PROGRAM MATERIALS
Program #36171
September 21, 2026

Thinking Like a Lawyer: Prompting Ethically, Securely and Safely

Copyright ©2026 by

- Jeffrey Allen, Esq. - Graves & Allen Attorneys and Counselors at Law
- Ashley Hallene - Demeter Land Development

All Rights Reserved.
Licensed to Celesq®, Inc.

Celesq® AttorneysEd Center
www.celesq.com

5301 North Federal Highway, Suite 150, Boca Raton, FL 33487
Phone 561-241-1919

Thinking Like a Lawyer, Prompting Like a Pro

Prompting Ethically, Securely, and Safely

Jeffrey Allen & Ashley Hallene

Copyright Jeffrey Allen and Ashley Hallene. All rights reserved.

1

Presenters



JEFFREY ALLEN
Graves & Allen
Jallenlawtek@aol.com
(510) 610-8777



ASHLEY HALLENE
ahallene@gmail.com
(281) 433-6220

Copyright Jeffrey Allen and Ashley Hallene. All rights reserved.

2

AI Is a Tool, Learn to Use It

CAPABILITY REQUIRES CONTROL

AI can assist.
AI can analyze.
AI can draft.
AI can make mistakes.

THE LAWYER REMAINS RESPONSIBLE.



Copyright Jeffrey Allen and Ashley Hallene. All rights reserved.

3

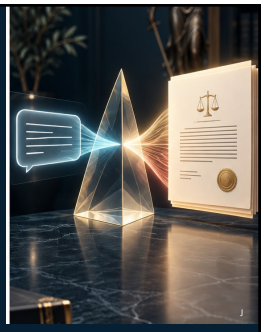
Why Prompting Matters

THE PROMPT SHAPES THE PRODUCT

AI responds to what you ask—not necessarily what you meant.

BETTER PROMPTS PRODUCE

- More relevant answers
- Better focused answers
- More useful work product



Copyright Jeffrey Allen and Ashley Hallene. All rights reserved.

4

Generative AI

Generative AI can:

- Analyze
- Summarize
- Draft
- Organize
- Brainstorm

FLUENCY ≠ ACCURACY



Copyright Jeffrey Allen and Ashley Hallene. All rights reserved.

5

Confidence Is Not Reliability

AI may sound certain while it:

- Gives a wrong answer
- Omits important information
- Makes unwarranted assumptions
- Invents information

CONFIDENCE DOES NOT PROVE RELIABILITY.



Copyright Jeffrey Allen and Ashley Hallene. All rights reserved.

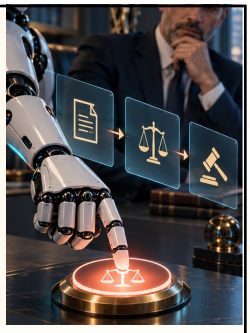
6

Agentic AI Changes the Equation

GENERATIVE AI
"Here is what I suggest."

AGENTIC AI
"I did it."

Errors can become actions, not merely bad answers.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 6

7

Keep a Human in the Loop

Review before the consequences become difficult to reverse:

- Filing or sending
- Purchasing or deleting
- Transferring information
- Committing the client



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 7

8



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 8

9



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 9

10



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 10

11



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 11

12



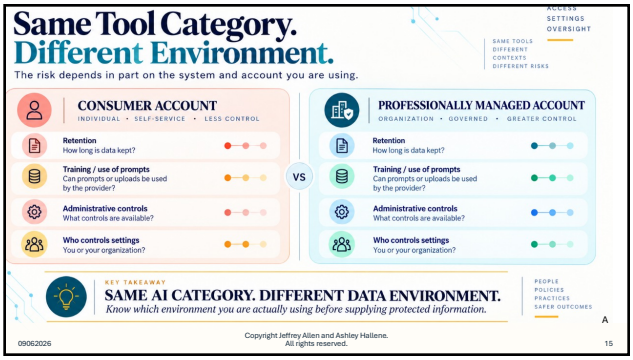
13



14



15



16



17



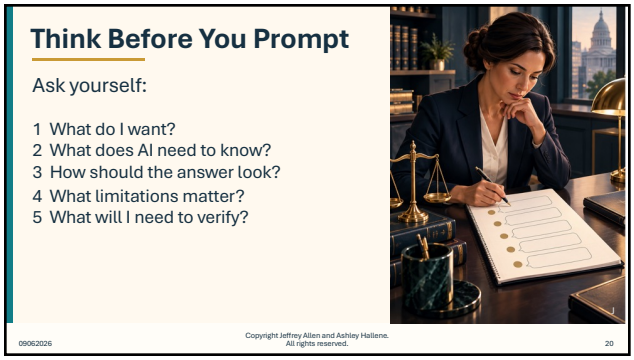
18



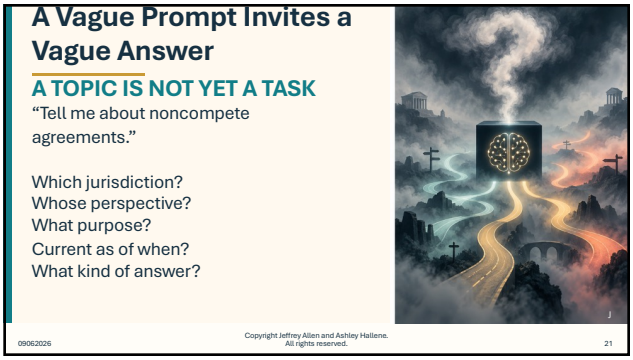
19



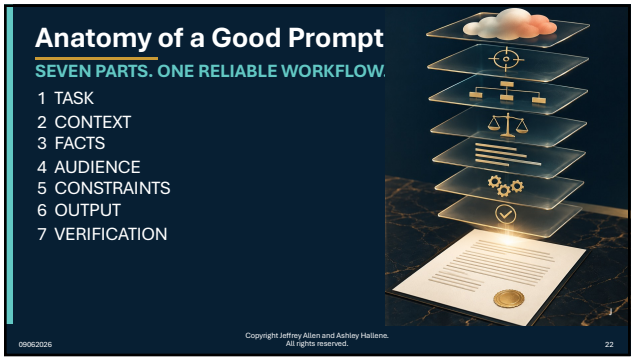
20



21



22



23



24

2. Supply Context

Tell AI:

- Who you represent
- Your objective
- The jurisdiction
- Relevant background
- What has already happened

Context changes the answer.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 24

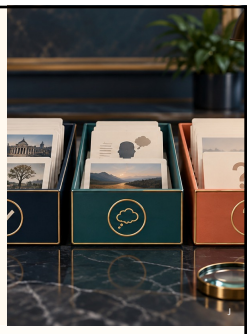
25

3. Give It the Facts

Include the material facts:

- Known facts
- Assumptions
- Missing information

Do not let AI silently fill in the blanks.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 25

26

4. Identify the Audience

Writing for:

- A client?
- A judge?
- Opposing counsel?
- Another lawyer?
- A nonlawyer employee?

Same issue. Different communication.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 26

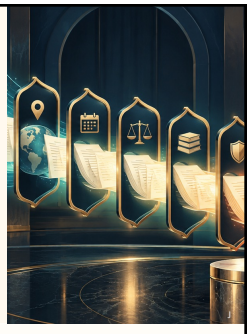
27

5. Set the Constraints

Tell AI the rules:

- Jurisdiction and date range
- Length and tone
- Permitted sources
- Assumptions
- Limitations

Be specific.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 27

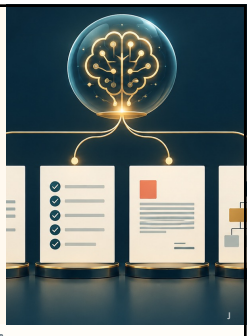
28

6. Specify the Output

What do you want?

- Memorandum or outline
- Table or checklist
- Letter
- Issues list
- Questions for further investigation

Do not make AI guess.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 28

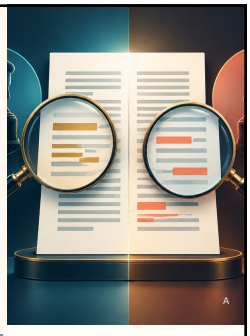
29

A Prompt Needs a Job and a Perspective

WEAK
"Review this agreement."

STRONGER
"Review this agreement for the buyer. Identify provisions that create unusual financial, termination, indemnification, or litigation risk."

The role changes what the model looks for and how it communicates.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 29

30

Give Context in Layers

1 **MATTER** — What kind of dispute or transaction?
2 **CLIENT** — Who do you represent?
3 **LAW** — Which jurisdiction and date govern?
4 **HISTORY** — What has already happened?
5 **PURPOSE** — How will the answer be used?



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 30

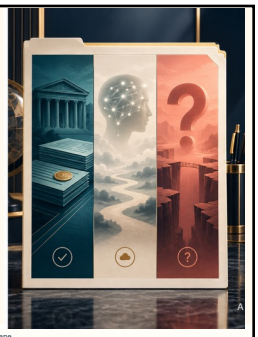
31

Separate Facts, Assumptions, and Gaps

FACTS
What the contract, client, and record establish

ASSUMPTIONS
What the model should treat as true for analysis

GAPS
What remains unknown and must not be inferred



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 31

32

Specify the Output and the Standard

OUTPUT
Client email, memorandum, comparison table, checklist, or witness questions

STANDARD
Plain language, neutral tone, no unsupported conclusions, flagged uncertainty, and cited propositions



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 32

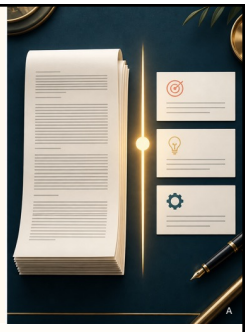
33

Use Delimiters for Long Source Material

Tell the model where instructions end and source material begins.

INSTRUCTIONS
Summarize the testimony. Identify contradictions and unanswered questions. Do not add facts.

SOURCE MATERIAL
<<< Paste the testimony here >>>



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 33

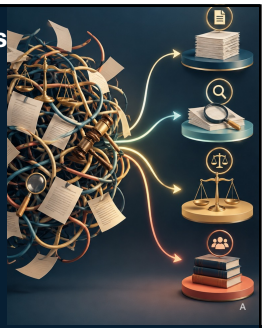
34

Avoid Compound Prompts

SEPARATE THE WORK TO CONTROL THE RESULT
One prompt. One principal task.

Break complex assignments into stages:

1 ANALYZE
2 CHALLENGE
3 DRAFT
4 VERIFY



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 34

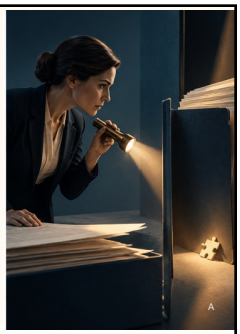
35

Ask What YOU Missed

TRY:

"What additional information would you need to give me a more reliable answer?"

AI can help identify gaps in your prompt.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 35

36

Protect Confidential and Sensitive Data

Before entering client or firm information:

- Use approved tools and accounts
- Follow firm policy and client restrictions
- Minimize or anonymize information
- Confirm retention, training, and sharing settings
- Stop if the data cannot be provided safely

Share no more than the task requires.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 36

37

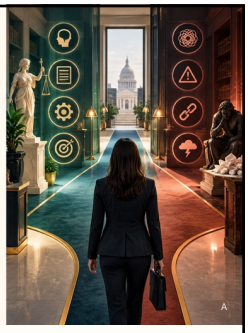
Prompting Do's and Don'ts

DO

Assign a clear job. Give relevant context. Separate facts from assumptions. Specify the output. Plan to iterate.

DON'T

Paste confidential data without approval. Bundle unrelated tasks. Invite guesswork. Treat tone as proof. Use output without review.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 37

38

A Reusable Prompt Blueprint

TASK — What should AI do?
CONTEXT — Who, why, where, and when?
MATERIAL — What facts or source text?
RULES — What limits and exclusions?
OUTPUT — What form, length, and tone?
CHECK — What must it flag or verify?



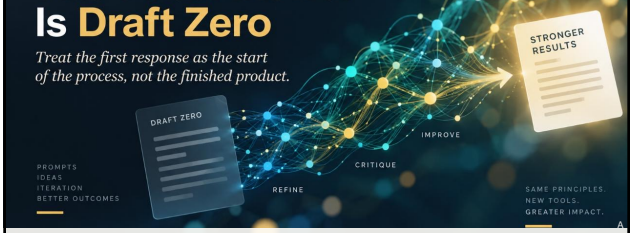
09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 38

39

ITERATION

The First Answer Is Draft Zero

Treat the first response as the start of the process, not the finished product.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 39

40

II

Iteration and Self-Critique

Working the draft. Then testing it.

PEOPLE
PRINCIPLES
TECHNOLOGY
A MORE CAPABLE PRACTICE

BETTER QUESTIONS. STRONGER OUTCOMES.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 40

41

LEGAL AI IN PRACTICE

An Iteration Example

SAME GOAL. A BETTER RESULT.

PEOPLE JUDGMENT A BRIGHTER TOMORROW

FIRST PROMPT

"Draft a short client email explaining the issue."

Result: usable, but generic.

FOLLOW-UP PROMPT

"Keep the tone friendly and lawyerly. Make it shorter. Add the missing deadline. Flag any assumption you are making."

TONE | LENGTH | MISSING FACTS | ASSUMPTIONS

KEY TAKEAWAY

Iteration is supervision in miniature: correct the draft, narrow the task, and keep judgment in the loop.

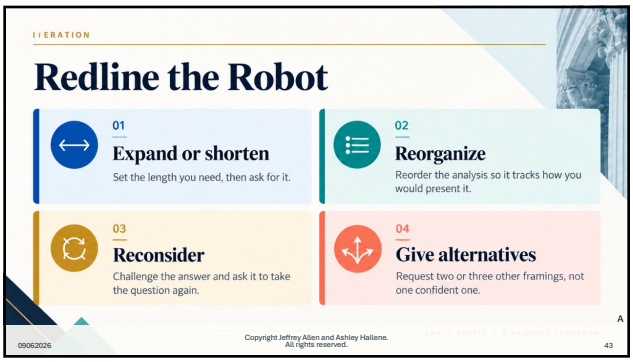
LAW MEETS POSSIBLE

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. 41

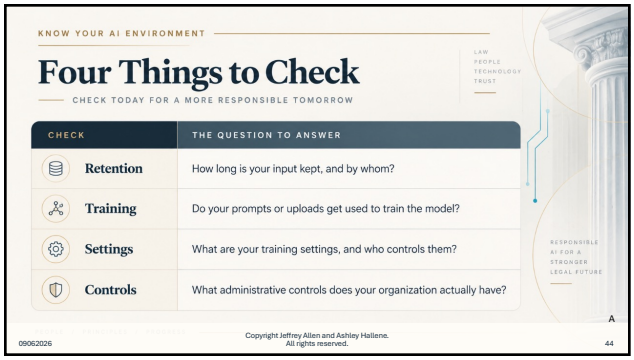
42



43



44



45



46



47



48



49



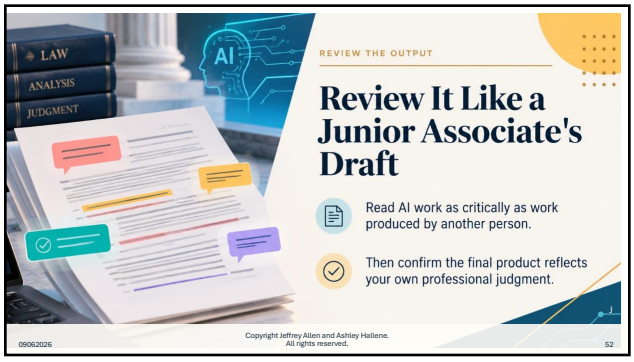
50



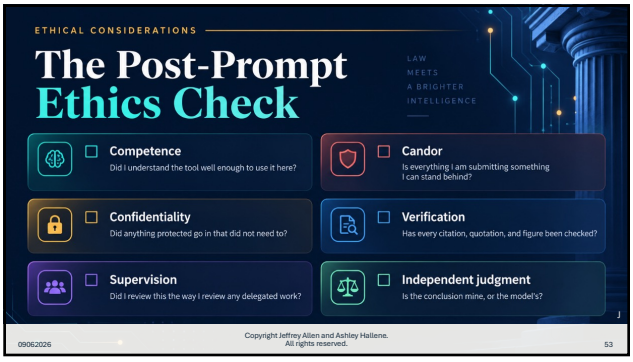
51



52



53



54



55



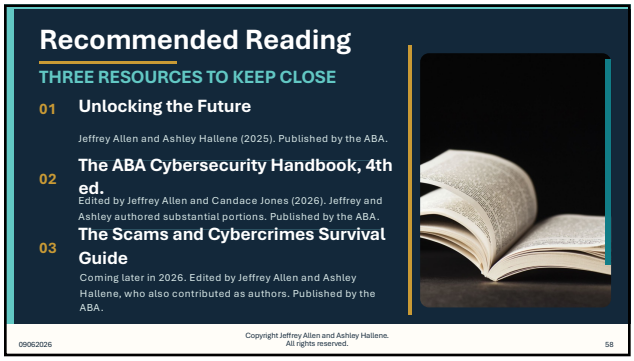
56



57



58



59



60



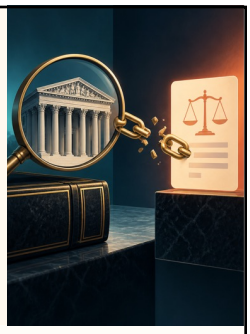
61

A Correct Citation Can Still Be Wrong

AI gives you:
✓ A real case
✓ A correct citation
✓ The right general subject

BUT THE CASE DOES NOT SUPPORT THE PROPOSITION.

