGPT	Claude	Gemini	Perplexity	NotebookLM
Legal Research & Analysis	Strong	Strong	Strong	Good	Limited
Document Analysis	Strong	Strong	Strong	Limited	Excellent
Legal Writing & Drafting	Good	Excellent	Good	Limited	Good
Citation Verification	Unreliable	Unreliable	Unreliable	Fair	N/A
Strategy & Argument	Strong	Strong	Strong	Good	Limited
Contract & Deal Analysis	Strong	Excellent	Strong	Limited	Good
Content Creation	Strong	Strong	Strong	Good	Excellent
Audio & Video Overviews	Limited	Limited	Limited	No	Excellent

⚠ IMPORTANT NOTE: Models tested: ChatGPT-4o, Claude Sonnet 4, Gemini 2.5, Perplexity Pro, NotebookLM. General chatbots show materially higher hallucination rates on legal queries than legal-tuned tools. Always verify AI-generated legal content. Individual experiences may vary.

1. PRACTICAL WAYS TO USE AI

Here are the highest-value use cases—and the prompts to get practitioner-grade results. Every prompt follows the Expert Prompt Formula:

The Expert Prompt Formula **ROLE + CONTEXT + TASK + FORMAT + CONSTRAINTS** Role: Who should the AI be? (“A seasoned M&A partner,” “A skeptical appellate judge”) Context: What’s the situation? (case posture, deal stage, audience, jurisdiction) Task: What exactly do you need? (analyze, draft, compare, critique, redline) Format: How should the output look? (table, numbered list, tracked changes, bullet points) Constraints: What guardrails apply? (no hallucinations, Bluebook citations, mark uncertainties)

A. Chained Analysis: Multi-Step Workflows

The most powerful AI results come from chaining prompts: use the output of one step as the input for the next.

Litigation

The Full-Record Case Autopsy (3-Step Chain)

Step 1 — The Fact Extractor: I’m uploading [describe documents]. Extract every material fact and organize them into a chronological master timeline. For each fact, note: (1) the exact source document and page, (2) which party the fact helps, (3) whether any other document in the set contradicts it. Flag any gaps in the timeline where events are referenced but never documented.

Step 2 — The Vulnerability Map: Using the master timeline you just created, act as opposing counsel. Identify the five most damaging factual patterns for my client. For each: (1) state the pattern in one sentence, (2) identify every document that supports it, (3) explain how opposing counsel will frame it in a summary judgment motion or closing argument, and (4) rate the severity from 1 to 5. Then identify the three factual patterns most helpful to my client that I may be under-leveraging.

Step 3 — The Strategic Brief Skeleton: Now act as lead trial counsel for [my side]. Using the vulnerability map, draft a detailed outline for a [motion type] that: (1) leads with our three strongest factual themes, (2) preemptively neutralizes the top two opposing patterns you identified, (3) sequences the arguments so each section builds on the last, and (4) identifies the two record citations most critical to win at each stage.

Transactional

The Contract Negotiation Engine (3-Step Chain)

Step 1 — The Risk-Allocation X-Ray: Act as a 20-year M&A partner. I’m uploading a [agreement type]. Map every risk-allocation mechanism into a table with six columns: (1) Section reference, (2) Risk being allocated, (3) Party bearing the risk, (4) Trigger event, (5) Cap or limitation (if any), (6) Whether this is market-standard, buyer-favorable, or seller-favorable for a [deal size] [industry] transaction. At the bottom, calculate the aggregate uncapped exposure for each party.

Step 2 — The Negotiation Playbook: Using the risk-allocation table, I represent [Buyer/Seller]. Rank the non-market provisions by commercial impact. For the top five, generate a negotiation strategy for each:

(1) our ideal position, (2) our fallback position, (3) the concession we can offer in exchange, (4) the business rationale that makes our ask persuasive, and (5) a one-sentence response to the most likely pushback. Step 3 — The Markup Generator: For each of the top five provisions, draft two versions of redline language: (a) an aggressive version that maximizes our protection, and (b) a compromise version we can live with. For each redline, include a brief explanatory comment in brackets that I can paste into the margin of the agreement.

B. Summarize and Analyze: Deep-Dives

The Disclosure Schedule Auditor

I'm uploading a set of disclosure schedules with the corresponding agreement. Cross-reference them against the representations and warranties. Produce a report that identifies: (1) every representation that references a disclosure schedule that either doesn't exist or is empty, (2) every disclosure that appears to qualify a representation so broadly it effectively negates the rep, (3) any disclosed items that should trigger additional due diligence, and (4) inconsistencies between schedules. Rank findings by deal risk.

