Relativity

Understanding the Promise and Limitations of **Agentic AI for Legal**

A Toolkit

Contents

Foreword 3

Section 1: Explaining Agentic AI: Key Terms 4

What is Agentic AI 6

A Note about Autonomy and Self-Direction 6

Tool: A Visual Guide on AI Safety 9

Section 2: Considerations and Ethics for Legal Practitioners **10**

The Responsible Practice of AI in Law 12

Build Responsible AI into Your Practice 14

Tool: Green-Flag Checklist for Evaluating Legal AI Tools 16

Section 3: Practical Use Cases **18**

AI Agents in the Broader World 20

Examples of Agentic AI for Legal Teams 20

Tool: Spot the Agent Reflection Exercise 22

Section 4: How to Get Started **24**

Develop an Agentic Mindset 26

A Word on Safe Sandboxes 27

Leveraging Agentic AI to Evolve Your Practice 27

Tool: A Prompt for Using Agentic AI In Your Professional Development Planning 29

Foreword

Aron Ahmadia

Vice President of Applied Science
at Relativity

Legal teams have moved from asking *whether* AI belongs in their practice to *how* to use it responsibly, defensibly, and at scale. "Agentic AI" is the next step in that journey, and also the most misunderstood. This toolkit is a plain-English guide to what "agentic" really means for legal work today, though you should know definitions vary even among experts.

At Relativity, when we say **aiR is agentic**, we mean it does more than generate text. It **perceives, thinks, and acts** across connected systems to pursue user-defined goals while keeping the right people firmly in control. In practice, that looks like structured multi-step planning, tool use, and orchestration, paired with **human-in-the-loop checkpoints** and **statistical validation** (think precision, recall, and elusion) before results drive outcomes. Agentic does not mean "unsupervised." In legal, it should not.

We also draw a clear line between **workflows** and **agents**. Workflows follow a pre-set path. Agents reason about the path, adapt, and choose tools to get the job done. Agentic tools like aiR should embrace that reasoning capability and embed safeguards like: transparent rationales you can review, controls to tune and test before you scale, and operational guardrails that favor defensibility over flash.

Why now? Because the volume, velocity, and variety of legal data demand systems that can plan and act with your guidance. Teams using agentic approaches are already accelerating first-pass review, issue tagging, data breach responses, and case strategy development, shifting human effort to judgment, advocacy, and client impact.

This toolkit offers three things: shared definitions, pragmatic cautions, and concrete starter patterns. Use it to evaluate claims in the market, align your internal standards, and design pilots that are ambitious and safe.

Our ask is simple: **keep engaging, keep learning, and test agentic AI responsibly**. Start small, measure everything, and insist on human oversight.

Let's cut through the noise and build what works.

Explaining Key Terms (and Their Nuances)