Verification means reading the authority.



09062026
Copyright Jeffrey Allen and Ashley Hallene.
All rights reserved.
Bonus 1

62

You Own the Work Product

Before using AI-generated work, ask:

☐ Is it accurate?
☐ Is it complete?
☐ Is it current?
☐ Have I verified it?
☐ Does it reflect my professional judgment?



09062026
Copyright Jeffrey Allen and Ashley Hallene.
All rights reserved.
Bonus 2

63

Keep the Lawyer in Charge

Prompt carefully.
Question the answer.
Follow up.
Verify what matters.

KEEP THE LAWYER IN CHARGE.



09062026
Copyright Jeffrey Allen and Ashley Hallene.
All rights reserved.
Bonus 3

64

Give AI a Job, Not a Topic

INSTEAD OF
"Tell me about noncompetes."

TRY
"Analyze whether this employee noncompete is enforceable under California law from the employer's perspective."



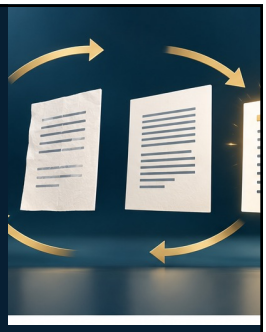
09062026
Copyright Jeffrey Allen and Ashley Hallene.
All rights reserved.
Bonus 4

65

The First Answer Is NOT the Last Answer

Prompting is a conversation.

READ IT → EVALUATE IT
↓
CORRECT IT → CHALLENGE IT
↓
REFINE IT



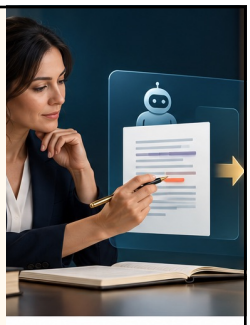
09062026
Copyright Jeffrey Allen and Ashley Hallene.
All rights reserved.
Bonus 5

66

Don't Like the Answer?

TELL AI WHY.

- "You focused on the wrong issue."
- "You overlooked this fact."
- "Make it less formal."
- "Explain your conclusion."
- "Give me three alternatives."
- "Reconsider your answer in light of…"



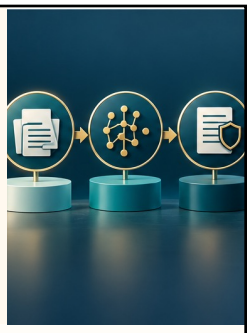
09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 6

67

Break Complex Work into Stages

1 Issue spotting
2 Questions and missing facts
3 Research plan
4 Draft analysis
5 Adversarial review

One prompt per stage makes errors easier to see and correct.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 7

68

7. Tell It What NOT to Do

TRY:

"Do not invent missing facts."
"Do not fabricate authorities."
"Identify assumptions."
"Tell me when you are uncertain."



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 8

69

Do Not Lead the Model to Your Answer

LEADING
"Explain why the clause is unenforceable."

NEUTRAL
"Analyze arguments for and against enforceability. Identify the strongest counterargument to each conclusion."

Ask what would change the answer.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 9

70

Put It Together

A GOOD PROMPT

TASK + CONTEXT + FACTS
+ AUDIENCE + CONSTRAINTS
+ OUTPUT + VERIFICATION

The more consequential the task, the more carefully you should construct the prompt.



09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 10

71

EXERCISE 1

A Client-Friendly Demand Letter Summary

THE PROBLEM

A client has received a six-page demand letter alleging breach of contract and threatening suit within ten days. The partner wants a short email that explains the allegations, deadlines, immediate risks, and next steps without giving a final merits opinion.

01

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 11

72

EXERCISE 1: PROMPT DRAFTING
A Client-Friendly Demand Letter Summary

WEAK PROMPT

"Summarize this demand letter for the client."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist counsel for the recipient. Using only approved attached materials, draft a plain-language client email of no more than 350 words. Identify: (1) allegations and requested remedies, (2) deadlines, (3) the sender's assertions versus facts supported by client documents, (4) urgent investigation, and (5) three next steps. Do not assess ultimate liability, invent facts, or quote more than one sentence. End with "What we need from you" and list what counsel must verify.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 12

73

EXERCISE 2
Contract Risk Review for the Buyer

THE PROBLEM

A buyer is considering a technology-services agreement. Counsel needs a focused review of termination rights, indemnity, liability limits, data security, automatic renewal, and provisions that shift unusual risk to the buyer.

02

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 13

74

EXERCISE 2: PROMPT DRAFTING
Contract Risk Review for the Buyer

WEAK PROMPT

"Review this contract and tell me what is wrong with it."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist buyer's counsel. Review only the approved attached agreement for termination, indemnity, liability limits, data security, automatic renewal, and dispute resolution. For each issue, give the clause reference, current effect, buyer risk, and a plain-English revision. Distinguish ordinary terms from buyer-unfavorable terms. Do not supply missing language or describe an absent clause. Flag uncertainty and unverifiable cross-references. End with what counsel must verify.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 14

75

EXERCISE 3
Deposition Contradictions and Follow-Up

THE PROBLEM

Counsel has a deposition transcript and several earlier witness statements. The goal is to locate material inconsistencies and prepare focused follow-up questions without overstating what the record proves.

03

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 15

76

EXERCISE 3: PROMPT DRAFTING
Deposition Contradictions and Follow-Up

WEAK PROMPT

"Find all the lies in this deposition."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist examining counsel. Compare the approved transcript and prior statements. Create a table showing topic, deposition testimony, earlier statement, exact source locations, type of inconsistency, and a neutral follow-up question. Distinguish contradictions from omissions, ambiguity, or changed recollection. Do not call the witness dishonest or invent an inconsistency. Flag uncertainty and quotations you cannot locate precisely. List what counsel must verify.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 16

77

EXERCISE 4
A Reliable Chronology from a Messy File

THE PROBLEM

A case file contains emails, letters, notes, and contracts with overlapping dates. The litigation team needs a chronology that separates confirmed events from disputed or inferred events.

04

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 17

78

EXERCISE 4: PROMPT DRAFTING

A Reliable Chronology from a Messy File

WEAK PROMPT

"Make a timeline of everything that happened."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist the litigation team. Using only the approved supplied file, prepare a table showing date, event, source, people involved, legal significance, and confidence. Label each entry confirmed, disputed, or inferred. If sources conflict, show both versions without resolving the conflict. Do not infer unsupported dates or facts. End with the five most important investigative gaps and the dates and source references counsel must verify.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 18

79

EXERCISE 5

Targeted Requests for Production

THE PROBLEM

After identifying gaps in the opposing party's document production, counsel wants draft requests aimed at communications, policies, transaction records, and electronically stored information while avoiding unnecessary breadth.

05

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 19

80

EXERCISE 5:

Targeted Requests for Production

WEAK PROMPT

"Draft discovery requests for this case."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist plaintiff's counsel. Using only the approved complaint, answer, discovery-gap memorandum, and form, draft no more than 12 requests for production. Organize them by communications, policies, transaction records, and electronically stored information. Define only needed terms, tailor dates to pleaded events, and avoid duplication. In an attorney note, link each request to a disputed issue. Do not invent facts or add boilerplate absent from the form. Flag assumptions. Counsel must verify scope, dates, definitions, and local-rule compliance before service.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 20

81

EXERCISE 6

A Compliance Checklist from a New Rule

THE PROBLEM

A new regulation affects client intake, required notices, record retention, and staff training. The firm wants an operational checklist that identifies effective dates and assigns responsibility.

06

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 21

82

EXERCISE 6

A Compliance Checklist from a New Rule

WEAK PROMPT

"Turn this regulation into a checklist."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist the firm's compliance lead. Using only the approved final regulation, create an implementation checklist. For each requirement, identify the section, required action, responsible role, deadline or frequency, evidence of completion, and open question. Separate mandatory duties from recommendations. Quote the effective-date language exactly. Do not fill regulatory gaps. Flag uncertainty. List every quotation, section, deadline, and interpretation counsel must verify against the official regulation.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 22

83

EXERCISE 7

An Initial Research Plan

THE PROBLEM

Counsel must evaluate whether a former employee's noncompete is enforceable. The agreement selects one state's law, the employee works in another state, and the employer seeks emergency relief.

07

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 23

84

EXERCISE 7

An Initial Research Plan

WEAK PROMPT

"Research whether this noncompete is enforceable."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist counsel with a research plan, not a legal conclusion. Evaluate the approved attached noncompete clause. Address choice of law, work location, statutes, current appellate authority, emergency-relief standards, protectable interests, and defenses. List threshold factual questions, primary-source searches, and useful search terms. Do not invent or cite unverified authority. Flag uncertainty and jurisdictional assumptions. Counsel must confirm each citation, holding, and proposition in a current authoritative source.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 24

85

EXERCISE 8

Negotiation Options for a Difficult Clause

THE PROBLEM

A vendor refuses to remove a broad indemnity clause. The client wants practical alternatives that reduce risk while preserving the deal and avoiding an unnecessarily adversarial tone.

08

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 25

86

EXERCISE 8

Negotiation Options for a Difficult Clause

WEAK PROMPT

"Rewrite the indemnity clause so it is fair."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist customer's counsel. Analyze the approved indemnity clause and propose preferred, compromise, and minimum-acceptable alternatives. For each, provide revised language, explain the risk reduction, identify the likely vendor objection, and suggest a response. Preserve defined terms. Do not revise unrelated provisions or assume outside facts. Flag uncertainty and related insurance, liability-cap, or defense-control terms. Counsel must compare each revision with the full agreement and verify cross-references.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 26

87

EXERCISE 9

Board Briefing on an AI Proposal

THE PROBLEM

Management proposes using an AI system to review customer complaints. The board needs a concise briefing on benefits, data risks, human oversight, vendor controls, and questions to ask before approval.

09

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 27

88

EXERCISE 9

Board Briefing on an AI Proposal

WEAK PROMPT

"Write a board memo about this AI product."

Why it fails: too little context, direction, and control.

WELL-DRAFTED PROMPT

Role: Assist counsel briefing the board. Using only the approved vendor proposal and internal policy, draft a two-page briefing. Cover the proposed use, expected benefit, data categories, material risks, required human review, vendor protections, and pilot conditions. Separate confirmed facts from vendor claims. Include six board-level questions and a recommendation with stated assumptions. Do not add undescribed capabilities or client-identifying information. Flag uncertainty and list what counsel must verify before circulation.

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 28

89

EXERCISE 10

Red-Team Review of a Draft Argument

THE PROBLEM

Counsel has drafted a motion argument and wants AI to expose vulnerabilities before filing, including missing authority, factual overstatement, weak counterarguments, and requested relief unsupported by the record.

10

A suitable task for AI-assisted legal work

09062026 Copyright Jeffrey Allen and Ashley Hallene. All rights reserved. Bonus 29

90


EXERCISE 10
Red-Team Review of a Draft Argument
WEAK PROMPT
"Tell me whether this argument will win."
Why it fails: too little context, direction, and control.
WELL-DRAFTED PROMPT
Role: Act as a skeptical reviewer for filing counsel. Evaluate only the approved draft argument and cited record. Identify unsupported premises, the strongest opposing argument, adverse authority or limits to research, overstated language, and relief that may exceed the record. For each criticism, cite the draft page and state what support or revision is needed. Do not predict the outcome or invent authority. Flag uncertainty. List every record citation, quotation, legal proposition, and authority counsel must verify in the originals.
09062026
Copyright Jeffrey Allen and Ashley Hallene. All rights reserved.
Bonus 30

91

Voice of Experience Voice of Experience: May 2026

How to Avoid Accidental Disclosure When Using AI

Jeffrey M Allen and Ashley Hallene

May 10, 2026 6 min read

Summary

- AI tools can improve legal drafting, research, and efficiency, but they also create client confidentiality risks when prompts and documents are processed on external cloud-based systems.
- Lawyers must treat AI use as a professional responsibility issue by understanding data retention, training practices, vendor terms, and the risk of accidental disclosure.
- Clear AI policies, secure platforms, staff training, and careful prompt drafting help law firms reduce disclosure risk while using artificial intelligence responsibly.



sportpoint via Getty Images

AI Column

Artificial intelligence tools are now part of everyday legal practice. Lawyers use them to draft documents, summarize contracts, research legal issues, and generate ideas. Used well, these tools can save time and improve workflow. Used carelessly, they can expose confidential client information.

That risk deserves attention now, not later. Many AI tools operate through cloud-based platforms that process user inputs on external servers. A lawyer who pastes facts into a prompt, uploads a draft agreement, or asks for help refining strategy may be sending sensitive information outside the firm without fully understanding where it goes, how long it stays there, or who may have access to it. Because confidentiality is a core professional duty, lawyers need to understand how these tools handle data before using them in client work.

It does not take a dramatic mistake to create a confidentiality problem. A lawyer might type what feels like a routine prompt, such as "revise this indemnity provision for our client's $20 million acquisition of a regional manufacturing company," without pausing to consider how much that single sentence reveals. In just a few words, the lawyer may have disclosed the type of deal, the size of the transaction, the client's position, and enough context to identify the matter. The real danger is not exotic misuse. It is ordinary convenience.

Confidentiality Reaches Far Beyond Privileged Communications

The duty of confidentiality is broad. It generally extends to all information relating to the representation of a client, not just information protected by the attorney-client privilege. That includes litigation strategy, draft contracts, financial information, internal legal analysis, and business plans. It can also include information that is publicly available, if it relates to the representation in a way that makes its use or disclosure problematic.

That matters in the AI context because lawyers may think they are safe as long as they avoid sharing a final document or a privileged email. But confidentiality risks can arise much earlier. A prompt that names the parties, describes a dispute, or summarizes a negotiation position may already reveal too much.

The governing principle is familiar. Lawyers must take reasonable steps to protect client information from unauthorized disclosure. Technology does not change that duty. It changes the setting in which the duty must be carried out.

AI Feels Private, but Often Is Not

AI tools are not inherently improper, but they create risks that differ from those that arise from more traditional office software. Many AI tools operate through external, cloud-based systems. When a lawyer enters text into an AI interface, the information may be transmitted to servers controlled by the provider rather than kept entirely within the firm's own environment.

What happens next depends on the platform. Information may be processed and deleted, stored for a period, reviewed by human personnel, or used to improve the system. Lawyers cannot safely assume that all AI tools handle data the same way or that a conversational interface is private simply because it feels private.

That point is easy to miss. AI tools often feel like digital research assistants. They are quick, polished, and conversational. Many function more like web-based services that receive, process, and sometimes retain user inputs outside of the organization. If lawyers treat these systems like private notebooks or internal colleagues, accidental disclosure becomes much more likely.

Routine AI Use Can Lead to Confidentiality Problems

Most confidentiality problems with AI will not come from reckless misconduct. They will come from ordinary, hurried work. Picture an ordinary workday. A lawyer is on deadline, the client wants revisions quickly, and the redline still needs work. To move faster, the lawyer drops part of the draft into an AI tool and asks for a tighter clause or a cleaner summary. It feels like a harmless productivity move, but if that text includes client names, transaction details, or strategy, a routine shortcut can become a confidentiality problem.

