

DON'T BRING A HAMMER WHEN YOU NEED A SCREWDRIVER:

Choosing The Right AI Tool And Building An AI Policy That Works

By: Crystal West Edwards, Esq., CELA*
Porzio, Bromberg & Newman, P.C.
Copyright 2026

All Rights Reserved

Every week brings a new AI product promising to transform legal practice. For most firms and businesses, the question is no longer whether to use AI, but which tool fits which job and what rules should govern its use. A general-purpose chatbot is not a legal research database, a contract review platform is not a timekeeping system, and a tool appropriate for polishing a newsletter may be entirely inappropriate for a client's confidential file. Selecting the wrong tool, or permitting lawyers and staff to select tools without guidance, exposes the firm and its clients to risks implicating confidentiality, privilege, candor to the tribunal, billing, and supervision.

This paper addresses the scope of AI and common misconceptions, the ethical framework governing responsible AI use, the reasons every firm should adopt a written AI policy, the components that policy should include, and the steps for deployment. It then reviews the major categories of AI products, their practical uses, and the guardrails necessary to maintain client confidentiality and comply with ethics obligations. Whether the reader practices in a large firm, a solo office, or an in-house department, the principle is the same: don't bring a hammer when you need a screwdriver.

UNDERSTANDING THE SCOPE OF ARTIFICIAL INTELLIGENCE

Before practitioners can select the right tool, they need a shared vocabulary. "AI" describes everything from spam filters to systems capable of drafting entire briefs. Think of AI as a toolbox rather than a single tool.

Predictive and Analytical AI

Lawyers who believe they have never used AI have likely been relying on predictive or analytical AI for years. Predictive or analytical AI identifies patterns in data and makes predictions based on those patterns. Spam filters operate on these principles.

Generative AI

Generative AI creates new content (text, images, or code) in response to a prompt. ChatGPT, Claude, Gemini, and Copilot are the products most people have in mind when they refer to AI today. These tools produce content by predicting the most likely next words in a sequence, which explains

both their fluency and their potential for confident inaccuracy. When generative AI produces authoritative-sounding but fabricated content, the result is called a "hallucination."

Agentic AI

Agentic AI is the newest category. These systems plan and carry out multi-step tasks on the user's behalf, such as browsing the internet, completing forms, or sending email. More autonomy means greater need for oversight; agentic tools require the closest supervision of any category discussed here.

Embedded AI

The category practitioners most often overlook is embedded AI., which has features built into software the firm already licenses. Word, Outlook, Zoom, Westlaw, Lexis, and many practice management platforms have added AI functionality, sometimes enabled by default. A firm may not have deliberately adopted AI, yet its lawyers and staff may use it daily. Any AI policy must account for embedded AI as well as standalone products.

SEPARATING MYTH FROM REALITY

Much of the hesitation and overconfidence surrounding AI in law practice stems from persistent misconceptions. Addressing them helps practitioners understand both capabilities and risks.

Myth: "AI Looks Up the Answer"

Generative AI is not a search engine. A large language model generates text that sounds right; it does not retrieve verified information from a database. It can invent cases, citations, and quotations with complete confidence.

Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023), brought this risk to the profession's attention. The plaintiff sued Avianca for injuries when a metal serving cart struck his knee during a flight. When Avianca moved to dismiss, plaintiff's counsel filed an opposition citing judicial opinions inclusive of fabricated quotations generated by ChatGPT. Judge P. Kevin Castel imposed a \$5,000 sanction under Rule 11¹, holding the firm jointly and severally liable. Rule 11(b) provides that by presenting a filing, an attorney certifies that the legal contentions are warranted by existing law or a nonfrivolous argument for extending it. A citation to a nonexistent case cannot satisfy that certification. Courts nationwide have since sanctioned lawyers for the same mistake.²

¹ See Rule 11, Fed. R. Civ. P.

² See *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 448 (S.D.N.Y. 2023)

Myth: "Legal-Specific Tools Don't Hallucinate"

Many practitioners assume that paying for a legal research product solves the hallucination problem. Legal tools grounded in real databases are much better than general-purpose chatbots, but better is not perfect. A 2024 Stanford study found that leading legal research AI tools still produced incorrect or unsupported answers roughly 17% to 33% of the time.³ Purchasing a legal-specific product does not relieve the lawyer of the obligation to verify output.