The Defined-Terms Auditor

Review the attached agreement and produce: (1) A complete index of every defined term, its section reference, and its definition. (2) Flag every defined term that is used but never actually defined. (3) Flag every defined term that is defined but never used. (4) Identify circular definitions. (5) Identify definitions that contain substantive obligations that should be in the operative provisions, not buried in the definitions section. (6) Flag any defined term whose definition materially changes meaning depending on which section it appears in.

The Summary Judgment Predictor

Act as a federal magistrate judge assigned to write a report and recommendation on the pending summary judgment motion. I'm uploading the motion, opposition, and reply. For each disputed issue: (1) state the legal standard, (2) identify what the movant must show, (3) assess whether a genuine dispute of material fact exists, and (4) state your tentative ruling with reasoning. Flag any issue where the result turns on credibility. Conclude with an overall recommendation and identify the single factual record cite most likely to determine the outcome.

C. War-Game: Adversarial Stress Testing

These prompts go beyond simple “devil’s advocate” by forcing the AI to adopt specific hostile personas and produce outputs you can directly use.

The Three-Judge Gauntlet

I’m uploading my appellate brief. Simulate a three-judge panel where each judge adopts a different posture: Judge 1 is a strict textualist who will press you on every statutory interpretation claim. Judge 2 is a pragmatist focused on real-world consequences and administrability. Judge 3 is skeptical of our standing and jurisdiction. Each judge should ask their two hardest questions. After all six questions, switch roles: act as my appellate advocate and draft a concise, persuasive answer to each question, noting where the brief should be strengthened.

The Hostile Witness Simulator

Act as a hostile corporate witness being deposed. I’ll play the deposing attorney. You are [role, e.g., the CFO of the defendant corporation]. You have been coached by aggressive defense counsel. Your goals: (1) give the narrowest possible answers, (2) claim lack of memory strategically, (3) redirect to documents when cornered, and (4) never volunteer information. I’ll ask questions one at a time. After every five questions, break character and give me candid feedback: Which questions were effective? Which let me off the hook? What follow-up should I have asked?

The Indemnity Stress Test

I’m uploading the indemnification article from a [agreement type]. Stress-test it against these five scenarios: (1) A third-party lawsuit filed two years post-closing, (2) A pre-closing liability discovered at month 18, (3) A regulatory investigation that doesn’t result in a formal enforcement action, (4) A claim that falls below the basket but aggregates with other claims to exceed it, (5) A dispute over whether the indemnifying party was given proper notice. For each scenario, walk through the operative language step by step and state whether the indemnified party is covered, partially covered, or exposed—and identify the exact language that controls the outcome.

The Dawn-Raid Simulator

Role-play as a team of [SEC / DOJ / FTC] investigators who have just received a whistleblower complaint about [describe the conduct]. You have subpoena power. Generate: (1) the five categories of documents you would subpoena first and why, (2) the three employees you would interview first and the key questions for each, (3) the two most likely theories of liability you would pursue, and (4) the factual predicates that would make you escalate from investigation to enforcement action. Then switch perspective: as outside counsel for the company, what proactive steps should we take this week?

D. Transform: Translate, Format, and Polish

Convert complex analysis into client-ready deliverables.

The Board-Ready Deal Summary

I'm uploading a [length]-page [agreement type]. Produce a board-ready executive summary structured as follows: (1) Transaction Overview (parties, structure, consideration, timeline—four sentences max), (2) Key Commercial Terms table (term, our obligation, counterparty obligation, risk level rated Green/Yellow/Red), (3) Top Five Risks with one-paragraph analysis of each, (4) Material Conditions to Closing as a checklist, (5) Recommended Board Actions in numbered sequence. Write for directors who are sophisticated businesspeople but not lawyers. Use plain English. Bold every defined term on first use.

The Jury-Ready Narrative

I'm uploading a [technical/scientific/financial] expert report. Translate the core findings into: (1) a three-paragraph narrative that a non-specialist juror could follow, using analogies and concrete examples, (2) a single-page visual outline showing the expert's logical chain from data to conclusion (describe the visual as boxes and arrows I can build), and (3) the three most powerful soundbites from the report that I should quote in closing argument. Preserve the scientific accuracy but eliminate all jargon.