That same pattern appears in other everyday situations. A lawyer may upload a contract for summary, paste facts into a prompt to test an argument, or ask an AI tool to draft a clause tailored to a specific transaction. A paralegal may use a free public tool to clean up a client communication. A junior lawyer may try a personal AI account for research because it is faster than approved systems.

Each of those choices may seem minor. Together, they show how easily confidential information can move outside of approved channels.

Several common situations deserve special care:

- Uploading client documents, including draft contracts, discovery materials, due diligence files, or internal memoranda
- Entering detailed prompts that include party names, dispute facts, transaction terms, or litigation strategy
- Using AI tools to draft documents that contain client names, settlement amounts, or deal-specific provisions
- Allowing lawyers or staff to use personal or unapproved AI tools for work-related tasks

The pattern is simple: The more specific the prompt or upload, the greater the risk. Lawyers do not need to attach a full document to create a problem. A detailed prompt alone may reveal the heart of the matter.

Four Habits That Lower the Risk of Disclosure

The good news is that accidental disclosure through AI is a manageable risk. It requires judgment, policy, and training, not panic.

First, adopt a clear policy for AI use. Lawyers and staff should not be left to guess which tools are allowed or what information may be entered into them. A sound policy should identify approved platforms, explain whether documents may be uploaded, define what information must never be entered, and require human review of AI-generated work before it is used in client matters.

Second, avoid entering identifying client details unless there is a clear reason to do so and an approved platform is being used. In many cases, a lawyer can get useful help from a generalized prompt. A request for an indemnification clause in a commercial acquisition agreement is usually safer than a prompt that names the parties and states the purchase price.

Third, choose tools with stronger confidentiality protections. Many organizations now use enterprise AI platforms that provide better safeguards than free public versions. Depending on the provider, those protections may include encryption, limits on retention, contractual confidentiality commitments, and settings that prevent user inputs from being used to train future models. Even then, lawyers still need to confirm how the platform works.

Fourth, read the provider's terms. Before using an AI tool for legal work, lawyers should know the answer to a few basic questions: Is user input stored? Can the provider access it? Is it used to train the model? How long is it retained? Those are threshold questions, not technical details to ignore.

Supervision and Training Cannot Be an Afterthought

Even a good policy will fail if no one enforces it. Law firm leadership and legal department management need to take an active role in supervising AI use. That means more than approving a platform once and moving on. It means setting expectations, training personnel, and monitoring whether people are following the rules in practice.

This is especially important for junior lawyers and staff, who are often the quickest to experiment with new tools. Their interest in efficiency is understandable and often helpful, but it also creates risk if they do not understand the confidentiality implications of a prompt, an upload, or a personal account. Supervising attorneys should ensure that AI use complies with firm policy, that confidential information is not entered into unapproved systems, and that a human carefully reviews AI-generated work.

Training is one of the best safeguards. Many avoidable mistakes happen because users do not understand how AI systems process data or assume that a polished interface means the tool is private and secure. Clear, practical training can prevent those assumptions from turning into disclosure problems.

AI May Be New, but the Duty Is Not

AI tools can deliver real benefits. They can help lawyers draft faster, organize information, and generate ideas more efficiently. Those benefits are worth pursuing. But they must be balanced against the lawyer's duty to protect client information.

That balance is not difficult to describe, even if it sometimes takes discipline to maintain. Lawyers should understand the technology they use, put reasonable safeguards in place, choose platforms carefully, and keep humans responsible for the final judgment. AI may change how legal work is performed, but it does not change the professional duty that sits underneath that work.

Client trust still depends on the same promise it always has—that lawyers will protect sensitive client information. AI may be new, but the lawyer's job is not: Protect the client, protect the information, and use good judgment.

Authors



Jeffrey M Allen

Graves & Allen

Jeffrey Allen is a principal in the law firm of Graves & Allen, in Oakland, California. He runs a general practice that, since 1973, has emphasized real estate and business transactions, receiverships and related...



Ashley Hallene

Ashley Hallene is an Attorney and Land Acquisition Specialist with Demeter Land Development. She is based in Houston, Texas. She is the co-author of several books, including Technology Tips for Seniors Volume 2.0 (2018),...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

ABA American Bar Association

https://www.americanbar.org/groups/senior_lawyers/resources/voice-of-experience/2026-may/client-confidentiality-and-ai-how-to-avoid-accidental-disclosure/

Prompt Engineering: The New Art of Asking Questions

Jeffrey M Allen and Ashley Hallene

Oct 12, 2025 6 min read

Summary

- Learn how prompt engineering improves AI accuracy for lawyers.
- Discover practical AI applications in research, contracts, and litigation.
- Reduce AI hallucinations with step-by-step legal prompt strategies.



skynesher via Getty Images

Lawyers have always thrived on the power of the right question. In the courtroom, a precise cross-examination can shift a jury's perception. In a deposition, a carefully framed question can expose the weakness of an opponent's case. Even in a client meeting, clarity of inquiry builds trust and guides decision-making. Perry Mason, the fictional defense attorney who captivated television audiences for decades, rarely won his cases through grand speeches. Instead, he triumphed through sharp, well-timed questions that unraveled the truth.

Today, lawyers confront a new arena where questions hold similar power: the world of artificial intelligence. As AI tools such as ChatGPT, legal research assistants, and document analyzers move into legal practice, the quality of the results depends heavily on the quality of the questions posed to them. This practice has a new name in

technological circles: prompt engineering. For lawyers, who already appreciate the art of inquiry, prompt engineering is less about learning something foreign and more about applying familiar skills in a new context.

Think about in the context of the Socratic Method. Just as Socrates guided his students toward understanding by asking simple, progressive questions, lawyers can guide AI tools toward accuracy by structuring their prompts in smaller, manageable steps.

What is Prompt Engineering?

Prompt engineering refers to the process of carefully crafting instructions for AI systems so that they deliver accurate, relevant, and useful responses. Instead of treating AI like a magic box that provides instant solutions, prompt engineering treats it like a junior associate who requires direction and supervision.

A poorly worded prompt is like asking a witness a vague question. The answer might be irrelevant or misleading. A well-crafted prompt narrows the focus, defines the terms, and makes the AI more likely to provide something that can actually help the lawyer. The practice has quickly become essential because of a growing recognition of how AI systems operate. These tools predict words and phrases based on patterns in massive datasets, which means their output is only as reliable as the way the request is structured.

The Problem of Hallucinations

AI tools often produce what are known as hallucinations. In plain terms, this means the system generates responses that sound convincing but are factually false. Some lawyers have already learned this the hard way, with cases in which AI-produced briefs contained entirely fabricated case citations. Courts have responded with sanctions, warnings, and strong reminders that lawyers cannot delegate their professional responsibility to a machine.

Hallucinations are the digital equivalent of a witness who testifies with confidence but lacks any foundation for their claims. If unchecked, they can damage a lawyer's credibility, harm a client's case, and potentially cross ethical lines. For lawyers, who should already understand the value of double-checking every fact, prompt engineering offers a path to minimize these risks.

Breaking Prompts into Small Steps

The most effective way to reduce hallucinations and increase accuracy is to break prompts into small steps. Instead of asking AI to complete a sweeping task such as "Write a client memo on liability under this contract," a lawyer can divide the request into a sequence of smaller questions.

For example:

1. Ask the AI to summarize the contract provisions in plain English.
2. Next, ask it to identify which provisions might create legal risk.
3. Finally, ask it to outline possible strategies for addressing those risks.

This step-by-step method mirrors how lawyers naturally work. When preparing for a cross-examination, you never begin by asking the witness to admit the ultimate fact. Instead, you build the foundation with smaller, targeted questions that create a logical path toward your conclusion. The same principle applies in prompt

engineering. By controlling the sequence, the lawyer controls the accuracy and transparency of the AI's reasoning process.

The Socratic Method offers another useful analogy. Socrates did not lecture his students with final answers. He asked them a series of small, carefully designed questions that forced them to examine their assumptions and reach conclusions step by step. Those who adopt this method for prompting AI will find that it leads to more reliable and verifiable results.

Benefits for Lawyers

When lawyers break prompts into smaller steps, the results often feel more reliable and far less frustrating. Accuracy is the most immediate benefit. By narrowing the scope of each request, the AI has less room to wander into irrelevant speculation or invent details that do not exist. What emerges is a response that stays closer to the facts at hand.

Transparency also improves. Instead of waiting for a finished product that may contain errors, lawyers can see the process unfold one stage at a time. If an AI tool stumbles on the summary of a contract clause, the lawyer can intervene before the analysis moves forward. The experience is not unlike reviewing the work of a junior associate: it is easier to correct a misstep early than to rewrite an entire draft later.

Smaller prompts also save time. Errors are inevitable, but they are far less costly when the task is broken into pieces. If one answer falls short, only that portion needs to be reworked. The rest of the process can continue intact. This efficiency makes the interaction feel more like a conversation than a one-off request, with the lawyer guiding the exchange toward a useful outcome.

Perhaps the most significant benefit is confidence. Many senior lawyers approach AI with caution, wary of overpromises and concerned about ethical risks. Prompt engineering gives them a way to stay firmly in control. Instead of acting as passive recipients of whatever the machine produces, they become active supervisors, shaping the output step by step. In that role, the lawyer retains the same authority he or she exercises in every other part of practice: asking sharp questions, checking the answers, and ensuring the work meets professional standards.

Practical Application in Law Practice

Prompt engineering already has clear and practical uses across many areas of law. In legal research, for example, lawyers often want to move from facts to holdings to application. Rather than asking an AI system to deliver a polished memorandum in one pass, it works better to stage the process. Begin by requesting a summary of the facts, then ask for the holding, and only afterward inquire how that holding might apply to the situation at hand. By breaking the task into steps, the lawyer can more easily spot errors before they multiply.

Communication with clients offers another strong opportunity. Lawyers know how difficult it can be to translate dense legal concepts into plain language. With AI, the process becomes easier if handled in stages. The lawyer might first ask for a simple summary that captures the essence of the issue, then refine that output into a version that is appropriate for client correspondence. This method ensures clarity without sacrificing accuracy or professionalism.

Contract review also benefits from this approach. Large agreements often overwhelm even seasoned attorneys, and an AI assistant can help by isolating specific clauses one at a time. A lawyer might direct the system to identify indemnification provisions first, then move on to termination clauses, and finally highlight governing law. This mirrors the methodical way senior attorneys train junior associates, focusing on one issue before moving on to the next.

Litigation preparation may be the area where prompt engineering feels most familiar. Building a case has always involved starting with a broad set of possible arguments, refining them into stronger points, and then supporting them with authority. AI can assist in the same way. By asking for an outline first, then narrowing to key arguments, and finally requesting supporting cases, the lawyer creates a stepwise path that echoes the process of drafting cross-examination notes: broad at the start, precise at the finish.

Lawyerly Tropes and Takeaways

Lawyers can lean on familiar tropes to understand prompt engineering.

- **Channel Your Inner Perry Mason**. Just as Perry Mason relied on carefully timed questions rather than flashy speeches, prompt engineering rewards lawyers who ask clear, deliberate questions in the right sequence.
- **Think Socratically.** The Socratic Method teaches that truth emerges from step-by-step questioning. Apply that principle to AI, treating it as a student whose learning requires guidance.
- **Practice Like Cross-Examination.** Never ask compound or rambling questions. Keep prompts tight, focused, and purposeful. Just as in court, clarity is your strongest weapon.

Prompt engineering is not a new trick for young technology enthusiasts. It is a natural extension of skills that lawyers already possess. Asking precise questions, structuring inquiry step-by-step, and supervising the work of others have always been part of the profession. The difference is that the “other” in this case is not a junior associate but an AI system.

By embracing prompt engineering, lawyers can reduce the risks of hallucination, increase the accuracy of AI tools, and maintain professional control. The lesson is simple but powerful: in the courtroom, the right question can win the case. In the digital age, the right prompt can win the answer.

Authors



Jeffrey M Allen

Graves & Allen

Jeffrey Allen is a principal in the law firm of Graves & Allen, in Oakland, California. He runs a general practice that, since 1973, has emphasized real estate and business transactions, receiverships and related...



Ashley Hallene

Ashley Hallene is an Attorney and Land Acquisition Specialist with Demeter Land Development. She is based in Houston, Texas. She is the co-author of several books, including Technology Tips for Seniors Volume 2.0 (2018),...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

ABA American Bar Association | https://www.americanbar.org/groups/senior_lawyers/resources/voice-of-experience/2025-october/prompt-engineering/

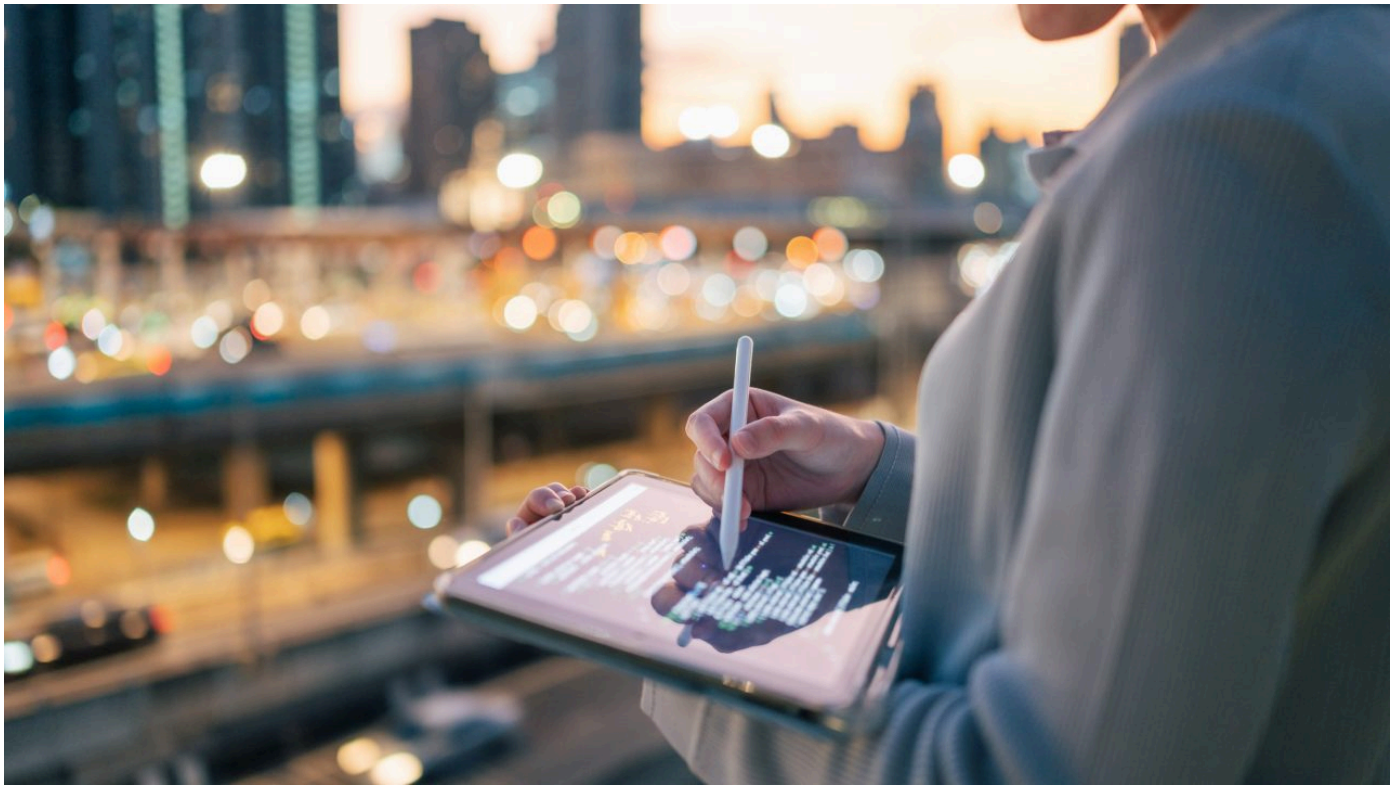
TAPAs: Advanced Prompt Engineering for Attorneys Using AI

Jeffrey M Allen and Ashley Hallene

Mar 03, 2026 13 min read

Summary

- **T**echnological **A**nd **P**ractice **A**dvice to help you become more efficient and effective. This month: advanced tips for using AI in your law firm for research, drafting, analysis, and other aspects of client service.
- This guide builds on foundational tips and offers next-level advice for achieving precision, reliability, and strategic value in GenAI interactions.
- At the end of this guide, we have provided a list of common pitfalls to avoid when drafting prompts, along with an appendix of legal prompt templates for attorneys.



iStock.com/primeimages

In the July 2025 issue of *GPSolo eReport*, we wrote an article about creating effective prompts for ChatGPT. That article looked at the basics of prompt engineering or prompt designing. Hopefully, you have worked on the skills we described there and have gained some proficiency in prompt engineering. This article picks up where the first one left off. Think of the initial article as "Prompt Engineering 1A." Think of this one as "Prompt Engineering 101A."