Myth: "AI Will Replace Lawyers"

AI replaces tasks, not judgment. First drafts, summaries, and document sorting are well suited to AI assistance. Professional judgment, the client relationship, and professional responsibility for the work remain with the lawyer.

Myth: "If We Don't Use AI, We Have No AI Risk"

AI features are now embedded in email, word processing, videoconferencing, and research platforms. Employees may also use free tools on personal accounts without telling anyone, which is called "shadow AI." A firm that has not adopted AI does not have an AI-free practice; it has unmanaged AI.

Myth: "My Chats with AI Are Private"

Consumer versions of many AI tools retain conversation history and may use inputs to train models unless the user opts out. *United States v. Heppner*, No. 25 Cr. 503 (JSR) (S.D.N.Y. Feb. 17, 2026), addressed the privilege consequences. The defendant, a former corporate executive indicted on securities and wire fraud, used the consumer version of an AI platform to generate approximately 31 documents analyzing his legal exposure and potential defenses. He did this on his own initiative, not at counsel's direction, which he later shared with his attorneys. The documents were seized when the FBI executed a search warrant.⁴

Judge Jed S. Rakoff ruled that the documents were not protected. An AI platform is not an attorney and cannot form an attorney-client relationship. The defendant's exchanges with a publicly available platform were not confidential communications with counsel. The documents were not work product because they were not prepared by or at counsel's direction. The decision treats the AI provider as a third party and underscores the risk of entering confidential information into consumer AI tools.⁵

³ Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools, *Journal of Empirical Legal Studies*, 2025; 0:1-27, https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf

⁴ *United States v. Heppner*, Guo, Elizabeth, *Harvard Law Review*, <https://harvardlawreview.org/blog/2026/03/united-states-v-heppner>

⁵ See Heppner

RESPONSIBLE AI USE IN LAW FIRMS

Responsible AI use does not require a new ethical framework. Instead, it requires applying existing rules of professional conduct to new tools.⁶ On July 29, 2024, the ABA Standing Committee on Ethics and Professional Responsibility issued Formal Opinion 512, *Generative Artificial Intelligence Tools*, its first comprehensive guidance on the subject. Opinion 512 maps generative AI use onto the ABA Model Rules.⁷

Competence (Rule 1.1)

Rule 1.1 requires competent representation, including the legal knowledge, skill, thoroughness, and preparation reasonably necessary for the representation. Comment 8 provides that lawyers should keep abreast of changes in the law and its practice, including the benefits and risks of relevant technology. Opinion 512 applies these principles: a lawyer need not be an expert, but must reasonably understand the capabilities and limitations of the tools used, including what a tool does with entered data and how reliable its output is.

Confidentiality (Rule 1.6)

Rule 1.6(a) prohibits revealing information relating to client representation absent informed consent or an applicable exception, and Rule 1.6(c) requires reasonable efforts to prevent inadvertent or unauthorized disclosure. Practitioners must evaluate where client data goes when entered into an AI tool.

Opinion 512 addresses "self-learning" tools (those that may use inputs to inform later outputs). Because information entered by one user could be disclosed in response to a later prompt, the opinion concludes that informed client consent is required before entering client information into such a tool. Consent must actually be informed: "merely adding general, boiler-plate provisions to engagement letters purporting to authorize the lawyer to use GAI is not sufficient." A meaningful conversation with the client about the tool and associated risks is required.

Communication and Fees (Rules 1.4 and 1.5)

Rule 1.4 requires lawyers to reasonably consult with clients about the means of accomplishing objectives and to explain matters as necessary for informed decisions. Depending on circumstances, clients may need to know when AI affects the representation.

⁶ See American Bar Association Model Rules of Professional Conduct, https://www.americanbar.org/groups/professional_responsibility/publications/model_rules_of_professional_conduct/model_rules_of_professional_conduct_table_of_contents/

⁷ American Bar Association, Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512, https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf

Rule 1.5 prohibits unreasonable fees or expenses. Under Opinion 512, a lawyer billing hourly must bill for time actually spent, not the time the task would have taken without AI. If AI reduces a three-hour task to thirty minutes, the bill should reflect thirty minutes. Lawyers generally should not charge for time spent learning a general-purpose tool, and any AI costs passed through must be reasonable.