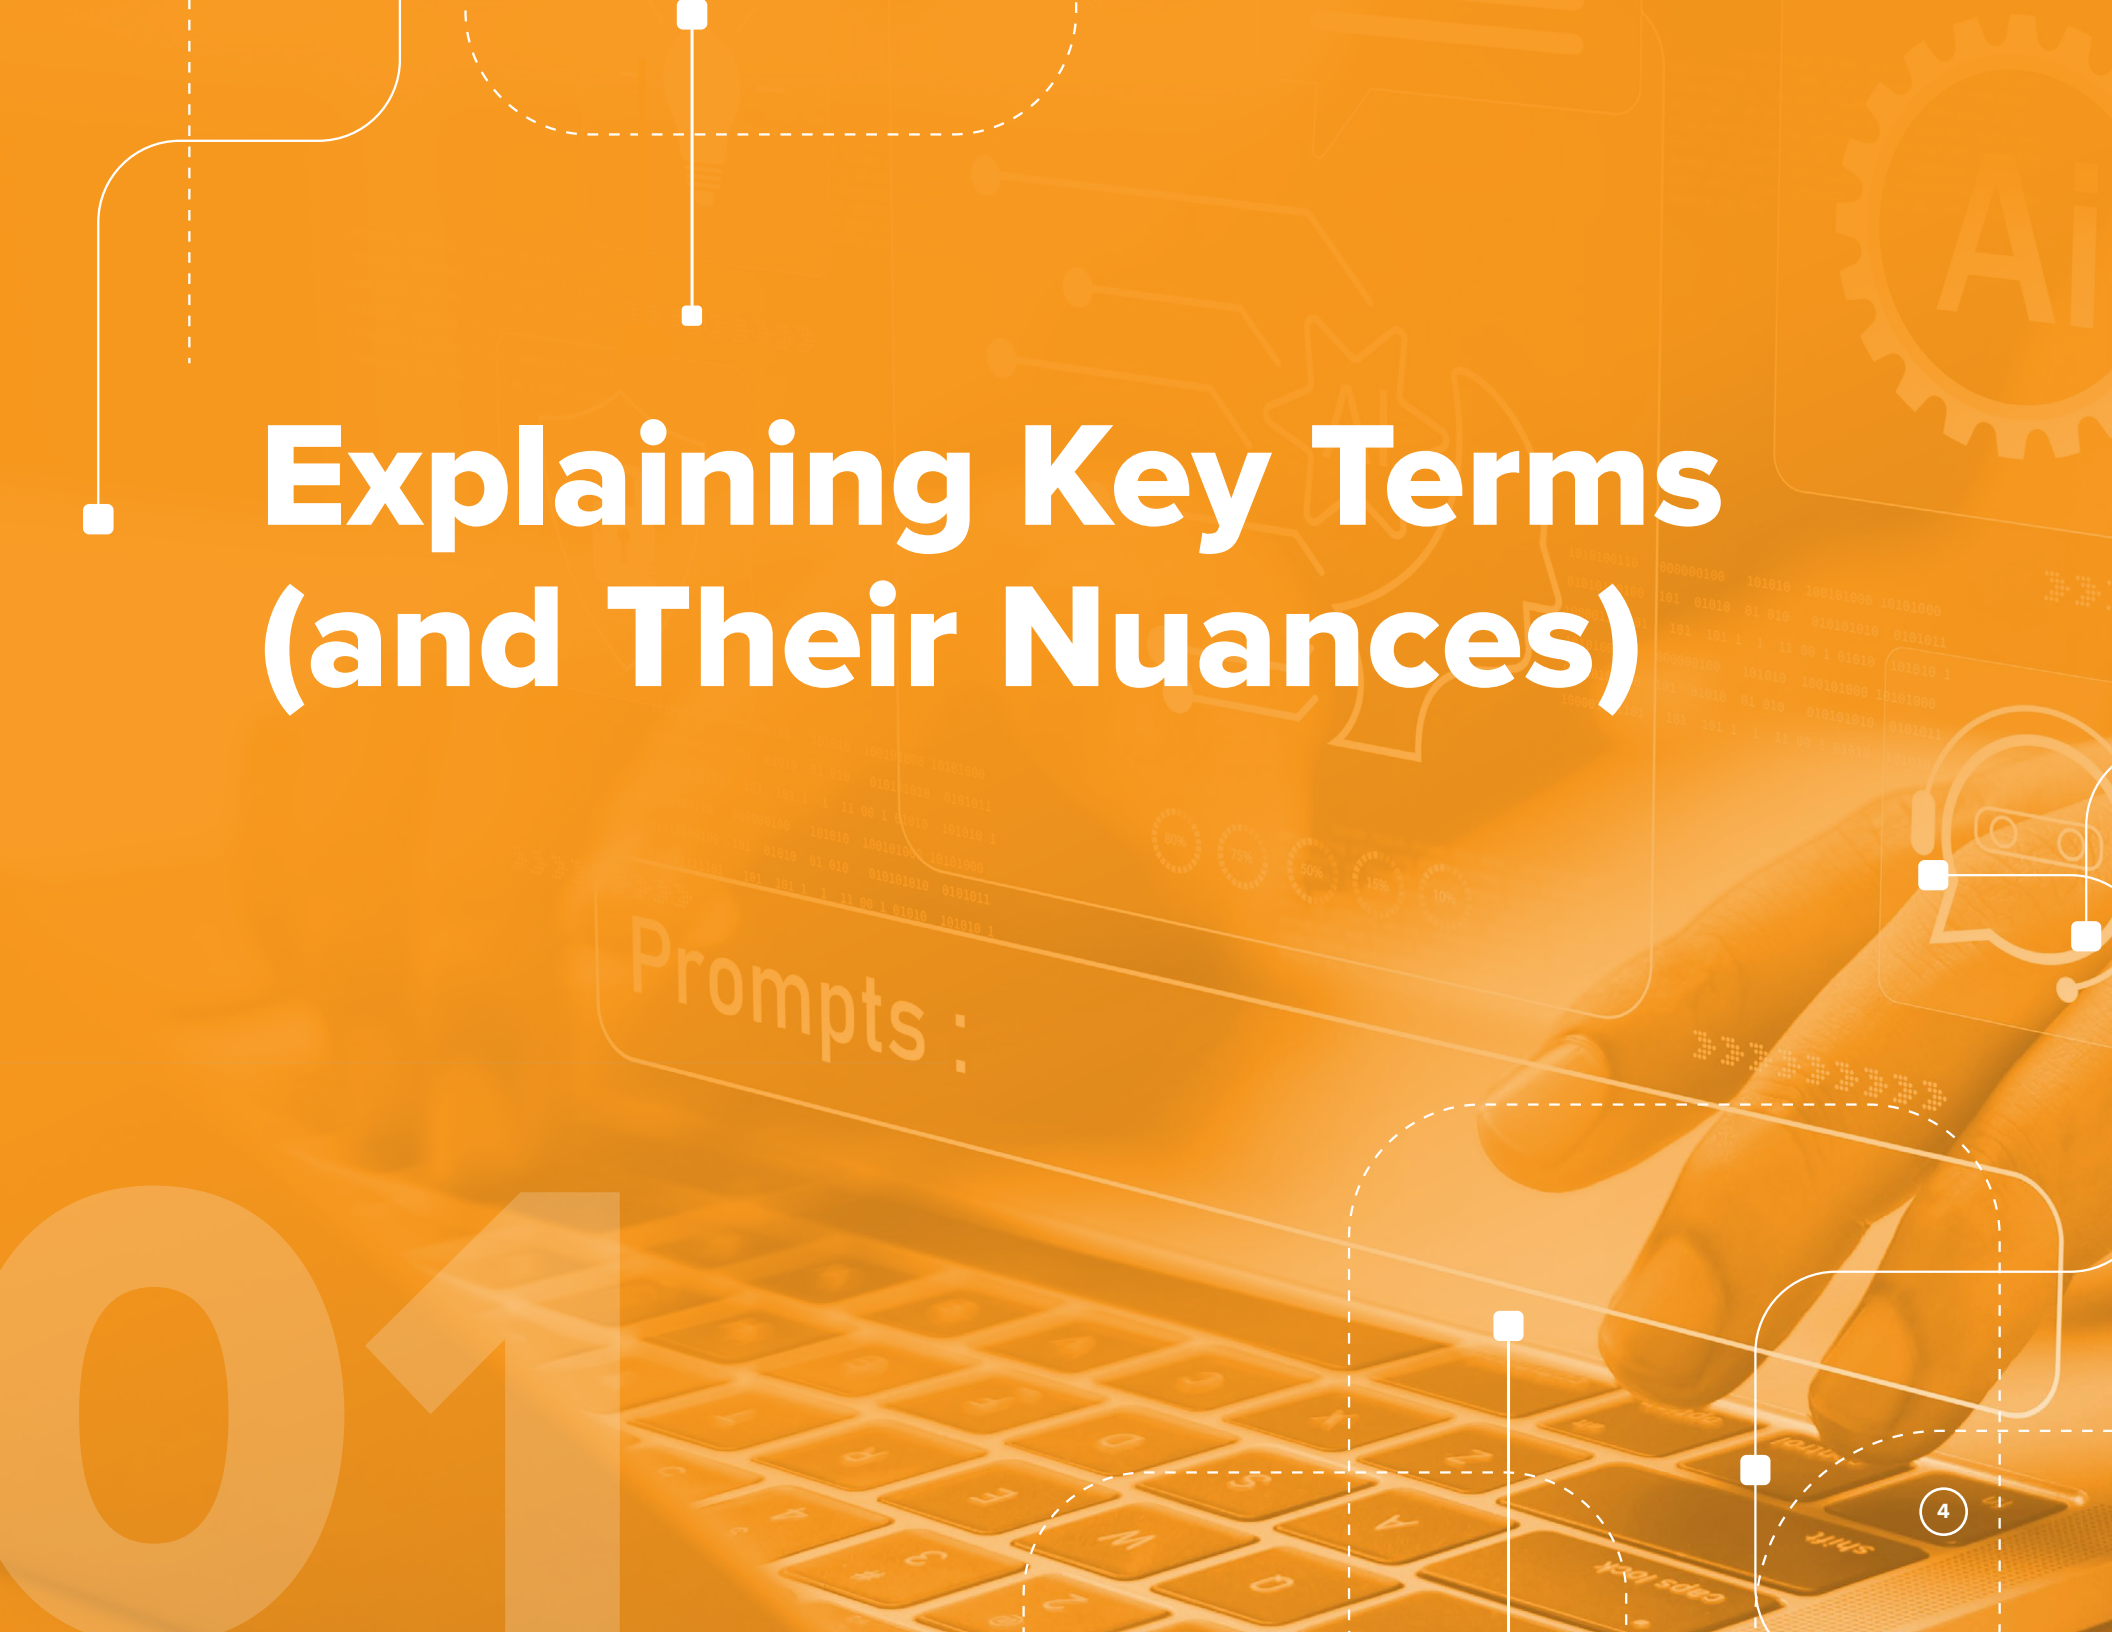

01

tl;dr

What You Should Know

- The AI space is evolving so rapidly, there's no consensus about a hard-and-fast definition of "agents"—even among experts.
- Still, here's a good baseline to start with: "Agentic AI encompasses systems that perceive, think, and act on their own to meet user-set goals."
- AI workflows and AI agents are not the same. One follows a pre-defined path of action; the other uses reasoning to plan and direct itself.
- Agentic AI can be built on a spectrum of autonomy, from entirely human-led to zero human supervision. Legal AI should always have a human in the loop to validate outputs and preserve the integrity of each project.

What You Should Do

- Look for legal AI that hinges upon human supervision and statistical validation—not free-for-all agents.
- Watch for forthcoming materials from Sedona Conference Working Group 13.
- Immerse yourself in AI conversations to start picking up the language: attend conferences, follow the experts' LinkedIn or newsletter content, and open a dialogue with tech-savvy colleagues and peers.

What is Agentic AI?

“If you see this question, ‘what is agentic AI,’ and you feel a little uncertain—you were just getting your head around generative AI, and now we’re already on agentic AI—well, you are in very, very good company,” Ray Mangum, partner at Redgrave, says.

“We’re all trying to get our heads around this. If you follow the news around AI, you’ll see lots of discussion: it’s a bit of a running joke, asking ‘what is an agent or agentic AI?’” he continues. “That said, I don’t think it’s a meaningless term.”

Mangum notes that he’s involved in a working group with the Sedona Conference, and part of their mission is to help set some formal definitions of these terms for the legal technology space. We’re all eager to see them; you can follow along on their page.

In the meantime, most definitions share some similarities: agentic AI does something on your behalf. It doesn’t only answer questions; it takes some action according to your instructions.