This article focuses on the use of generative artificial intelligence (GenAI) in legal practice. Nevertheless, its lessons will help you in every use of GenAI. Continue working on the skills we addressed in the first article. They give you the foundation for those addressed in this article. Attorneys need to acquire the skill of prompt engineering. It has become one of the most essential skills you must learn to deal effectively and efficiently with GenAI. GenAI has assumed a significant role in legal practice. You cannot succeed by playing ostrich and burying your head in the sand. GenAI will still be around when you pull your head out of the sand, but it will have a bigger and more significant presence. Moreover, while your head resides in the sand, you will lose the benefits GenAI offers.

In this article, we focus on how to achieve precision, reliability, and strategic value in legal AI interactions.

GenAI models are rapidly transforming the legal landscape, offering attorneys new ways to research, draft, analyze, and serve clients. The true power of these tools depends on how attorneys use them. Prompt engineering does not just address getting answers—it also shapes the quality, relevance, and strategic impact of those answers. This guide introduces advanced techniques for attorneys seeking to leverage GenAI for research, drafting, analysis, and other aspects of client service.

Tip 1. Deconstruct Legal Complexity into Modular Prompts

Legal tasks often involve layered issues—statutory interpretation, factual analysis, and risk assessment. Attempting to address all aspects in a single prompt can overwhelm GenAI and reduce the quality of its output. Instead, break down your requests into discrete modules:

Overloaded: *"Draft a memo on data privacy, cite statutes, analyze recent cases, and recommend compliance steps."*

Modular sequence:

1. *"Summarize the statutory framework for data privacy in California."*
2. *"List three recent cases interpreting CCPA enforcement."*
3. *"Draft compliance recommendations for a mid-sized tech firm."*

This approach ensures depth at each stage, allowing you to review and refine each component before assembling the final work product. Modular prompting also makes it easier to update or expand specific sections as laws evolve or new facts emerge.

Be sure you do not give these instructions at the same time; give them serially. Even though better drafted, they can still overwhelm GenAI if you provide them concurrently. As AI has been known to hallucinate (make things up and cite them as factually accurate), you should endeavor to mitigate that by including limitations requiring the citation of documentable or verifiable facts and requiring the AI to show its work by giving you accurate citations of authority for each fact it asserts. If you already have a prompt leaning toward the complex, you can do a cleanup at the end by having it review its response for any factually inaccurate content. Another approach is to limit the universe AI searches to a specific source or class of sources. In an abundance of caution, you might run the responses of an AI through another AI for fact-checking. Of course, it should go without saying that you need to review all sources cited by your AI to verify accuracy of citation and analysis before passing the conclusions along to your client or to a court.

Tip 2. Specify Legal Standards, Jurisdictions, and Timeframes

GenAI thrives on specificity. The context-dependence of legal analysis requires that you define the legal standard, jurisdiction, and relevant dates. For example:

- *"Explain the 'reasonable expectation of privacy' standard under U.S. Fourth Amendment law, focusing on cases since 2020."*
- *"Compare the duty of care in negligence under New York and California law, with examples from the past two years."*

By anchoring your prompt in a particular legal framework, you guide GenAI to deliver tailored, actionable insights. That result has special importance when dealing with multi-jurisdictional matters or rapidly changing areas of law, such as data privacy or employment regulations.

Tip 3. Assign Professional Roles and Legal Writing Styles

Direct GenAI to adopt the voice of a legal expert, such as a litigator, transactional attorney, or judicial clerk. The model can employ the tone, structure, and rhetorical strategies of different legal professionals:

- *"You are a senior appellate attorney. Draft a persuasive argument challenging the admissibility of digital evidence."*
- *"Write a client advisory in plain English, suitable for small business owners."*
- *"Act as a judicial clerk and summarize the majority and dissenting opinions in* Smith v. Jones *(citation)."*

Role-based prompts produce contextually appropriate output, so that the writing style matches the intended audience and purpose. That outcome will prove particularly valuable whether preparing documents for court, clients, or internal review.

Tip 4. Provide Case Context, Client Background, and Strategic Objectives

Providing more relevant, practical background information enables GenAI to generate better output. Include the client's industry, the matter's purpose, and strategic goals. For example:

- *"Draft a risk assessment memo for executive review for a client in health care facing HIPAA compliance issues."*
- *"Prepare a litigation strategy outline for a breach of contract dispute involving SaaS agreements."*
- *"Rewrite this paragraph for a client memo about non-compete agreements, targeting a small business owner without legal training."*

Rich context guides GenAI to deliver accurate results, aligned with your client's needs and business realities. It also helps avoid generic or irrelevant responses that may miss critical nuances.

Tip 5. Request Structured Legal Output

Legal work benefits from structure: outlines, tables, timelines, and checklists. Structured prompts save time and enhance clarity, making it easier to review, edit, and share results. For example:

- *"Create a table comparing indemnification clauses in three sample contracts."* [Provide the contracts to AI for the most useful response.]
- *"Outline the steps for responding to a subpoena, with deadlines and responsible parties pursuant to California law."*
- *"Summarize the article in PowerPoint slides with three bullet points per slide."*

By specifying the desired format, you ensure usable output that reduces the need for post-processing. Structured outputs can prove especially helpful for presentations, client deliverables, and internal training materials.

Tip 6. Supply Exemplars: Model Clauses, Brief Excerpts, or Preferred Tone

Show GenAI what "good" looks like by providing sample clauses, argument styles, or tone. Even a single sentence can set the standard for the rest of the output:

- *"Match the tone of this opening statement: 'The facts are clear, the law is settled.'"*
- *"Use this template for a confidentiality clause: [insert model text]."*
- *"Here's the tone I like: 'The truth is simple. The law is not.' Please match this tone in the rest of the paragraph."*

Examples drive consistency and quality, allowing you to maintain firm standards and client expectations across multiple drafts and documents.

Tip 7. Set Boundaries: Word Limits, Citation Requirements, and Jargon Controls

Control the scope and style of GenAI's output by setting clear boundaries:

- *"Limit the summary to 300 words."* [NOTE: We have found that AI's talents do not include counting; if you tell an AI that you want 1,500 words, you might get any number of words—usually less than 1,500—so it can help to ask for more than you need.]
- *"Cite only published U.S. Court of Appeals and U.S. Supreme Court decisions."*
- *"Avoid technical jargon; this is for a lay audience."*
- *"Summarize each court case in no more than five sentences."*

Boundaries ensure relevance and accessibility, helping you tailor the output to specific audiences, such as clients, judges, or colleagues. They also prevent GenAI from producing overly verbose or tangential responses and may reduce the incidence of hallucinations.

Tip 8. Iterate: Refine, Challenge, and Validate

The iterative nature of legal analysis suggests that we use GenAI in cycles—draft, critique, revise. Each round of feedback sharpens the output:

1. *"Draft a summary of the holding in* Smith v. Jones*."*
2. *"Now add a counterargument based on dissenting opinions."*
3. *"Strengthen the analysis with policy considerations."*
4. *"Cite two recent cases or authorities if available."*

Iterative prompting allows you to test different approaches, challenge assumptions, and validate the reasoning. It also helps you adapt the output to new facts, legal developments, or client feedback.

Tip 9. Demand Source Transparency and Fact-Checking

Good practice requires a focus on accuracy. Instruct GenAI to avoid fabricating citations (hallucinations) and to flag unverifiable claims:

- *"Provide a summary of the GDPR's impact on U.S. law firms. Include supporting authority for each statement. Do not invent citations; I will verify sources."*
- *"List only cases available on Lexis or Westlaw."*
- *"Provide a general overview of the GDPR's impact on U.S. law firms. No fake citations. I will verify sources separately."*

Source control helps protect your credibility and increases the likelihood that the output meets professional standards. When brainstorming, you can request "no citations" to focus on ideas rather than verification.

Tip 10. Label Iterations and Track Versions

For ongoing matters, name each draft and specify changes. Clear labeling streamlines collaboration and review:

- *"This is version 3 of the client memo on employment law updates. Expand section 2 with recent EEOC guidance."*
- *"This is draft 2 of a slide deck on mobile cybersecurity for lawyers. Please expand slide 5 with an example of a recent phishing attack targeting legal professionals."*

Tracking versions helps you and GenAI stay aligned across sessions, making it easier to manage revisions and maintain a clear record of changes.

Common Pitfalls in Legal Prompt Engineering

1. Vague Prompts

- **Pitfall:** Prompts that lack specificity lead to generic, incomplete, or irrelevant responses.
- **Example:** *"Write about contracts."*
- **Why it's a problem:** GenAI cannot determine which aspect of contracts you want—formation, enforcement, breach, or a specific jurisdiction.

- **How to fix:** Be precise about the topic, scope, and context.
- **Improved prompt:** *"Summarize the key elements required for contract formation under Florida law, with examples from recent cases."*

2. Overloaded Prompts

- **Pitfall:** Trying to address multiple complex issues in a single prompt overwhelms GenAI and dilutes output quality.
- **Example:** *"Analyze five statutes and draft three clauses in one response."*
- **Why it's a problem:** The model may skip steps, provide shallow analysis, or miss important details.
- **How to fix:** Break the task into modular prompts and tackle each issue separately.
- **Improved sequence:**
    1. *"Summarize the five relevant statutes."*
    2. *"Draft a clause for each statute, one at a time."*

3. No Context Provided

- **Pitfall:** Prompts without background information or clear objectives result in unfocused or irrelevant answers.
- **Example:** *"Fix this."*
- **Why it's a problem:** GenAI doesn't know what needs fixing—grammar, legal reasoning, formatting, or compliance.
- **How to fix:** Provide context, including the legal issue, the audience, and the desired outcome.
- **Improved prompt:** *"Revise this paragraph to clarify the enforceability of non-compete agreements for a client memo intended for small business owners."*

4. Ignoring Audience Awareness

- **Pitfall:** Using technical legal jargon for a nonlawyer audience can confuse or alienate readers.
- **Example:** *"Draft a memo on indemnification using standard legal terminology."*
- **Why it's a problem:** Clients or business stakeholders may not understand the memo, reducing its effectiveness.
- **How to fix:** Specify the audience and adjust the tone and complexity accordingly.
- **Improved prompt:** *"Write a plain English summary of indemnification clauses for startup founders with no legal background."*

5. Unclear Boundaries and Requirements

- **Pitfall:** Failing to set word limits, citation requirements, or formatting instructions can result in overly verbose or unusable output.
- **Example:** *"Summarize recent cases on data privacy."*
- **Why it's a problem:** You may receive a lengthy, unfocused summary or citations that are unverifiable.
- **How to fix:** Set clear boundaries and requirements.
- **Improved prompt:** *"Summarize three recent appellate cases on data privacy in no more than 300 words, citing only published decisions."*

6. Not Iterating or Validating

- **Pitfall:** Accepting GenAI's first response without refinement or fact-checking can lead to errors or missed opportunities for improvement.
- **Example:** *"Draft a client advisory on HIPAA compliance."*
- **Why it's a problem:** The initial draft may lack nuance, miss recent developments, or contain inaccuracies.
- **How to fix:** Iterate: review, critique, and request revisions.
- **Improved workflow:**
    1. *"Draft a summary of HIPAA compliance requirements."*
    2. *"Add recent enforcement actions from 2024."*
    3. *"Clarify implications for telehealth providers."*

7. Assuming the Accuracy of AI Output

- **Pitfall:** Trusting GenAI to provide correct citations, facts, or legal analysis without verification leaves you vulnerable to hallucination and errors.
- **Example:** *"List supporting cases for this argument."*
- **Why it's a problem:** AI may fabricate citations or misinterpret legal standards.
- **How to fix:** Always verify sources and instruct GenAI to flag unverifiable claims.
- **Improved prompt:** *"List only cases available on Lexis or Westlaw. Flag any sources you cannot verify."*

Prompt Engineering as Legal Craft

Just as legal writing demands clarity and purpose, prompt engineering requires attorneys to frame questions with intent, context, and rigor. The quality of GenAI's output mirrors the quality of your prompts. Treat each interaction as an opportunity to advance your client's interests and your professional excellence.

Prompt engineering is not a one-time skill but an ongoing craft. As GenAI evolves, so, too, must your approach—experiment with new techniques, refine your instructions, and stay attuned to the ethical and professional standards of legal practice.

Important Warning: Never Trust AI to Get It Right. Even the Best AI Can Hallucinate

Promoters may tell you how reliable and dependable they have made their AI. Still, to protect itself, the company will likely have a disclaimer and/or waiver about hallucinations buried in its paperwork for you to sign. It's fine to ask an AI to find supporting information for points in an article or a brief. Often, you will get usable results, sometimes even good ones. Sometimes, you will not.

We recently ran a test by stripping the footnotes from an article and giving it to an AI with instructions to find supporting information for the article and format it in bluebook form. To avoid overwhelming the AI, we broke the article down into sections. We told the AI not to make up citations and to use only verifiable sources delivered in Bluebook form; we further instructed the AI that it need not limit itself to cases and statutes but could also draw from news broadcasts, newspapers, and other articles available online. The AI delivered 97 footnotes in fair Bluebook form. The article originally had 54 footnotes. We thought the 97 represented an improvement (although it included a number that we did not consider necessary) until we started verifying the footnotes. Of the 97 footnotes, 94 had errors or represented fabrications. Some of these citations had nonexistent sources, and others had broken links (several of which we could easily fix).

This AI model was not built specifically for lawyers or legal use. Nevertheless, the AI correctly cited relevant statutes and appropriately referred to cases and the American Bar Association Model Rules of Professional Conduct. It also included incorrect citations for some of the cases; it cited news articles that, as far as we could tell, did not exist. In fairness, we have had better results on other occasions, and sometimes we have had absolutely correct results. Still, we would submit nothing generated by AI without first verifying its accuracy. We do not care how good GenAI gets; we will continue that practice. Attorneys have an ethical duty to do so under the ABA Model Rules and likely the rules adopted by your state bar.

Takeaway: Try, try again. Think of your work with AI as a conversation and not as a one-and-done. If you don't get exactly what you want from the AI, tell it what you want it to change and how you want it modified. The AI won't mind or take offense to your additional requests, but it will modify the response to comply with your fine-tuning.

Appendix: Legal Prompt Templates for Attorneys

How to Use These Templates

- **Customize** the bracketed sections ([...]) for your specific legal issue, jurisdiction, audience, or requirements.
- **Combine** templates for multistep tasks (e.g., research, drafting, and client communication).
- **Iterate** by refining your prompts based on GenAI's output and your review.