Supervision and Candor (Rules 5.1, 5.3 and 3.3)

Rule 5.1 requires partners and lawyers with managerial authority to ensure the firm has measures giving reasonable assurance that all lawyers conform to the rules. Rule 5.3 imposes a parallel obligation for nonlawyer assistance. Opinion 512 applies these directly: "[m]anagerial lawyers must establish clear policies regarding the law firm's permissible use of GAI, and supervisory lawyers must make reasonable efforts to ensure that the firm's lawyers and nonlawyers comply with their professional obligations when using GAI tools."

Rule 3.3(a)(1) prohibits knowingly making a false statement to a tribunal or failing to correct one previously made. With Rule 11, discussed in the context of *Mata*, the lesson is straightforward: the lawyer owns every citation filed with the court, and every filing must be verified.

State Guidance

Many states have issued their own guidance. Florida Bar Ethics Opinion 24-1 (January 19, 2024)⁸ concludes that lawyers may use generative AI but must protect confidentiality, provide accurate and competent services, avoid improper billing, and comply with advertising restrictions. The New Jersey Supreme Court's *Preliminary Guidelines on the Use of Artificial Intelligence by New Jersey Lawyers* (January 24, 2024)⁹ emphasize that AI does not change fundamental duties but lawyers must be aware of new technology challenges. Practitioners should review current guidance in each jurisdiction where they practice.

Quality Control and Human Review

Regardless of how capable a tool may be, it must never substitute for critical thinking or produce finalized work product on its own. AI may assist with drafting and analysis, but it is a support tool, not a decision maker. No finalized material should be submitted to a colleague, client, court, or agency without thorough verification and review for accuracy, completeness, and legal sufficiency.

⁸ <https://www.floridabar.org/etopinions/opinion-24-1/>

⁹ <https://www.njcourts.gov/sites/default/files/notices/2026/03/n260330c.pdf>

Attorneys are also responsible for adhering to court or client requirements regarding AI disclosure. Many judges have standing orders on this issue, and many corporate clients address AI in outside counsel guidelines.

Bias and Content Standards

AI-generated content can reflect and amplify biases embedded in training data, including stereotyping, exclusionary language, or overemphasis on certain perspectives. Tools used in hiring, lending, tenant screening, or risk assessment can produce discriminatory outcomes, but the same concern applies to everyday drafting. Prompts must be drafted carefully, lawyers and staff should never use AI to generate offensive or discriminatory content, and all output should be reviewed with a critical eye for bias.

WHY YOUR FIRM NEEDS AN AI POLICY

Whether the practice is a large firm, a solo office, or an in-house department, a written AI policy is quickly becoming a basic expectation for three principal reasons.

First, AI is already in the office. Staff may be using free tools on personal devices without approval. The question is not whether AI is being used, but whether with guidance or without. A clear policy brings "shadow AI" into the open.

Second, the rules require reasonable measures. Rules 5.1 and 5.3 place partners and managers on the hook for ensuring compliance. Opinion 512 specifically identifies clear policies and training as part of meeting that duty.

Third, clients, courts, and carriers are asking. Corporate clients add AI provisions to outside counsel guidelines. Judges have issued standing orders on AI use and disclosure. Malpractice and cyber insurers increasingly ask about AI governance.

A policy also protects the firm less obviously: by giving clear permission to use the right tools, it often stops people from using the wrong ones.

WHAT YOUR AI POLICY SHOULD INCLUDE

The components below are drawn from actual firm policies and the ethics guidance above. Each firm should tailor them to its size, practice areas, and technology environment.

Scope and Purpose

The policy should define what counts as an AI tool, identify who is covered (attorneys, staff, contractors, and vendors), and explain why the policy exists. A well-drafted purpose statement acknowledges both sides: AI platforms can be useful for brainstorming, outlining, stylistic tasks, drafting, summarizing documents, and preparing materials. At the same time, they pose risks to data

security and confidentiality, can result in privilege waiver, and do not always produce accurate or unbiased output.