The Client-Alert Generator

A [describe legal development: new regulation, Supreme Court decision, enforcement action] just happened. Draft a client alert in three sections: (1) What Happened (three sentences, plain English), (2) What It Means for You (practical impact organized by client type—public companies, private companies, financial institutions, individuals), (3) What to Do Now (numbered action items, most urgent first). Tone: authoritative but accessible. Length: one page. End with a one-sentence call to action directing readers to contact us.

E. Precision Legal Writing

These prompts target the specific writing skills that separate competent briefs from exceptional ones. Each is grounded in techniques from Point Made.

The Lede Lab

Act as a legal writing coach who has analyzed the openings of 500 Supreme Court merits briefs. I'm pasting the Introduction from my brief below. Generate five alternative opening paragraphs, each using a different technique: (1) The Story Lede: open with a single concrete, human moment from the record, (2) The Thesis Lede: state our winning theory in one declarative sentence, then unfold it, (3) The Contrast Lede: juxtapose what the law requires with what the opposing party did, (4) The Stakes Lede: frame the broader consequences of ruling against us, (5) The Question Lede: open with the single question the court must answer, framed to favor our side. After all five, recommend which lede works best for this case and explain why.

The Cohesion Audit

Act as an expert legal editor. Review the attached brief one paragraph at a time. For each paragraph, assess: (1) Does the first sentence state the paragraph's point, or does it bury the conclusion? (2) Does every sentence in the paragraph serve that point, or are there stowaways? (3) Does the last sentence transition to the next paragraph's topic? Produce a table: Paragraph Number | Current First Sentence | Recommended First Sentence | Sentences to Cut or Move | Transition Fix. Process the first ten paragraphs of the Argument section.

The Legalese Exorcist

I'm pasting a section of my brief. Apply these edits in sequence and show me the before and after for each change: (1) Convert every nominalization back to its root verb ('made a determination' → 'determined'). (2) Rewrite every passive sentence to assign a clear actor. (3) Break any sentence over 35 words into two. (4) Eliminate every 'it is [adjective] that' construction. (5) Replace every 'the fact that' with a tighter phrasing. Present the final clean version at the end, then count how many words you saved and by what percentage.

The Rhetoric Scorecard

Act as a debate coach who specializes in appellate advocacy. I'm uploading my brief's Argument section. Score each major argument (i.e., each section heading) on five dimensions, 1–10 scale: (1) Clarity of Rule Statement, (2) Strength of Record Support, (3) Effectiveness of Case Analogies, (4) Preemptive Rebuttal of Counterarguments, (5) Persuasive Framing and Word Choice. For any dimension scoring below 7, provide a specific, concrete suggestion to improve it. End with an overall Assessment and the three highest-ROI revisions I could make in two hours.

The Authority Weaver

Review the attached section and identify every instance where a case citation appears in a way that buries the authority or disrupts the sentence flow. For each, rewrite the sentence using one of these techniques: (1) Lead with the court's action ('The Seventh Circuit held...'), (2) Integrate the citation parenthetically so the sentence reads as a clean assertion, (3) Pair the authority with a punchy parenthetical that does real argumentative work—not just a quote but a framing of why the case helps us. Show before and after for each rewrite.

F. Meta-Prompts: Controlling AI Behavior

Prepend these to any task to get more reliable, structured, and trustworthy outputs from any AI tool.

The Structured Output Enforcer

For the task I'm about to give you, follow these output rules strictly: (1) Begin every response with a Confidence Assessment: state High, Medium, or Low confidence for each major claim, and explain what would change your confidence level. (2) Mark any statement that relies on your training data rather than the documents I've provided with [TRAINING DATA]. (3) Mark any statement you're generating by inference rather than direct textual support with [INFERENCE]. (4) End with an Uncertainty Register listing everything you're unsure about, ranked by how much it matters to the analysis.

The Self-Critique Cycle

After you complete the task I'm about to give you, do not stop. Immediately: (1) Re-read your output and identify the three weakest points—the places where a skeptical partner would push back hardest. (2) For each weak point, either strengthen it with better reasoning or acknowledge the limitation explicitly. (3) Re-read the improved version and identify anything you stated with more confidence than the evidence warrants. Dial it back. (4) Produce the final revised version with a brief changelog showing what you improved and why.

The Five-Lens Analysis

Analyze [issue/document/strategy] from five distinct professional perspectives, and produce a separate assessment from each: (1) The Litigator: What are the dispute risks? What would you want different if this ends up in court? (2) The Regulatory Lawyer: What compliance exposures exist? What would a regulator flag? (3) The Business Executive: Does this make commercial sense? What flexibility does it preserve or surrender? (4) The Opposing Counsel: Where would you attack if you represented the other side? (5) The Appellate Judge: Is the legal reasoning sound? Would you affirm or reverse, and why? Conclude with a synthesis identifying the areas where all five perspectives converge—those are the true priorities.