1. Legal Research

- **Jurisdiction-specific statute summary:** *"Summarize the statutory framework for [legal issue] in [jurisdiction]."*

- **Case law search:** *"List three recent cases interpreting [law/statute] in [jurisdiction] since [year]."*
- **Comparative analysis:** *"Compare the duty of care in negligence under [jurisdiction 1] and [jurisdiction 2] law, with examples from the past [number] years."*

2. Drafting Legal Documents

- **Contract clause drafting:** *"Draft a [type of clause, e.g., confidentiality clause] for a [type of contract, e.g., technology licensing agreement], including provisions for [specific requirements]."*
- **Memo preparation:** *"Write a client advisory in plain English, suitable for [audience, e.g., small business owners], about [legal topic]."*
- **Litigation document:** *"Outline the key arguments for a motion to dismiss based on [legal basis, e.g., lack of personal jurisdiction]."*

3. Client Communication

- **Risk explanation:** *"Explain the risks of [legal issue, e.g., non-compete agreements] to a [client type, e.g., small business owner] in plain English."*
- **Regulatory update:** *"Summarize the impact of recent [regulation, e.g., data privacy regulations] for [industry, e.g., health care providers].*

4. Internal Training and Knowledge Management

- **Checklist creation:** *"Create a checklist for onboarding new associates in the [practice area, e.g., litigation department]."*
- **Newsletter summary:** *"Summarize recent developments in [area of law, e.g., intellectual property law] for the monthly firm newsletter."*

5. Structured Output Requests

- **Table comparison:** *"Create a table comparing [legal concept, e.g., indemnification clauses] in three sample contracts."*
- **Timeline/checklist:** *"Outline the steps for [legal process, e.g., responding to a subpoena], with deadlines and responsible parties."*

6. Role-Based Prompts

- **Judicial clerk summary:** *"Act as a judicial clerk and summarize the majority and dissenting opinions in [case name]."*
- **Senior attorney argument:** *"You are a senior appellate attorney. Draft a persuasive argument challenging the admissibility of [evidence type, e.g., digital evidence]."*

7. Boundary and Format Controls

- **Word limit:** *"Limit the summary to [number] words."*
- **Citation requirement:** *"Cite only published appellate decisions."*
- **Jargon control:** *"Avoid technical jargon; this is for a lay audience."*

8. Iteration and Validation

- **Refinement sequence:**
  - 1. *"Draft a summary of [legal holding or issue] in [case name]."*
  - 2. *"Now add a counterargument based on dissenting opinions."*
  - 3. *"Strengthen the analysis with policy considerations."*
  - 4. *"Cite two recent cases or authorities if available."*

9. Source Transparency

- **Verification prompt:** *"List only cases available on [legal database, e.g., Lexis or Westlaw]. Flag any sources you cannot verify."*

Authors



Jeffrey M Allen

Graves & Allen

Jeffrey Allen is a principal in the law firm of Graves & Allen, in Oakland, California. He runs a general practice that, since 1973, has emphasized real estate and business transactions, receiverships and related...



Ashley Hallene

Ashley Hallene is an Attorney and Land Acquisition Specialist with Demeter Land Development. She is based in Houston, Texas. She is the co-author of several books, including Technology Tips for Seniors Volume 2.0 (2018),...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

TAPAs: Creating Effective Prompts for ChatGPT

Jeffrey M Allen and Ashley Hallene

May 16, 2025 4 min read

Summary

- **T**echnological **A**nd **P**ractice **A**dvice to help you become more efficient and effective. This month: tips for creating effective prompts for generative artificial intelligence (GenAI) tools such as ChatGPT.
- If you ask for too much in one prompt, GenAI may underdeliver or skip key points. Dividing complex tasks improves clarity and output quality.
- If your prompt lacks context, GenAI will make assumptions. If incorrect, those assumptions may cause an irrelevant or incomplete result.
- Don't expect perfection on the first try. Better GenAI tools respond well to feedback and follow-up instructions.



pixdeluxe via Getty Images

Generative artificial intelligence (GenAI) continues to grow more sophisticated, and we use it for an ever-increasing variety of things. We use it increasingly to help with writing. This article focuses on that use and assumes you have one of the better GenAI programs, such as the ChatGPT app or Copilot. These tips should help

with any GenAI model, but we still hope you have a good one. Despite its power, as with any tool, the quality of the result you get with GenAI depends heavily on how you use it. Think of prompt writing as giving instructions to a capable but literal assistant: The more clearly you define what you want, the more likely you'll get it.

We will depart from our usual format to give you better organization of the information and follow our own advice in Tip 1. Here's a breakdown of ten core tips, each with expanded guidance and examples.

Tip 1. Break Big Tasks into Smaller Ones

If you ask for too much in a single prompt, GenAI may underdeliver or overlook key points. Dividing complex tasks improves clarity and output quality.

Ineffective:

- "Write an article about AI, cybersecurity, and the legal system, with examples, ethics, citations, and a glossary."

Better Sequence:

1. "Outline an article on AI and cybersecurity in law."
2. "Expand the section on ethical implications of AI in evidence."
3. "Add real-life examples from the past five years."
4. "Write a glossary of five AI terms mentioned in the article."

Tip 2. Be Specific and Direct

GenAI can respond to broad or vague questions, but it works better when your prompt describes *exactly* what you want. General requests often produce generic results. The more precise your language, the better the answer.

Examples:

- **Don't say:** "Write about contracts."
- **Instead, say:** "Explain the difference between unilateral and bilateral contracts with two examples, each under 100 words."

Tip 3. Define a Role, Style, or Tone

GenAI can adjust its voice based on the role or tone you assign it. You can instruct it to sound like a journalist, teacher, trial attorney, therapist, or corporate consultant; the better iterations of GenAI usually adapt accordingly.

Examples:

- "You are a legal writing instructor. Rewrite this brief to be more persuasive."

- "Adopt a conversational tone suitable for a newsletter aimed at small firm lawyers."
- "Act as a product reviewer and compare two legal research platforms."

The more clearly you define the speaker or audience, the more tailored the output.

Tip 4. Provide Context or Background

If your prompt lacks context, GenAI will make assumptions. If incorrect, those assumptions may cause an irrelevant or incomplete result. When writing prompts, include:

- The purpose of the task (e.g., educational, marketing, persuasive).
- The intended audience (e.g., seniors, attorneys, clients, children).
- Relevant details that shape the request.

Examples:

- **Don't say:** "Make this simpler."
- **Instead, say:** "This paragraph is for a client memo about non-compete agreements. Please rewrite it in plain English, targeting a small business owner without legal training."

Tip 5. Ask for Your Preferred Format or Structure

If you have a preferred output format—outline, bullet list, table, paragraph, slide summary—*say so.* This improves clarity and saves you editing time.

Examples:

- "Create an outline with five main sections and two bullet points under each."
- "Present the comparison in a table with three columns: Feature, Option A, Option B."
- "Summarize the article in PowerPoint slides with three bullet points per slide."

Structured prompts lead to structured results.

Tip 6. Give Examples or Templates

When you show GenAI what "good" looks like, it can mimic the structure, rhythm, or logic you like. Even one or two lines of sample text can shape the model's style.

Example:

- "Here's the tone I like: 'The truth is simple. The law is not.' Please match this tone in the rest of the paragraph."

Templates Also Help:

- "Use this format: Term—Short definition—One sentence explanation of why it matters. Now do that for ten AI terms."

Tip 7. Set Limits or Parameters

You control the length, number of items, and the type of content you want. Tell the model how long you want the product to be, what it should look like, and what to include or not include.

Examples:

- "Keep your answer under 250 words."
- "List three pros and three cons."
- "Avoid legal jargon—this is for the general public."
- "Summarize each court case in no more than five sentences."

Tip 8. Use Iteration: Prompt, Refine, Repeat

Don't expect perfection on the first try. Better GenAI responds well to feedback and follow-up instructions. Start with a rough version, then improve it step-by-step.

Workflow Example:

1. "Write a brief description of AI in criminal law."
2. "Now add an example involving predictive policing."
3. "Make it sound more persuasive."
4. "Cite two recent cases or authorities if available."

By refining your output in stages, you gain control over style, substance, and relevance.

Tip 9. Clarify That You Want Fact-Checking or Source Verification

Accuracy matters; say so, and note whether you will verify citations yourself. GenAI may fabricate details unless instructed not to do so.

Example:

- "Provide a general overview of the GDPR's impact on U.S. law firms. No fake citations. I will verify sources separately."

When needed, ask for accurate and *verifiable* sources or request "no citations" if you're only brainstorming.

Tip 10. Name Your Project or Version (for Iterative Work)

Label your drafts clearly. It helps you and the AI stay aligned across sessions.

Example:

- "This is draft two of a slide deck on mobile cybersecurity for lawyers. Please expand slide five with an example of a recent phishing attack targeting legal professionals."

Bonus Tip: Common Prompting Mistakes to Avoid

- **Too vague:** "Write about AI."
- **Overloaded prompt:** "Do ten different things in one response."
- **No context:** "Fix this" (without telling the AI what you're doing or why).
- **No audience awareness:** Writing legal content for the general public using legalese.

Final Thought: Prompt Like a Lawyer, Not a Tourist

Ask deliberate, specific, and well-informed questions. Be clear about purpose, audience, tone, and format. As with cross-examination or brief writing, the outcome depends on how you frame the question.

Authors



Jeffrey M Allen

Graves & Allen

Jeffrey Allen is a principal in the law firm of Graves & Allen, in Oakland, California. He runs a general practice that, since 1973, has emphasized real estate and business transactions, receiverships and related...



Ashley Hallene

Ashley Hallene is an Attorney and Land Acquisition Specialist with Demeter Land Development. She is based in Houston, Texas. She is the co-author of several books, including Technology Tips for Seniors Volume 2.0 (2018),...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

Experience Experience April-May 2026

What Attorneys Need to Know About Ag AI

Jeffrey M Allen and Ashley Hallene

Apr 29, 2026 10 min read

Summary

- ChatGPT helps lawyers think and write, and agentic AI acts. That shift from assistance autonomous action fundamentally changes risk, supervision duties, and professional responsibility.
- GenAI errors stay on the page; agentic AI errors can cause real-world harm.
- Attorneys remain fully accountable for ethics, confidentiality, and outcomes, regardles sophisticated or autonomous technology becomes.



iStock.

Jump to:

Introduction →

What Is ChatGPT/GenAI? →

What Is Agentic AI? →

Detailed Comparison →

Why the Distinction Matters →

What Each Tool Does Well →

What Neither Tool Does Reliably →

Risk: Similar Roots, Different Scale →

Risks and Ethical Considerations →

Client Confidentiality and AI Risks →

Economic Impact on Law Firms →

Regulatory Landscape and Compliance Challenges →

AI in Litigation Strategy and Risk Management →

Conclusion →

Introduction

Just When You Thought You Could Safely Take Your Head Out of the Sand...

Over the past few years, many lawyers have become at least casually familiar with Generative AI (GenAI), sometimes called ChatGPT (imprecisely). ChatGPT, a form of GenAI, has opened GenAI technology to many uses in and out of the legal profession. Because of ChatGPT's impact, many people use the term interchangeably with GenAI.

Some use GenAI regularly, growing increasingly comfortable with it as a partner or an assistant. Others experimented, shrugged (perhaps with disdain), and moved on. Still others saw it as interesting but not something they trusted with anything important. We consider all those reactions reasonable. Some looked at it as a high-tech form of Frankenstein, stuck their heads in the sand, and hoped it would go away. We don't see that as a reasonable approach, but to each their own...

We have come to accept artificial intelligence (AI) (some of us reluctantly), and it has planted itself in our lives and practices. It continues to grow and expand its scope of competence. Many of the sand people have started to emerge, thinking we have some reasonable grasp of AI now and acknowledging that it has come to stay.

But just when they thought they could safely emerge...

Now, a new term has entered the conversation: **Agentic AI**. GenAI raised issues of accuracy, ethics, confidentiality, and professional judgment. Agentic AI continues to grapple with those problems and adds even more complex concerns.

What Is ChatGPT/GenAI?

A generative AI model developed by OpenAI, ChatGPT, reads prompts and produces human-like text or other media. Because of its ubiquitous presence, many began using the terms “GenAI” and “ChatGPT” interchangeably, much as “Kleenex” became the generic term for tissues.

It may prove helpful to think of GenAI as an advanced autocomplete system—impressive, helpful, and sometimes far too confident for its own good. In most consumer and professional deployments, GenAI operates reactively, responding to prompts rather than initiating actions on its own.

In our experience, GenAI can help attorneys draft contracts, summarize case law, generate marketing content, and translate documents. These tasks save time, reducing repetitive work. For example, an attorney can ask a GenAI application to summarize a 50-page case into a one-page brief. The tool delivers a readable summary in seconds.

GenAI, however, has limits. It cannot act autonomously, guarantee factual accuracy, or reliably provide legal advice. It may hallucinate facts or cite non-existent cases. GenAI has no internal moral compass. It does not perceive right from wrong as humans do. GenAI does not think the way humans do. It predicts language. When you ask a question or give it a task, it responds based on patterns it learned during training. Those responses can come across as impressively fluent, and they often prove helpful. They can also end up confidently wrong. GenAI also does not verify facts the way lawyers expect verification to work. It does not “look up” cases, unless connected to a tool designed to retrieve them, and even then, it may misinterpret what it finds. OpenAI has publicly acknowledged that large language models can hallucinate—producing statements that sound plausible but are completely false. That risk does not disappear just because the writing sounds authoritative.

From a lawyer's perspective, GenAI works best as a language tool. It drafts, rewrites, summarizes, explains, and organizes information quickly. It can turn a rough idea into a polished paragraph, translate legal jargon into plain English, or condense a long document into something quickly readable.

What Is Agentic AI?

Agentic AI represents the next evolution of AI—yes, already, whether we feel ready for it or not. Agentic AI can act independently. Agentic AI can pursue predefined goals, plan intermediate steps, *and execute tasks* with minimal ongoing human intervention. Imagine an AI agent tasked with filing documents near a deadline, using outdated instructions copied from a prior matter. Or it misunderstands your instructions and sends confidential files to the wrong party—without waiting for your approval. You only discover the error after the damage has occurred.

Agentic AI behaves like a digital worker. It integrates with external tools, adapts to changing conditions, and completes multi-step workflows. Agentic AI can also learn from feedback and adjust its strategies on the fly. Agentic AI builds on the same language-model foundation as ChatGPT, but it adds something new: *the ability to plan, decide, and act across multiple steps without continuous human prompting*. This autonomy makes it a powerful assistant but also introduces new risks. For lawyers, this means that *errors no longer stay on the page. They can turn into actions*.

We like the idea of AI continuing to expand, but we are not quite ready to unleash the full power of Agentic AI on our practices. The concept of AI acting autonomously keeps us up at night. Yet, we know it has arrived and will continue to expand its impact across our practices and personal lives.

Detailed Comparison

While GenAI reacts, Agentic AI moves proactively. In our experience, this shift means lawyers must rethink how they supervise technology. GenAI focuses on generating text and other media. Agentic AI *executes tasks*. GenAI has limited

integration and static learning. Agentic AI can dynamically adjust its actions within a task and connect to multiple systems. These differences matter because they affect risk, governance, and ethical responsibilities. Another key difference lies in accountability. With GenAI, the attorney controls each interaction (although sometimes poorly). Agentic AI can act without explicit approval, raising questions about liability. Firms must design workflows that include checkpoints and human review to mitigate these risks. Rethinking workflows only represents a preliminary step. Agentic AI requires that we re-examine our entire risk-management approach.

Why the Distinction Matters

Lawyers have always relied on tools that increase efficiency. Dictation machines, word processors, legal research databases, and document assembly software all changed how legal work gets done. Agentic AI crosses a line that earlier tools did not.

When a word processor formats something incorrectly, you notice and fix it. When a research database returns irrelevant cases, you refine the search. When you discover that an agentic system makes a mistake, it may already have done something with consequences—sent information to the wrong person, filed the wrong document, updated a system incorrectly, or taken other actions on its own. That difference changes how lawyers must think about supervision, accountability, and risk management of AI. That supervisory role brings with it obligations that resemble those associated with junior staff, outsourced services, or automated filing systems—but with one critical twist: agentic systems operate faster, at greater scale, and with fewer natural pauses than those of us with natural intelligence. What happens when Agentic AI acts faster than we can supervise?

Remember, when something goes wrong, the responsibility does not rest with the AI; it rests with the lawyer or the firm that deployed it.

What Each Tool Does Well

GenAI shines at tasks involving language. It can help lawyers brainstorm arguments, draft correspondence, summarize complex material, and explain legal concepts in ways clients can understand. It works particularly well as a conversational partner, allowing exploration without commitment. We confess to having many conversations with our AI (not all of which were pleasant).

Agentic AI excels at process-oriented work. It handles multi-step workflows that follow defined rules. It can move information from one place to another, monitor systems over time, and coordinate tasks that would otherwise require repeated manual intervention.

When used appropriately, agentic systems can reduce administrative burdens. Misused, they can magnify small mistakes into significant problems—faster than you can say “algorithm.”

What Neither Tool Does Reliably

Neither ChatGPT nor Agentic AI understands professional responsibility as we do. They don’t recognize privilege instinctively. They can’t grasp the reputational consequences of a poorly timed email or an ill-considered filing. They do not exercise judgment about what they *should not* do. They also do not provide much support for accountability.

Agentic systems: they often require access to credentials, tokens, and permissions to function. Once granted, we must manage that access carefully. A system that can read email can often send email. A system that can retrieve documents may also attach them. Every permission expands the potential impact of an error or misuse. Many (most) firms have not considered, let alone properly prepared to handle, that level of supervision. It makes us nervous to think about the consequences.