Approved Tools and the Approval Process

The firm should maintain a published list of tools vetted by IT and security, easily accessible to all. The policy should establish a clear process for requesting and approving new tools, including advance approval for tools carrying additional cost. Many firms also monitor usage and revoke underutilized licenses.

Prohibited Inputs and Data Protection

The policy should specify what may never be entered into an AI tool, and where. General instructions such as "be careful with client data" are insufficient; people need concrete examples. The following language is adapted from an actual firm policy:

Employees are prohibited from entering into AI platforms any information that is sensitive, confidential, privileged, or protected by law, except for in firm-approved GAI Tools that guarantee prompts will not be captured outside the firm's technical boundaries. This prohibition includes any information that contains personal data, i.e., any information that can be linked to an identifiable person (such as name, location, date of birth, an online device identifier, etc.), trade secrets, intellectual property, financial information, or other sensitive data of our clients, our firm and employees, or third parties.

This language identifies both the prohibited categories and the limited exception for firm-approved tools that keep data within the firm's control.

Human Review and Verification

The policy should make explicit that AI assists but does not decide, and that no AI output leaves the firm until a responsible attorney or staff member has reviewed and verified it.

Disclosure and Client Communication

The policy should address when the firm will tell clients and courts about AI use, assign responsibility for tracking court rules and client guidelines, and decide when it will proactively disclose AI use to clients.

Billing for AI-Assisted Work

Consistent with Opinion 512, the policy should require billing for time actually spent, prohibit charging for time learning general-purpose tools, and require disclosure of any AI costs passed through. Firms should also consider whether value-based or flat fees make more sense for AI-assisted work.

Training, Incident Reporting and Enforcement

The policy should require training before access to a tool and tell people what to do if something goes wrong. Reportable incidents include inadvertent inclusion of confidential information in a prompt, or the realization that AI has produced incorrect output used in work product. Incidents and policy violations should be reported immediately to IT and security personnel.

The policy should state that violations may result in disciplinary action, including access limitation or termination. It should make clear that all users are responsible for protecting confidential data and that proper use supports ethical obligations and client trust. The policy should supplement existing policies on confidentiality, privacy, and professional conduct. Finally, the firm should commit to reviewing the policy regularly, because the technology will not stand still.

DEPLOYING YOUR AI POLICY

A policy in a drawer protects no one. Deployment is where most firms succeed or fail. The following roadmap works for firms of any size.

1. **Form an AI governance team.** In a large firm, this may include IT, security, general counsel, risk management, and practice group leaders. In a small firm, it may be two partners and the IT consultant. Either way, include someone who does the day-to-day work.
2. **Inventory how AI is already being used.** Find out what is happening and where the pain points are. A short anonymous survey is often eye-opening.
3. **Define the jobs to be done.** Identify where people lose time (i.e. research, drafting, intake, or billing). That list should drive tool selection, not the other way around.
4. **Vet vendors and pilot.** Start small. Never pilot with live client confidential information until the tool has been approved; use non-client or anonymized data.
5. **Launch with training.** Roll out the policy with training and a published list of approved tools. People need to know what is allowed, not just what is forbidden.
6. **Monitor and update.** Track usage, audit outputs, and revisit the policy at least annually.

Match the Tool to the Task

This is the heart of the matter. Too many firms purchase the tool everyone is talking about, then look for ways to use it. The better approach is to flip that sequence and begin with the problem. If the firm's greatest pain point is time entry, a research platform will not help. If the need is contract review, a general-purpose chatbot is wrong. The best tool is the one built for the job.

Once the job has been defined, vet the vendor and then pilot the tool. This means confirming how data is stored, used, and protected before anyone uploads a single document. A pilot with a small

group on realistic, non-client scenarios reveals whether the tool saves time and how often its output requires correction. Time saved and error rates should be measured before firm-wide rollout.

Vendor Due Diligence

The following questions should be put to every vendor, with answers obtained in writing.