The Hidden Assumption Audit

Review the analysis you just provided. Identify every assumption you made—explicitly or implicitly—that I did not instruct you to make. List each assumption, explain why you made it, and state how the analysis would change if the assumption were wrong. Group the assumptions into three categories: (1) Assumptions that almost certainly hold, (2) Assumptions that are reasonable but could go either way, and (3) Assumptions that are aggressive or optimistic. For category 3, redo the relevant portion of the analysis under the opposite assumption.

2. WHERE AI FAILS in Legal Writing

Before learning how to use AI effectively, lawyers must understand where it commonly goes wrong. General-purpose AI tools are confidently wrong at an alarming rate when it comes to legal content.

Hallucinations and Fabricated Citations

General-purpose AI tools (ChatGPT, Claude, Gemini) frequently fabricate case names, invent statute numbers, and generate holdings that do not exist. Several lawyers have already faced sanctions for filing briefs with AI-generated fake citations, most notably in *Mata v. Avianca, Inc.* (S.D.N.Y. 2023).

⚠️ CRITICAL WARNING: General AI tools will fabricate cases, statutes, and holdings with complete confidence. These tools do not access legal databases—they predict plausible-sounding text. Never rely on AI-generated citations without rigorous verification through Westlaw, Lexis, or court databases. Even legal-specific AI tools can produce errors.

Common Misstatements

- **Incorrect legal standards:** AI may apply the wrong standard of review, conflate state and federal tests, or misstate burden-of-proof rules.
- **Outdated law:** Training data has a cutoff. AI may cite overruled cases or superseded statutes without flagging the change.
- **Jurisdiction confusion:** AI frequently blends rules from different jurisdictions without distinguishing them.
- **Overconfident analysis:** AI presents nuanced, unsettled questions as though they have clear answers, suppressing uncertainty.
- **Fabricated quotations:** AI may attribute language to a judge, statute, or treatise that was never written.

Red Flags to Watch For

- Citations with incorrect reporters or unusual case name formats.
- Statutes with section numbers that don't exist in the actual code.
- Holdings that perfectly support your position without any nuance.
- References to cases from jurisdictions where they don't exist.
- Inconsistent dates or procedural history that doesn't make sense.

3. ETHICS AND PROFESSIONAL RESPONSIBILITY

AI tools raise significant professional responsibility considerations. Lawyers must understand their obligations before integrating AI into their practice.

ABA Formal Opinion 512

The ABA's 2024 formal opinion on generative AI addresses three core obligations:

- **Duty of Competence (Rule 1.1):** Lawyers must understand AI's capabilities and limitations well enough to use it competently. Blind reliance on AI output is incompetent.
- **Duty of Confidentiality (Rule 1.6):** Inputting client information into AI tools may constitute disclosure. Lawyers must evaluate whether a tool's data practices protect confidentiality.
- **Duty of Supervision (Rules 5.1, 5.3):** Lawyers who direct others to use AI, or who use AI-generated work product from associates or staff, must supervise the output.

Candor to the Tribunal

Filing AI-generated content that contains fabricated citations or misstatements of law violates Rule 3.3. The lawyer, not the AI tool, bears full responsibility for every assertion in a filing.

Emerging Court Rules

Many federal and state courts now require disclosure of AI use in filings. Lawyers should check local rules and standing orders before using AI in any case.

Best Practices for Ethical AI Use

- **Verify everything:** Treat AI output as a first draft from an unreliable associate. Check every citation, holding, and legal standard.
- **Know your tools:** Understand whether your AI tool stores inputs, trains on your data, or shares information with third parties.
- **Use secure platforms:** Enterprise tools with data isolation (like BriefCatch) are safer than consumer chatbots for client work.
- **Document your process:** Maintain records of how AI was used, what was verified, and by whom.
- **Start with non-confidential tasks:** Build familiarity before using AI on sensitive client matters.

4. WHY EDITING-BASED AI IS SAFER

Not all AI tools carry the same risk profile. Understanding the difference between drafting-based AI and editing-based AI is critical for responsible adoption.