Risk: Similar Roots, Different Scale

ChatGPT and agentic AI share certain core risks. Both can hallucinate. Both can mishandle sensitive information if used carelessly. Both can perform

"confidently," even when relying on hallucinated "facts." The AI's strong presentation, combined with a natural human indolence, encourages our over-reliance on AI, especially when the output looks polished and professional.

Agentic AI adds a new layer of risk by coupling GenAI's weaknesses with autonomy. A hallucinated fact in a draft memo can prove embarrassing. A hallucination instruction that triggers an automated action poses a potentially far more serious operational danger.

Agentic AI, like GenAi, can work for good or for evil. Consider the social engineering dimension. GenAI can generate convincing, but false messages. Agentic AI systems can generate *and send* compelling, but false, messages at scale. That capability creates obvious concerns about fraud, impersonation, and manipulation—concerns serious enough that OpenAI and other providers have issued specific policy guidance about agent-based misuse.

As if that lacked sufficiency, Agentic AI complicates audits and defensibility. Lawyers need the ability to explain what happened, why it happened, and who approved of it. That requires logging, transparency, and human checkpoints—features we must learn to design into these systems from the beginning.

Risks and Ethical Considerations

GenAI creates risks around data privacy, misinformation, and overreliance. Depending on the platform and configuration, *GenAI tools may retain and use user-submitted dat***a**, particularly in consumer (or in "free") versions. Many enterprise and professional deployments now allow data-use opt-outs or contractual prohibitions on training. Always look for an opt-out option and take it, but do not rely on that. Anonymize all information you upload to AI entities, and encrypt the data you store with password protection to protect confidentiality in the event of inadvertent mistransmission.

GenAI can fabricate citations or misinterpret statutes, creating ethical violations and malpractice risks if unchecked.

Agentic AI carries all the risks of GenAI and introduces autonomy risks. When a

system acts without direct oversight, unintended consequences can multiply. Security vulnerabilities, such as prompt injection attacks, can compromise workflows. Compliance becomes more difficult due to the complexity of auditing autonomous decisions. Ethically, attorneys face accountability questions. Who is responsible if an AI agent makes a mistake? ABA Model Rule 1.1 requires lawyers to maintain technological competence, which now includes understanding AI in its various permutations.

We've found that technical errors do not represent the most significant risk—it's the temptation to trust a system that sounds authoritative but may be confidently wrong.

Client Confidentiality and AI Risks

Client confidentiality remains a cornerstone of legal ethics. AI tools pose a challenge to this principle because they process sensitive data and may misuse it. Agentic AI compounds the risk by accessing multiple systems and autonomously transmitting data.

Economic Impact on Law Firms

AI changes the economics of legal practice. GenAI reduces drafting time. Agentic AI automates administrative tasks, cutting overhead costs. These efficiencies can increase profitability, but they can also disrupt traditional billing models. Flat fees and subscription services may ultimately replace hourly billing to address these issues.

Technology investments require capital. Smaller firms may struggle to afford advanced AI systems. This gap could widen inequality in legal services, creating competitive pressure.

Regulatory Landscape and Compliance Challenges

Regulators worldwide continue to debate AI governance. The EU's AI Act imposes strict requirements on high-risk AI systems. U.S. states explore transparency and accountability rules. At some point, the federal government may (and should) act in this area. Attorneys must monitor these developments because they affect client counseling and firm operations. Compliance failures can lead to sanctions, reputational harm, economic damages, and malpractice claims.

AI in Litigation Strategy and Risk Management

Attorneys can use GenAI to draft persuasive arguments, summarize case law, and identify trends in judicial decisions. Agentic AI can analyze large bodies of prior rulings to identify statistical patterns and correlations that may inform strategy. Imagine an AI agent that reviews discovery documents, flags privileged material, and suggests deposition questions—all without constant supervision. This level of autonomy saves time but introduces additional risk. Risk management becomes essential. Attorneys must maintain ultimate responsibility for strategic decisions, ensuring compliance with ethical standards and procedural rules.

Attorneys must remain the final decision-makers, no matter how advanced the technology becomes.

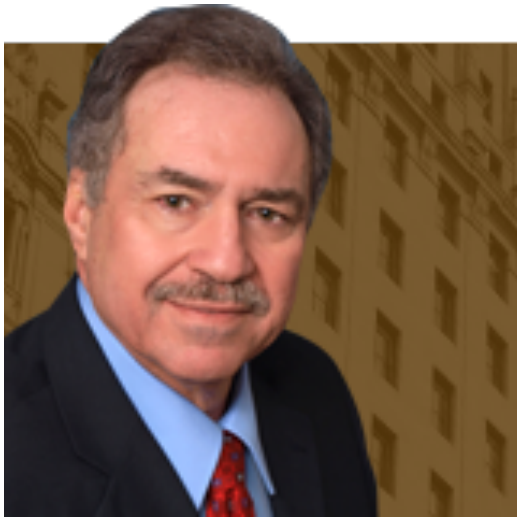
Conclusion

For most attorneys today, the safest and most productive use of AI remains **assistive, not autonomous**. ChatGPT and similar tools can provide real value when used as drafting aids, explanatory tools, and brainstorming partners—so long as lawyers verify the output and retain control. Agentic AI requires a more cautious approach. While potentially valuable for tightly controlled internal workflows, particularly where actions remain reversible and subject to review, it should not operate unsupervised in client-facing or high-stakes contexts. Those safeguards include limited permissions, approval gates before taking external actions, clear audit logs, and an easy way to stop or override the system when something looks wrong. We need to treat agentic AI less like software and more like a junior staff member who never sleeps and never improvises wisely.

GenAI and agentic AI do not compete; neither serves as a replacement for the other. They present different tools built on the same foundation, each designed for a different role. These impressive technologies come with risks. We can manage those risks if attorneys understand what these systems do, where they fall short, and how and why autonomy changes everything.

We're genuinely excited by what AI can do for us personally and professionally. But every time we see a new tool, we remember the lessons learned from early missteps. Our advice? Treat these systems as helpful partners but never let go of the steering wheel. The best outcomes result from technology and human judgment working side by side.

Authors



Jeffrey M Allen

Graves & Allen

Jeffrey Allen is a principal in the law firm of Graves & Allen, in Oakland, California. He runs a general practice that, since 1973, has emphasized real estate and business transactions, receiverships and related...



Ashley Hallene

Ashley Hallene is an Attorney and Land Acquisition Specialist with Demeter Land Development. She is based in Houston, Texas. She is the co-author of several books, including Technology Tips for Seniors Volume 2.0 (2018),...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

American Bar Association |
https://www.americanbar.org/groups/senior_lawyers/resources/experience/2026-april-may/what-attorneys-need-know-about-agentic-ai/

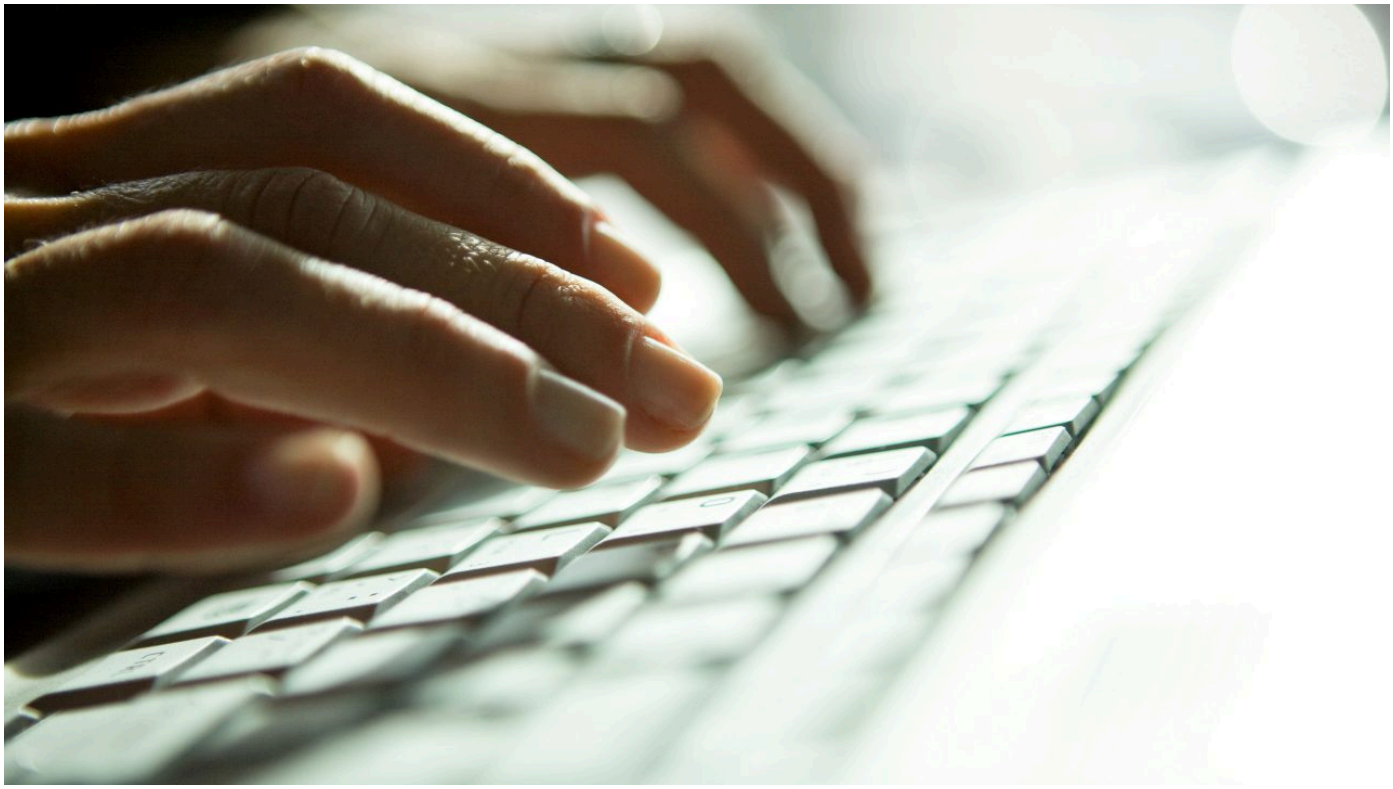
AI and You: An Artificial Intelligence Primer

Jeffrey M Allen

Jun 20, 2024 11 min read

Summary

- Welcome to *GPSolo eReport*'s new column, which discusses the latest developments in artificial intelligence. This month, we present an overview of AI, complete with a glossary of common AI terms.
- Many of the new programs being offered for attorneys to assist in litigation, document control, drafting, legal research, etc., have some version of generative artificial intelligence (GenAI) at their core.
- We need to ensure the ethical and responsible development of AI. This includes creating transparent algorithms, avoiding bias, and ensuring human oversight.



webphotographeer via Getty Images

GPSolo eReport's Editorial Board has recognized the growing importance of artificial intelligence (AI) to our professional and personal lives. They have asked us to create a new AI column that will appear in each issue of the *eReport* to help members of the Division better understand AI, what it does, and what problems it creates. For the first article in the series, we thought that a primer on AI would serve readers best.

Artificial Intelligence Basics

AI has gained significant prominence in recent years. It has started to reshape industries, improve our daily lives, and raise important ethical questions. This primer will summarize AI, its history, key concepts, current applications, challenges, and prospects. We will touch upon several topics in this primer that future columns will explore in greater depth.

AI refers to the capability of a machine or computer program to perform tasks that typically require human intelligence. Such tasks include problem-solving, learning, understanding natural language, recognizing patterns, and data-based decision making.

We have had and used AI in our lives for many years, during which AI's capabilities have improved and expanded, causing substantial but incremental growth in our use of and reliance upon AI, both in our professional and personal lives. You probably have seen this in the evolution of the software offered for searching case law and statutes. The software has grown more and more sophisticated as AI has evolved. As a result, we have increasingly used it to search for specific information relevant to our cases. Examples of AI in our personal lives include virtual assistants, such as Siri, Alexa, and Google Assistant, which employ AI to understand and respond to voice commands.

ChatGPT

In 2022, the tech company OpenAI introduced a program called ChatGPT. ChatGPT marks a significant turning point in the evolution of AI, much like the movement of life from the sea to land marked a significant turning point in the evolution of life as we know it. ChatGPT represents a different form of AI. The development of ChatGPT rapidly accelerated the growth of AI, ramping up its significance to us personally and professionally and opening the door to many new applications. We call this new form of AI "generative AI" (GenAI) as it can generate things such as text, music, diagrams, pictures, and much more.

ChatGPT is an artificial intelligence chatbot that uses natural language processing to create humanlike conversational dialogue. ChatGPT creates humanlike images, text, or videos in response to prompts or instructions we provide. The acronym "GPT" stands for "Generative Pre-trained Transformer," referring to how the program processes requests and formulates responses. ChatGPT trains with reinforcement learning through human feedback and reward models that rank its responses. This process allows the augmentation of ChatGPT with machine learning to help improve future responses. The language model responds to questions and can compose written content of a variety of types for personal and professional uses, including articles, social media posts, essays, code, and emails.

Two recent advances have played a critical part in GenAI going mainstream: transformers and the breakthrough language models they enabled. "Transformers" (*not* the movie version) are a type of machine learning that allows researchers to train increasingly large models without labeling all the data in advance. New models train on billions of pages of text, resulting in answers with more depth and comprehension. Transformers unlocked a new notion called "attention," which enables the ability to track connections between words across pages, chapters, and books rather than only in individual sentences. Transformers can also use their ability to track connections for scientific analysis of such things as DNA.

Rapid advances in large language models (LLMs) having billions or even trillions of parameters opened the door for GenAI models to write engaging text, paint photorealistic images, and even create somewhat entertaining sitcoms on the fly. The rapid evolution of the capabilities of ChatGPT has expanded the use of AI, changing the nature of the beast and its role in our society. The evolution of AI in the wake of ChatGPT has created a tsunami of new applications and put AI in our homes, our cars, our health and medical care, and our offices. For example, in health care, AI aids in diagnosing diseases, predicting patient outcomes, and discovering drugs. In finance, AI

can assess risk, detect fraud, and monitor high-frequency trading. In transportation, self-driving cars and AI-based traffic management systems may improve safety and efficiency.

Many industries have started considering the trade-off between human workers and ChatGPT-powered AI. Some have actually replaced human workers in some positions with AI-driven bots.

To give you some perspective, the move from strictly relying on natural intelligence (that's what we have and use without the help of computers) to using AI could be compared to the move from walking to riding a bicycle. Similarly, the evolution of AI before GenAI could be compared to the evolution from the most basic bicycles (effectively a wooden scooter with no pedals) in the early 1800s to the first commercial car (a one-cylinder Winton, in 1898). Continuing the analogy, from that point, the evolution from pre-generative AI to GenAI would equate to the evolution from the Winton to the modern Ferrari.

AI's Future

AI continues to improve and to learn and apply new skills as it improves. Nothing suggests that its rate of growth will diminish in the near future. An early version of ChatGPT (ChatGPT-3) took and passed law school exams without ever attending a day of law school. It did not, however, pass the bar exam administered to it. But many people who finish law school also fail the bar the first time they take it. A later iteration, ChatGPT-4, however, successfully passed the Uniform Bar Exam with a very respectable score of 297, which bested 90 percent of the human examinees who went to law school. ChatGPT is not focused exclusively on the legal profession. It also passed the CPA exam on its second try and has passed or nearly passed all three components of the U.S. Medical Licensing Exam. Does this mean we will soon have bot lawyers, bot doctors, bot accountants? Does it mean that we will create separate licensing categories for artificial intelligence in these areas? Does it mean we will not bother to license the bots and simply let them work as assistants to the various professions? Time will tell whether licensing boards will license a bot or restrict licenses to those of us limited by natural intelligence (which may be better in some respects but certainly much slower than GenAI).

Collaboration

Expect to see increased collaboration between humans and AI, with AI bots assisting in many professions, including law, accounting, and medicine. In truth, we have already seen bots doing significant work in all those fields and many others. Many of the new programs being offered for attorneys to assist in litigation, document control, drafting, legal research, etc., have some version of GenAI at their core. In fact, many would not exist but for the capabilities created by GenAI and ChatGPT.

We also will see bots supported by GenAI playing a bigger role in our personal lives. Expect to see many doctors using AI as part of our health care, accountants using AI to help prepare our taxes, and our cars using AI to drive us around, likely more safely than we drive ourselves, as natural intelligence has a greater susceptibility to distraction than AI. (The last of these raises an interesting question: If the car is driving itself and you are in the back seat legally intoxicated, are you operating a vehicle under the influence, or does the AI driving the car count as a designated driver?)