Data Use and Retention. The most important question is whether firm inputs will train the vendor's models. Enterprise agreements usually say no; consumer terms often say yes unless the user opts out. Also ask how long prompts and outputs are retained and whether they can be deleted.

Security and Access. Ask about encryption, independent audit reports (such as SOC 2 Type II or ISO 27001), single sign-on, and access controls. Determine where data is hosted and who at the vendor can access it.

Contract Protections. The agreement should include confidentiality obligations, prompt breach notification, limits on subcontractors, and confirmation that the firm and client own inputs and outputs. If the firm handles health information, a business associate agreement under HIPAA may be required.

One Firm's AI Toolbox

Matching tools to tasks rarely results in a single product. My firm did not choose one AI product; it chose several, each for a specific job:

- **Copilot**, an enterprise AI assistant embedded in Microsoft tools, for secure drafting, summarization, and workflow support;
- **Spellbook**, a contract-focused tool designed to enhance drafting accuracy, identify risks, and improve review efficiency;
- **Harvey**, a legal AI platform supporting research, analysis, and complex drafting tailored to legal professionals;
- **Laurel**, an AI-driven timekeeping system that automates time capture and improves billing accuracy; and
- **Blue J**, an AI-powered tax research tool that analyzes case law and predicts likely outcomes based on fact patterns.

Each tool went through IT and security review, each has its own training, and each appears on the firm's approved list. That is the screwdriver-not-hammer approach in action.

CHOOSING THE RIGHT TOOL: TYPES OF AI PRODUCTS

With the policy framework in place, practitioners can turn to the toolbox itself. For each category below, this paper identifies examples, best uses, and the guardrails practitioners should observe. Product names are examples, not endorsements, and the market changes monthly.

General-Purpose AI Chatbots

General-purpose chatbots, like consumer versions of ChatGPT, Claude, Gemini, and Copilot, are the most accessible AI tools and the most frequently misused. Often free or low cost, they are well suited for brainstorming, explaining a concept in plain language for a client newsletter, outlining a presentation, or tightening a paragraph containing no confidential information.

Treat a consumer chatbot like a public space. Do not enter client names, facts, or documents into a consumer plan. Review privacy settings, including whether conversations are used for training, and never rely on a chatbot for legal authority without independent verification. *Heppner* bears repeating: the court treated a consumer AI platform as a third party. This is a reason to counsel clients about using these tools to analyze their own legal matters.

Enterprise AI Assistants

Enterprise AI assistants, like Microsoft 365 Copilot, ChatGPT Enterprise, Claude for Enterprise, and Gemini for Google Workspace, are licensed to the organization, typically with commitments not to train on organizational data and with administrative controls. They are most useful inside the firm's environment: summarizing email threads, drafting responses, or locating documents.

The key guardrail is permissions. An enterprise assistant can see whatever the user can see. If document management or SharePoint permissions are loose, the AI will expose that weakness. For example, by summarizing a confidential HR file for someone who technically had access but never should have. Clean up file permissions and ethical walls before rollout, confirm that the contract bars training on firm data, and confirm retention settings, audit logging, and which features are covered by the enterprise terms.

Legal Research AI

Legal research tools, like Lexis+ AI, Westlaw with CoCounsel, vLex Vincent, Bloomberg Law, and specialty tools such as Blue J for tax, connect a language model to a curated database of cases, statutes, and secondary sources. That grounding makes them far more reliable than general-purpose chatbots, well suited to finding authority, summarizing case law, and preparing first drafts of research memoranda.

Reliable is not infallible. As the Stanford study demonstrates, these tools can misstate holdings or cite sources that do not support the proposition for which they are cited. The guardrail is simple and non-negotiable: pull and read every authority cited, confirm it says what the tool claims, and confirm it is still good law. Neither Rule 3.3 nor Rule 11 contains an AI exception. Check any applicable standing order on AI use and disclosure before filing.

Drafting and Contract Review AI

Drafting and contract review tools, like Spellbook, Harvey, and the AI features of document assembly and contract management platforms, work inside Word or on a dedicated platform. They can generate first drafts, compare clauses, redline contracts against the firm's preferred positions, and flag unusual or missing provisions.