Drafting-Based AI: Higher Risk

Tools like ChatGPT and Claude generate text from scratch. This means:

- Hallucination risk is inherent: The model predicts plausible next words, which means fabricated citations and misstatements are a feature of the architecture, not a bug.
- No anchor text: Without your draft as a starting point, the AI is filling a blank page—and filling it with whatever sounds right.
- Confidentiality exposure: Generating a full draft often requires inputting detailed client facts, increasing disclosure risk.
- Voice and judgment are lost: AI-generated drafts lack your analytical voice, your strategic choices, and your knowledge of the case.

Editing-Based AI: Lower Risk

Tools like BriefCatch work on your existing draft. This changes the risk calculus:

- You control the substance: The legal analysis, citations, and arguments are yours. The AI focuses on clarity, concision, and persuasion.
- No fabrication risk: An editing tool that suggests “change ‘made a determination’ to ‘decided’” cannot hallucinate a case citation.
- Minimal confidentiality concern: BriefCatch processes text locally within Microsoft Word. Your drafts don’t travel to external servers.
- Preserves your voice: The AI improves your writing without replacing it. You stay in the driver’s seat.

The Practical Takeaway

Use drafting-based AI for brainstorming, research, and strategy—but always with verification. Use editing-based AI for the final mile: tightening prose, catching legalese, improving persuasion, and verifying citations. This combination gives you the best of both worlds while managing risk.

BriefCatch: Editing-Based AI Built for Legal Writing Deterministic rules catch legalese, wordiness, and passive voice with no hallucination risk. AI-powered holistic editing analyzes your entire document for structure, flow, and persuasion. Citation Checker verifies citations against real legal databases—not AI guesswork. Works inside Microsoft Word. Your text stays in your workflow. Learn more at BriefCatch.com.

5. HOW FIRMS ARE INTEGRATING AI

Leading law firms and legal departments are adopting AI thoughtfully—balancing the productivity gains against the professional and reputational risks.

Common Integration Strategies

- Start with low-risk use cases: Firms typically begin with internal brainstorming, research summaries, and first drafts of non-filing documents before expanding to client-facing work.
- Layer tools by risk: Use general AI (ChatGPT, Claude) for ideation and strategy. Use specialized tools (BriefCatch, Westlaw AI, Lexis+ AI) for citation-dependent and client-facing work.
- Create AI-use policies: Most AmLaw 200 firms now have written guidelines governing which tools are approved, what data can be inputted, and what disclosures are required.
- Train across all levels: Effective firms train partners, associates, and staff together—not just the early adopters.

Managing Risk While Maintaining Standards

- Human-in-the-loop: Every AI output should be reviewed by a lawyer who takes ownership of its accuracy and quality.
- Verification workflows: Build citation-checking and fact-verification into your workflow before any filing or client delivery.
- Secure platforms: Use enterprise-grade tools with data isolation—not consumer-facing free tiers.
- Track ROI: Measure time savings and quality improvements so you can scale what works and abandon what doesn't.

The BriefCatch Approach

BriefCatch serves 50+ AmLaw 200 firms and federal courts by combining deterministic, rule-based editing (zero hallucination risk) with AI-powered holistic analysis and citation verification. This layered approach lets firms adopt AI for legal writing with confidence.

- BriefCatch Classic: Deterministic rules catch legalese, wordiness, passive voice, and grammar issues—predictable, explainable, and hallucination-free.
- BriefCatch Next: AI-powered holistic editing analyzes entire documents for structure, flow, and persuasion alongside the deterministic rules.
- Citation Checker: Verifies citations against real legal databases. No AI guesswork.
- BriefChat: A 24/7 AI legal writing advisor grounded in Point Made methodology.

Other Specialized Legal AI Tools

- Clearbrief: Confirm the accuracy of facts and statements.
- Westlaw Precision: AI-powered research within the Westlaw ecosystem.
- Lexis+ AI: Grounded with linked sources; verify before filing.
- Harvey AI: Enterprise-grade AI for law firms.

Conclusion: Start with One Prompt This Week

AI is transforming legal practice—but it requires informed, disciplined use. Understand where it fails. Know your ethical obligations. Use editing-based AI for the final mile and drafting-based AI for brainstorming and strategy. Build verification into your workflow.

Pick one prompt from this guide and try it on a non-confidential task this week. Then gradually expand your AI toolkit.

⚠️ REMEMBER AI is a powerful assistant, not a replacement for legal judgment. Always verify AI-generated content. Editing-based AI is safer than drafting-based AI for client-facing work. Maintain your professional responsibility to your clients and the courts.

• • •

For more information:

BriefCatch.com

Ross Guberman, Founder