We have started to see the evolution from limited but useful virtual assistants such as Siri and Alexa toward far more sophisticated and capable virtual assistants, such as Microsoft's Copilot, which is frankly amazing, even to a jaded technophile like me. If you have not yet experienced Copilot, take a good, hard look at it and what it can do to help you. You can test it out with the free version, but it has some serious limitations, the most important of which is the lack of interface with Microsoft Office. I chose the Pro version ($20 per month) to get its extra features, including the interface with the Microsoft Office Suite. The creation of assistants such as Copilot does not render Siri or Alexa obsolete—yet—as they currently focus on different things. But I predict that those earlier

assistants are nearing end-of-life as they now exist and will either evolve into something more complex and capable or go the way of the dodo as more modern bots replace them using increasingly sophisticated GenAI.



Governance

AI, like most of technology, comes as a double-edged sword, bringing risks and benefits to us. Because of its power, both risks and benefits have increased significance. Governments and organizations have started developing guidelines and regulations for AI to ensure safety, fairness, and accountability in its deployment. Those processes remain in their infancy and do not yet offer any significant protection. More regulations will come in time, but for now, we are in the wild, wild west of the 1800s.

Ethics

We need to ensure the ethical and responsible development of AI. This includes transparent algorithms, avoiding bias, and ensuring human oversight. Ethical issues related to AI should concern us at all levels of personal and professional use. AI has no moral compass and no sense of right or wrong. It is not immoral; it is truly amoral. This should particularly concern us in connection with its use as a legal assistant or when entrusting it with confidential, personal, medical, or financial information. AI has facilitated several attorneys running afoul of their legal and ethical obligations and the laws under which they operate as attorneys. GenAI actually writes as well as or better than most attorneys I have encountered in my now 50-plus years practicing law. Unfortunately, nobody told it not to make up cases or misapply cases, citing them for propositions they do not support.

Some attorneys have had AI bots do their research and draft briefs or memoranda of points and authorities and then turned them in to the court without verifying the legitimacy of the references or the accuracy of the citations. As if that was not a poor enough example of natural intelligence, some of these attorneys also had the bright idea of misrepresenting to the court what they did, so they would not get into trouble. We should learn several lessons from these examples: (1) as an attorney, you need to take responsibility for your actions; (2) never lie to the court; (3) never file or lodge with the court something generated by AI without first having natural intelligence verify the accuracy and legitimacy of every citation. In the words of President Reagan, “trust, but verify.”

There is also the matter of protecting the confidential data entrusted to us by our clients. AI systems collect and analyze vast amounts of data. We need to ensure the safeguarding of such data, which may include personal and case-related information you input about your clients. FYI, the U.S. House of Representatives has banned Microsoft Copilot from all House staff’s computers due to the concerns of the Office of Cybersecurity that Copilot might provide confidential House data to unexpected and non-House approved cloud services.

Other Ethical Considerations and Dangers

In addition to issues of legal ethics noted above, the use of AI also raises more general concerns:

- **Fairness.** AI can inherit biases intrinsic to data on which it trains. Those biases can influence the outcomes selected by the AI model.
- **Automation and job loss.** Automation of tasks through AI will likely lead to job displacement in some industries. We need to prepare the workforce and address what effects that might have on our economy and our social structure. As a society, we need to consider what we will do about the displaced workers, their potential futures, and economic security.
- **Existential risks.** AI's increasing power and flexibility create concerns about the potential for misuse or unintended consequences that could pose existential risks to our society and our population. Industry and governments have raised these concerns and the need for regulation to mitigate these risks. Doing so will require responsible AI development and the implementation of effective monitoring and governance.

Glossary of AI-Related Terms

I do not intend this glossary to be an all-inclusive list of every term you might encounter in connection with AI. It includes the terms you will most likely encounter in articles and discussions about AI.

Algorithm. A step-by-step set of instructions or rules for solving a specific problem or performing a specific task. AI uses algorithms to train models and make predictions.

Artificial intelligence (AI). Artificial intelligence refers to the development of computer systems capable of performing tasks that typically require human intelligence.

Autonomous. We describe a machine as autonomous if it can perform its task or tasks without human intervention.

Big data. Large and complex datasets not easily managed or analyzed with traditional data processing tools. AI often leverages big data for training and decision-making.

Chatbot. A program designed to communicate with people through text or voice commands in a way that mimics human-to-human conversation.

Data mining. Seeking and discovering patterns, trends, and insights in large datasets.

Deep learning. A subfield of machine learning using artificial neural networks, specifically deep neural networks, to model and solve complex problems, often achieving state-of-the-art performance in tasks such as image and speech recognition.

General (strong) AI. AI with human-level intelligence capable of performing any intellectual task a human can perform. True general AI remains a long-term goal.

Generative AI (GenAI). Artificial intelligence systems that create new content such as text, images, audio, and video. GenAI models data to build its own representations, allowing the creation of brand-new, realistic, and often highly customized outputs.

Hallucination. Hallucination refers to outputs generated by AI systems that are fabricated rather than grounded in reality. For example, a text generator could hallucinate fictional events and present them as fact.

Large language model (LLM). Large language models (LLMs) are a subset of generative AI specifically designed to process, understand, and generate humanlike text. They're based on deep learning techniques and trained on massive datasets, usually containing billions of words from diverse sources such as websites, books, and articles. This extensive training enables LLMs to grasp the nuances of language, grammar, context, and even some aspects of general knowledge.

Machine learning (ML). A subfield of AI that focuses on developing algorithms and models enabling computers to improve performance on a task through experience, without explicit programming.

Mathematical model. An abstract description of a concrete system using mathematical concepts and language.

Metrics. Tools used to evaluate the performance of supervised learning models in tasks such as classification and regression.

Narrow AI. Also known as weak AI, narrow AI is designed for specific tasks, such as virtual personal assistants (Siri, Alexa), recommendation systems, and chatbots.

Natural language processing. A field of AI focusing on enabling computers to understand, interpret, and generate human language.

Neural networks. Computer models inspired by the structure and function of the human brain, consisting of interconnected nodes (neurons).

Prompt. A prompt provides context to generative AI systems to produce a desired output. A prompt may be a text description, an image, or even samples of desired audio or video. Carefully engineered prompts guide the creative process while leveraging the model's knowledge.

Reinforcement learning. A form of machine learning in which bots, virtual assistants, or other GenAI creatures learn to make decisions by interacting with an environment and receiving rewards or penalties.

Sentiment analysis. Analyzing text to determine the sentiment or emotional tone; often used in social media monitoring and customer feedback analysis.

Supervised learning. A form of machine learning in which the model trains on a labeled dataset, learning to make predictions or decisions based on input data.

Unsupervised learning. A form of machine learning in which the model learns patterns and structures in data without labeled examples, often used for clustering and dimensionality reduction.

Parts of this article are based on "A Primer on Artificial Intelligence," written by Jeffrey Allen and Ashley Hallene for the ABA Senior Lawyers Division periodical Voice of Experience, January 2024.

Author



Jeffrey M Allen

Graves & Allen

Jeffrey Allen is a principal in the law firm of Graves & Allen, in Oakland, California. He runs a general practice that, since 1973, has emphasized real estate and business transactions, receiverships and related...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

ABA American Bar Association | https://www.americanbar.org/groups/gpsolo/resources/ereport/2024-june/ai-you-artificial-intelligence-primer/

GPSolo eReport GPSolo eReport May 2026

AI and You: How Agentic AI Changes Lav Practice - From Assistants to Actors

Jeffrey M Allen

Apr 29, 2026 12 min read

Summary

- Once AI moves from assistance to execution, the legal risk landscape changes dramati shift demands a recalibration of judgment, supervision, and responsibility.
- Unlike generative artificial intelligence, agentic AI can set goals, plan steps, and execut independently.
- Supervision must start *before* deployment, not after something goes wrong. Unlike hu agentic systems do not raise their hand when uncertain. They proceed
- Where agentic AI is used, it must be tightly constrained, carefully supervised, and desi failure in mind.



MTStock Studio vi

Jump to:

Why the Shift to Agentic AI Feels Different →

Generative AI: A Powerful Assistant With Built-In Friction →

Agentic AI: When the System Stops Waiting for You →

The Risk Multiplier: Speed, Scale, and Silence →

Supervision Reimagined: From Tool User to Process Owner →

Regulatory Trends: AI Governance and Legal Practice →

Ethical Obligations: Navigating AI Risks →

Where Agentic AI Can Add Real Value →

Training, Culture, and the Myth of "Set It and Forget It" →

Conclusion: Autonomy Changes Everything →

Appendix 1: The Solo Lawyer's Practical Guide for Using AI Responsibly →

Appendix 2: The Small Law Firm's Practical Guide for Using AI Responsibly →

For many lawyers, artificial intelligence has quietly crossed a psychological threshold. What once felt experimental now feels ordinary. Generative AI (GenAI) tools draft emails, clean up prose, summarize cases, and reorganize thoughts with impressive fluency. Some lawyers use them daily. Others use them occasionally. A few still avoid them—but even they know that clients, colleagues, and competitors use those tools.

Lawyers have always adapted to tools that promise efficiency. We learn them, use them, supervise them, and eventually stop thinking about them.

Therin lies the danger.

Just as the profession began feeling some comfort with GenAI as a tool in the normal tool kit of practice, a new category of systems has emerged: *agentic AI.* These systems look familiar on the surface. They speak the same language as GenAI. They rely on the same underlying models. Vendors often describe them as "smarter," "more integrated," or "more powerful"—all accurate, but misleading, terms.

Intelligence does not make agentic AI stand out as different. *Agency does.*

Generative AI helps lawyers think and write. *Agentic AI acts.*

Once AI moves from assistance to execution, the legal risk landscape changes—dramatically. That shift demands a recalibration of judgment, supervision, and responsibility.

Why the Shift to Agentic AI Feels Different

Lawyers older than 50 have lived through multiple waves of legal technology disruption. Younger lawyers entered a profession already shaped by many of

those changes. Lawyers have good reason to feel confident in their ability to adapt as technology continues to evolve. But agentic AI does not represent just another iteration of familiar tools.

Earlier technologies changed *how fast* work happened. They did not change *who (or what) did it.*

- Word processors replaced typewriters, but lawyers still authored the work.
- Research databases accelerated access to cases, but lawyers still chose authorities.
- E-discovery systems reduced manual review, but lawyers still defined relevance and privilege.

Agentic AI represents the next evolutionary step for AI. Unlike GenAI, agentic AI can set goals, plan steps, and execute tasks independently. It integrates with external systems, adapts to changing conditions, and completes multistep workflows with minimal human intervention. This autonomy introduces new efficiencies—and new risks. Errors can quickly escalate from mere drafting mistakes to operational mishaps, such as sending confidential documents to unintended recipients. Agentic AI behaves like a digital worker, capable of learning from feedback and adjusting its strategies. It builds on the same language-model foundation as GenAI but adds the ability to plan, decide, and act across multiple steps without continuous human prompting. Do not let its talents mislead you; like its progenitors, agentic AI does not think as we do. Nor is its intelligence true intelligence in the sense that humans have intelligence. We think of human intelligence as a general mental capability that includes the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, and learn quickly from experience. Human intelligence is not merely book learning but rather an adaptive, cognitive function that allows individuals to understand their environment, apply knowledge, and adjust to new situations. While AI has some of those features, it does not have all of them, at least not yet.

In a nutshell, artificial intelligence represents a human-engineered, silicon-based system that processes data using algorithms, excelling in speed, scalability, and pattern recognition without emotional bias, while natural intelligence refers to the biological, carbon-based cognitive capability of humans, characterized by intuition, adaptability, consciousness, and emotional understanding.

Generative AI: A Powerful Assistant With Built-In Friction

GenAI remains, for most lawyers, the safest and most productive place to start. Primarily a language tool, GenAI has limited integration. These tools work *reactively*. They wait for prompts. They respond. They stop.

Used properly, GenAI can:

- Draft and revise language.
- Digest and summarize large volumes of material.
- Translate legal analysis into client-friendly explanations.
- Brainstorm arguments, outlines, and issue lists.
- Organize information for human review.

Genn AI introduces *natural friction*. The lawyer must ask, review, decide, and act. That friction slows mistakes and keeps humans in the loop.

Yes, generative AI hallucinates. Yes, it can come across as confidently wrong. Yes, it requires verification. But its errors typically remain confined to drafts. They do not propagate in legal practice unless a lawyer allows them to. That containment is structural, not accidental.

Agentic AI: When the System Stops Waiting for

You

Agentic AI systems remove that friction by design.

Agentic AI adapts continuously, connects with multiple systems, and executes tasks across platforms. This adaptability makes Agentic AI powerful but also increases the complexity of risk management.

Rather than responding to discrete prompts, you can assign a goal to an agentic system—explicitly or implicitly—and then let the AI determine how to accomplish the goal. That may include:

- Breaking goals into tasks.
- Selecting tools to use.
- Executing actions across systems.
- Monitoring outcomes.
- Adjusting behavior over time.

This can mean reading email, sending messages, moving documents, updating records, scheduling events, or filing materials—often faster than a human can meaningfully supervise in real time.

No longer science fiction, we already have these systems, with more on the way.

This evolution has simple, but profound, implications: *errors no longer stop at the screen.*

The Risk Multiplier: Speed, Scale, and Silence

Lawyers regularly manage risk created by human error: Staff misunderstand instructions, associates miss details, vendors make mistakes. These risks look familiar—and remain manageable—because they come with warning signs.

Agentic AI changes the math.

- **Speed.** Systems act instantly.
- **Scale.** Systems can repeat mistakes hundreds of times.
- **Silence.** Systems do not hesitate unless designed to do so.

A misunderstood instruction can turn into an automated action before anyone notices. A minor configuration error can cascade across matters, clients, and systems.

A correctable draft error in a GenAI context can, when agentic AI is involved, become a confidentiality breach, a missed deadline, or an improper filing.

A human assistant hesitates. An agentic system executes.

Supervision Reimagined: From Tool User to Process Owner

Recognizing that supervision of agentic AI does not look like supervision of people—or even traditional software—makes for a difficult adjustment for lawyers. You cannot simply "keep an eye on it."

Effective supervision requires *design-level decisions*, including:

- What the system may do.
- What the system may not do.
- When the system must stop and ask for approval.
- What actions require logging and review.
- How humans can interrupt or override the system.

Supervision must start *before* deployment, not after something goes wrong. This mirrors the way lawyers supervise automated filing systems, outsourced services, or junior staff—but with less margin for intuition. Unlike humans, agentic systems do not raise their hand when uncertain. They proceed.

Confidentiality and Access: The Permission Problem

Client confidentiality has always required careful handling of access. Agentic AI raises the stakes because access is often *bidirectional.*

- If a system can read email, assume it can send email.
- If it can retrieve documents, assume it can attach them.
- If it can monitor deadlines, assume it can act on them.

Every permission expands the blast radius of a mistake. Lawyers and firms must resist the temptation to over-authorize systems for convenience. Least-privilege access does not simply represent a cybersecurity principle—it constitutes a part of ethical compliance.

Accountability Does Not Automate Away

No matter how autonomous a system becomes, responsibility remains human. Courts, regulators, and clients will not accept "the AI did it" as an excuse. Rules governing competence, supervision, confidentiality, and judgment apply, regardless of whether a human, a vendor, or an algorithm performed a particular task.

This has practical implications:

- Lawyers must understand how systems work, not just what they promise.
- Firms must have the ability to explain decisions made by automated processes.

- Documentation and audit trails become essential, not optional.

GenAI fits that pattern. It felt like an unusually articulate junior associate who never slept but also never acted without instruction. But, like a junior associate, GenAI must be supervised: It cannot guarantee factual accuracy, may hallucinate facts, and must be used with careful verification.

Regulatory Trends: AI Governance and Legal Practice

Global Developments

Regulators worldwide continue grappling with how to govern AI. The European Union's AI Act imposes strict requirements on high-risk systems, including transparency, accountability, and human oversight. U.S. states are exploring their own rules, focusing on transparency, data privacy, and ethical use. The federal government is considering broader action, and attorneys must monitor these developments, as compliance failures can lead to sanctions or reputational harm.

Professional Standards

The American Bar Association (ABA) and state bar associations are issuing guidance on AI use. Attorneys must ensure that AI tools do not compromise client confidentiality, privilege, or other ethical requirements and standards.