Because these tools work on client documents by design, they must go through vendor vetting. Depending on how the tool handles data, Opinion 512 may require informed client consent before uploading client documents. The lawyer remains responsible for every word of the final draft; AI can miss context a lawyer would catch, such as a side agreement or the client's real business goals.

E-Discovery and Document Review AI

E-discovery is where AI in law has its longest track record. Courts have accepted technology-assisted review for more than a decade, and platforms such as Relativity aiR, Everlaw, and DISCO now layer generative AI on top. These tools are best used to prioritize large document sets, screen for privilege, and summarize productions.

Guardrails center on defensibility. Validate results through statistical sampling to confirm the tool is finding what it should. Protective orders and ESI protocols may restrict where data can be sent or require disclosure of the review methodology; follow them. Be prepared to explain and defend the process if challenged.

AI Meeting Notetakers and Transcription

Meeting assistants, like Zoom AI Companion, Microsoft Teams recap, and Otter, are spreading rapidly, often arriving uninvited when another participant's bot joins a call. They are useful for internal meeting notes, action items, and dictation.

Three guardrails apply. First, privilege: a third-party service recording a privileged conversation raises questions similar to *Heppner*; use only approved tools, or none, on client calls. Second, recording laws: several states, including Florida, require all-party consent, so announce the recording and obtain consent. Third, retention: transcripts are discoverable records. Know where they are stored, how long they are kept, and incorporate them into records retention.

Practice Management, Billing and Timekeeping AI

Tools such as Laurel, Clio Duo, Smokeball, and the AI features of practice management platforms address the business side of practice. Automated time capture can recover otherwise-lost time, and AI can draft billing narratives and help organize matters and deadlines.

The bill is the lawyer's representation to the client. Review every auto-generated time entry for accuracy and reasonableness before billing, consistent with Rule 1.5, and ensure narratives do not

reveal privileged details if invoices are shared or produced. Deadlines are malpractice territory; AI may suggest dates, but never calendar a deadline without human confirmation.

Client-Facing Chatbots and Intake Tools

Website chat assistants, automated intake forms, and AI answering services can handle scheduling, collect basic intake information, and answer general FAQs. These tools carry unique risks because the AI speaks for the firm to people who are not yet clients.

The tool should not give legal advice (unauthorized practice concerns) and should make clear that no attorney-client relationship yet exists and that the user is interacting with AI. Rule 1.18 provides that information from a prospective client may disqualify the lawyer from adverse representations, so limit what the tool collects and route that information into conflict checks. Anything the tool says is a communication about the lawyer's services; Rule 7.1 prohibits false or misleading communications, and Florida Bar Ethics Opinion 24-1 requires compliance with advertising restrictions. For practices serving vulnerable clients, ensure a human is easy to reach.

GUARDRAILS THAT TRAVEL WITH EVERY TOOL

The guardrails above vary by category, but four principles apply to every tool.

Know Where the Data Goes. Read the terms. Who stores the data, who can see it, and is it used for training? If the practitioner cannot answer, the tool is not ready for client information.

Get Informed Consent When Needed. Opinion 512 is clear: before confidential client information enters a tool that could disclose it, a genuine conversation with the client may be required.

Verify Before You Rely. Every citation, fact, and number must be checked. Every sanctions decision involving AI comes down to someone skipping this step; the signature on the filing belongs to the lawyer.

Supervise People and Tools. Train staff, monitor use, and treat AI the same as nonlawyer assistance under Rule 5.3. The lawyer is responsible for the work that goes out.

Before you use any AI tool, ask yourself: Is this the right tool for this job? Is this tool approved by my firm? Am I entering confidential client information, and can I anonymize it? Do I need the client's informed consent? Does the court or client require disclosure? Who will verify the output before it is used?

Crystal West Edwards, Esq. is a Principal with Porzio, Bromberg & Newman, P.C. in Parsippany, New Jersey. She is Certified as an Elder Law Attorney (CELA) by the ABA Accredited National Elder Law Foundation, a Fellow of the American College of Trust and Estate Counsel (ACTEC), a Member of the Special Needs Alliance, and a former Director of the National Academy of Elder Law Attorneys (NAELA).