Compliance Challenges

Agentic AI complicates audits and defensibility. Lawyers need the ability to explain what happened, why it happened, and who approved it. That requires logging, transparency, and human checkpoints—features that must be designed into these systems from the beginning.

Ethical Obligations: Navigating AI Risks

ABA Model Rule of Professional Conduct 1.1 requires competence, including technological competence. That obligation now includes understanding how AI systems function, where they fail, and how they create risk. Ignorance does not offer a defensible position.

Model Rule 1.6 demands confidentiality. Lawyers must understand what data an AI provider stores, how it processes information, whether it retains information, and whether it uses your information for training. With agentic systems, lawyers must also understand how data moves between connected systems.

Model Rule 5.3 requires supervision of nonlawyer assistants. An agentic AI system, functionally, behaves like a nonlawyer assistant—except faster, more scalable, and without intuition. The duty of supervision does not disappear simply because the assistant is software. And we must not forget the issue of professional judgment.

This raises uncomfortable but unavoidable questions:

- Who bears responsibility when an AI agent makes a mistake?
- How do lawyers ensure informed consent when clients might not understand the technology involved?
- What does reasonable supervision look like when actions occur automatically?

The answers will continue to evolve, but we have a clear direction: *deploying agentic AI without understanding its operational boundaries represents a competence risk.* AI systems do not grasp context. They do not appreciate reputational nuance. They do not instinctively recognize when silence is wiser than response. They do not know when a technically permissible action may not make good sense strategically. Law practice often turns on those subtleties.

There is also a subtle ethical risk we do not discuss enough: over-trust. AI systems

communicate with fluency and confidence. They sound authoritative. That tone can erode skepticism. Lawyers who would never accept a junior associate's draft without review may accept AI output more readily because it "sounds right."

That psychological effect is real. It requires discipline to counteract.

Where Agentic AI Can Add Real Value

Despite the risks, agentic AI is not inherently inappropriate for legal environments. It can be valuable when used in *narrow, well-controlled contexts,* particularly for internal operations.

Examples include:

- Monitoring systems for predefined conditions.
- Coordinating internal workflows.
- Managing low-risk administrative processes.
- Handling repetitive, reversible tasks

The closer an action gets to a client, a court, or an opposing party, the more it requires human review.

Training, Culture, and the Myth of "Set It and Forget It"

One of the most dangerous misconceptions about AI is that once deployed, it fades into the background like earlier tools. Agentic AI does not.

These systems evolve. They interact with changing environments. New data, new rules, and new integrations affect them. Ongoing training and review are essential —not just for the systems, but for the lawyers who oversee them. AI governance

is an ongoing professional obligation.

Conclusion: Autonomy Changes Everything

AI—both GenAI and Agentic AI—offers transformative potential for attorneys. The safest and most productive use remains assistive, not autonomous. Lawyers must verify outputs, retain control, and treat AI as a partner. Agentic AI requires cautious integration, with safeguards, oversight, and clear audit trails. Where agentic AI is used, it must be tightly constrained, carefully supervised, and designed with failure in mind. By understanding the distinctions, risks, and opportunities, attorneys can harness AI's power while protecting clients, firms, and their own professional reputations.

AI provides the promise and the fact of actual assistance. It also brings the promise and the fact of real risk. The profession can manage both—but only if lawyers understand that once technology starts acting on its own, *you no longer have a choice; you must keep your hand on the wheel.*

Appendix 1: The Solo Lawyer's Practical Guide for Using AI Responsibly

This checklist is designed for solo practitioners who want the benefits of AI without sleepwalking into confidentiality, competence, or supervision problems. Treat AI output as a draft, not a source of truth, and assume that any tool you authorize can act faster than you can correct. Start small, document your workflow, and expand only when you can explain—and supervise—what the system is doing.

1. Pick the Right Use Case

- Start with low-risk work: outlines, checklists, client-friendly explanations, and first drafts.
- Avoid high-impact automation at first (sending email, filing, calendaring)

unless you can confirm every step.

2. Choose Tools You Can Explain

- Prefer tools that offer clear terms on data use/retention and a way to disable training on your inputs.
- If you use agentic features, keep permissions narrow (read-only where possible) and require confirmation before any external action.

3. Build a Simple Human-in-the-Loop Workflow

- Use a standard review sequence: verify facts and citations, check names/dates, then check tone and client objectives.
- Save prompts/outputs in the matter file when appropriate (or document key decisions) so you can show what you relied on.

4. Protect Client Information by Default

- Assume anything you paste may be stored somewhere; redact identifiers unless you've confirmed confidentiality safeguards.
- Use matter-specific placeholders (e.g., "Client A," "Opposing Counsel") during drafting, then swap in real names in Word.

5. Track Ethics Guidance and Court Rules

- Keep a one-page "AI notes" file with your state bar guidance, key vendor terms, and any court disclosure/citation requirements you encounter.

6. Write (and Actually Use) a Personal AI Policy

- Decide what you will not do with AI (e.g., no direct client communications or court filings without full review).
- Create a simple data rule: what you can paste, what must be redacted, and what is never entered into any tool.
- Set a review cadence (e.g., quarterly) to recheck settings, vendor terms, and any new features you've enabled.
- Write down what you'll do if something goes wrong (stop use, preserve logs/outputs, notify affected parties as required, and consult ethics counsel/carrier).

Appendix 2: The Small Law Firm's Practical Guide for Using AI Responsibly

This checklist is designed for small firms that want consistent, defensible AI use without slowing their practice. The goal is not central control, but shared standards—so speed does not come at the cost of confidentiality, supervision, or professional judgment.

As firm size increases, the checklist changes. A small firm version shifts from "personal habits" to *repeatable, role-based controls,* for example:

- **Ownership and roles.** Assign who approves tools, who administers settings, and who reviews outputs (even if it's just one partner plus one office manager).
- **Standard workflows.** Define a firm-wide review checklist (citations/facts, confidentiality, tone) and when a second set of eyes is required.
- **Access control by role.** Separate permissions for partners/associates/staff; limit agentic actions; require approval gates for

sending, filing, or sharing.

- **Centralized vendor management.** One intake process for new tools (terms, data retention/training, security questionnaire, insurance implications).
- **Training and consistency.** Short internal training plus a document listing "approved prompts/templates" so everyone uses the tools the same way.
- **Incident response.** A simple playbook (who to notify internally, how to preserve logs/outputs, when to notify clients/court, insurance carrier notice).

1. Define Approved Use Cases (Firm-Wide)

- Identify *approved categories* (drafting, summaries, client explanations, internal checklists).
- Identify *restricted actions* (direct client communications, filings, calendaring, document sending without review).
- Make this a short list—one page, not a policy manual.

2. Approve Tools Centrally (Even If Lightly)

- Designate one partner (or small committee) to approve AI tools.
- Confirm, at minimum:
  - Data retention and training terms.
  - Ability to limit permissions.
  - Basic security posture.

- Maintain a simple *"approved tools" list.*

3. Standardize the Human-in-the-Loop

- Use a *firm-wide review sequence:*

    1. Facts, citations, names, dates.
    2. Confidentiality and privilege.
    3. Tone, strategy, client objectives.

- Require a *second human review* for:
    - Anything client-facing.
    - Anything court-bound.
    - Anything automated beyond drafting.

4. Assign Role-Based Access and Permissions

- Partners: broader drafting access; limited agentic actions.
- Associates: drafting and analysis only; no autonomous actions.
- Staff: no AI tools touching client content unless explicitly approved.
- Agentic features (if any*): opt-in only,* never default-on.

5. Centralize Ethics and Court Guidance

- Maintain one shared "AI Guidance" file with:

    - State bar guidance.
    - Court disclosure/citation rules.
    - Firm decisions on disclosure (when required).

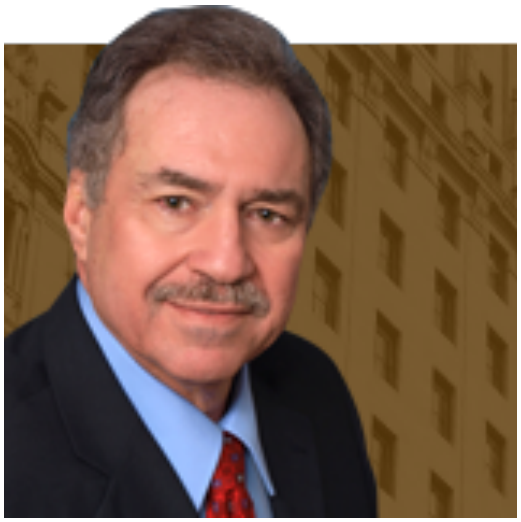
- Update it when rules change—not annually by habit.

6. Create a Firm AI Playbook (Not a Policy Binder)

Include, in plain language:

- What AI is *not* used for.
- What information may never be entered.
- How prompts/outputs are documented.
- What happens if something goes wrong:
    - Stop use.
    - Preserve logs/outputs.
    - Notify supervising partner.
    - Evaluate client/court notice.
    - Notify insurance carrier if appropriate.

Author



Jeffrey M Allen

Graves & Allen

Jeffrey Allen is a principal in the law firm of Graves & Allen, in Oakland, California. He runs a general practice that, since 1973, has emphasized real estate and business transactions, receiverships and related...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

ABA American Bar Association |
https://www.americanbar.org/groups/gpsolo/resources/ereport/2026-may/ai-you-how-agentic-ai-changes-law-practice-assistants-actors/

Can My Computer Get Infected with Malware Just by Surfing the Web?

Wells Howard Anderson and Ashley Hallene

Dec 19, 2023 5 min read

Summary

- Ransomware criminals are targeting legal professionals by poisoning search results on the web.
- Solo and small law firms have modest budgets for cybersecurity and may not have the protection against website landmines.
- Remote browser isolation (RBI) technology can automatically isolate suspicious email links to prevent users from being exposed to potentially malicious web content.
- Aria is an artificial intelligence browser tool created by the browser developer Opera, which can be trained to write in your style.



Steve Cicero via Getty Images

Welcome to the latest installment of our monthly Q&A column, where a panel of experts answers your questions about using technology in your law practice.

This month, we answer readers' questions about how you might be putting your confidential data at risk just by surfing the web, and how to train AI to write documents in your own style.

Q: I've Heard That My Computer Could Get Infected Just by Surfing the Web. What Should I Watch Out For?

A: In a new 2023 campaign, ransomware criminals began targeting legal professionals by poisoning search results on the web. Let's say you search for useful forms using a legal term and click on a top-ranking search result. You risk downloading both the form and ransomware.

Attacking Legal Organizations of All Sizes

The Blackcat ransomware gang floods the Internet with malicious web pages and infected advertisements. They stuff their phony pages with 3.5 million occurrences of targeted search words and phrases. Most of them are legal terms.

The gang effectively poisons legal search results. Their pages rank high in search engines such as Google. So, legal organizations of all sizes become targets, not just the big law firms whose data breaches make the headlines.

For example, clicking on a search result takes a legal professional to what appears to be a forum page. On the page is a download link offering a document relevant to the search term. Clicking that link triggers the malware Gootloader. It brings down a ZIP file containing a hidden JavaScript file. The hidden file launches ransomware or opens a back door to the computer.

Targeting Legal Professionals

Half of the Gootloader attacks strike the legal sector.

Joe Stewart, a principal security researcher at eSentire, observed, "This [is] what I call a landmine approach. They're just mining the entire web with these search keywords and just waiting for somebody in the legal profession, or somebody who needs this legal document, to just stumble on it and open it up. . . ."

The vast majority of files dropped by Gootloader set off ransomware.

Small Firm Risks

Ilia Kolochenko, chief architect at ImmuniWeb, observed that law firms are often small, composed of one or two people, so they lack the cybersecurity knowledge of the larger firms. "Solo practitioners and small law firms are usually poorly protected, having very modest budgets for cybersecurity," said Kolochenko.

Protection Against Malicious Web Pages

Antivirus companies do their best to keep up with the criminals, quarantining known bad files and blocking malicious programs based on behavior. You need antivirus protection to detect those known dangerous files and behaviors. But, sadly, the hackers keep winning. Large numbers of their continually altered files make it through antivirus services to cause harm.

Cloudflare, a web performance and security company, advises, "remote browser isolation (RBI) technology . . . can automatically isolate suspicious email links to prevent users from being exposed to potentially malicious web

content."

RBI integrated into web browsers offers the same protection against both infected search results and phishing emails. If you click on a bad link, whoa!, you see a warning screen and maybe an option to safely view a screenshot of the dangerous webpage.

The features and pricing of RBI products vary. Researchers at the RBI companies are constantly updating and expanding their analytic technologies. They identify never-before-seen threats based on the techniques used by criminals to design their phony websites.

Examples of RBI subscription services include:

- Ericom Zerotrust Web Isolation, Ericom Solutions, Jerusalem, Israel
- Zscaler Browser Isolation, Zscaler, Inc., San Jose, California
- CrimeBlocker, SecureMyFirm Inc., Minneapolis, Minnesota [Disclaimer: This is the author's company]

Techie: Wells H. Anderson, JD, *GPSolo eReport* Contributing Technology Editor and CEO of SecureMyFirm, 952/922-1120, www.securemyfirm.com—we protect small firms from cyber threats with affordable, multiple layers of defense.

Q: How Can I Train AI to Write Documents in My Own Style?

A: Did you know that Aria, an artificial intelligence browser tool created by the browser developer Opera, can be trained to write in your style? You can ask Aria to write emails, tweets, or any other type of text in the specific style you want. Aria's My Style feature allows users to submit to Aria examples of their writing that reflect their tone and voice. The AI then processes them as it would text defined as Shakespeare, Poe, or other authors, giving the AI program the ability to reproduce an individual's natural style for tasks such as reviews or professional emails. If you prefer not to upload samples of your own writing style, you can specify the style that you would like Aria to adopt.

To do this, open your Opera browser. On the left edge of the browser, you will see a column of menu icons; click the Aria icon. In the top right corner, you will find a bluish-purple box with a "+" in the center; click that to start a new chat. At the bottom of the chat window is a box waiting for your prompt. To the left of it is a pen icon. Click on that to adjust the Compose settings. Here, you can define the writing task as a blog post, email, essay, presentation, social post, speech, or article. Next, you will put a description in for the writing task (similar to the prompt that you enter in the main window). Then, you can select the tone you are going for (formal, informal, neutral, academic, business, or funny). The next option is the customized My Style setting. You will see the words Define My Style followed by a "+" icon on the same line. Clicking the "+" opens the My Style window. Here, you can train Aria AI by completing these tasks:

1. Write a formal letter of complaint (to a restaurant, hotel, cinema, etc.) expressing your displeasure in five to ten sentences.
2. Write a product review about an item you recently purchased in four to eight sentences.
3. Write a casual text message to a friend about your plans for the weekend.

Copy and paste the requested number of sentences that reflect your style and click Save. You may have to do this again when you reset the AI, so you should save your answers in a place where you can easily copy them again to

paste them back into the window. You can also set the length of what you are looking for (short, medium, or long.)

When everything is ready, click Generate and be amazed! (But be sure to proofread the resulting text carefully. AIs are known to "hallucinate," making up statements that are false.)

Techie: Ashley Hallene, JD, *GPSolo eReport* Editor-in-Chief (hallenelaw@outlook.com).

What's YOUR question?

If you have a technology question, please forward it to Managing Editor Rob Salkin (robert.salkin@americanbar.org) at your earliest convenience. Our response team selects the questions for response and publication. Our regular response team includes Jeffrey Allen, Wells H. Anderson, Ashley Hallene, Al Harrison, and Matthew Murrell. We publish submitted questions anonymously, just in case you do not want someone else to know you asked the question.

Please send in your questions today!

Authors



Wells Howard Anderson

Wells H. Anderson, J.D., CEO of SecureMyFirm Inc., has over 40 years applying technology to the practice of law as a practicing attorney, independent consultant, and cybersecurity business owner. His business protects solos...



Ashley Hallene

Ashley Hallene is an Attorney and Land Acquisition Specialist with Demeter Land Development. She is based in Houston, Texas. She is the co-author of several books, including Technology Tips for Seniors Volume 2.0 (2018),...

Published by the American Bar Association ©2026. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association.

ABA American Bar Association
https://www.americanbar.org/groups/gpsolo/resources/ereport/2023-december/can-my-computer-get-infected-malware-just-surfing-web/