Put simply, **agentic AI encompasses systems that perceive, think, and act on their own to meet user-set goals.**

Such a system has access—via APIs or other connectors—to the platforms it needs to work within, whether it’s enterprise software, online resources, or other services. When a user assigns an agentic system a task, it can autonomously follow those instructions in a multi-step way to deliver the requested result.

A Note about Autonomy and Self-Direction

Generally, a large language model is the “brain” of the agentic system, enabling the reasoning skills and natural language processing required to interact with the user, analyze solutions to the prompt, and deliver outcomes.

It’s worth noting that AI workflows can appear to yield similar results as AI agents, but there is a crucial difference: in the case of AI workflows, a predefined path orchestrates an LLM and other tools to accomplish particular tasks, whereas in the case of AI agents, the LLM plans and directs itself.

Both approaches are useful and can deliver excellent outcomes, but when building an understanding of how the underlying technology works and what safeguards it should have, it’s important to grasp the technological difference.

One such safeguard comes in the form of human supervision. An agentic system can be built with several levels of autonomy, ranging from full human control and oversight to entirely self-directed and -completed task execution.

Both approaches are useful and can deliver excellent outcomes, but when building an understanding of how the underlying technology works and what safeguards it should have, it's important to grasp the technological difference.

Aron Ahmadia, vice president of applied science at Relativity, says the concept of fully autonomous AI agents went from a fairly abstract idea to real, in-play technology in less than a year.

"We've seen fully autonomous coding agents pop up left and right in 2025, and that's pretty scary," he tells us.

For the team at Relativity, fully autonomous agents do not belong in legal tech. Chris Brown, the company's chief product officer, noted on *The Relativity Blog*: "We believe it is critical for users to guide, validate, and maintain control over AI's decisions. Maintaining human intervention ensures legal AI remains ethical, practical, and defensible."

From Relativity's perspective, an agentic system built to tackle first-pass document review, issue tagging, case strategy development, data breach response, and similar legal data intelligence projects, should have a human brought into the loop at various points to input and refine prompt criteria, validate outputs, evaluate the agentic system's reasoning, and ensure consistency.

AI Workflows	AI Agents
Predefined paths that orchestrate LLMs and tools	LLM **plans** and **directs** its own tool usage and iteration, controlling its own completion
Predictable and consistent for well-defined tasks	Flexible for open-ended, unpredictable tasks
Common patterns include prompt chaining, routing, and parallel tasks	Common patterns include research loops for question-answering, autonomous agents, and coding agents

Agentic Level	Description	Term	Example Code	Who's in Control?
☆☆☆☆	Model has no impact on program flow	Simple processor	`print_llm_output(llm_response)`	Human
★☆☆☆	Model determines basic program flow	Router	`if llm_decision(): path_a() else: path_b()`	Human: How functions are done; System: When
★★☆☆	Model determines how functions are executed	Tool call	`run_function(llm_chosen_tool, llm_chosen_args)`	Human: What functions are done; System: How
★★★☆	Model controls iteration and program continuation	Multi-step agent	`while should_continue(): execute_next_step()`	Human: What functions exist; System: Which to do, when, how
★★★★	Model creates & executes new code	Fully autonomous agent	`create_code(user_request); execute()`	System

Table 1. Levels of AI Agent: Systems using machine-learned models can have different levels of agency. They can also be combined in "multiagent systems," where one agent workflow triggers another, or multiple agents work collectively toward a goal. Levels adapted from (Roucher et al., 2024).

Source

In short, these human touchpoints lay a crucial framework for validation, transparency, and trustworthiness, while allowing the AI to do what it is designed to do.

“These steps we add to the system do not make it less agentic; they make it safer,” Ahmadia says.

This safety is critical for legal use cases for AI, where highly sensitive data, judicial and regulatory scrutiny, professional liability, and real-world outcomes that impact real people, are all intrinsic components of each case.

Recommended Resources

The best way to learn a new language is immersion. Explore these resources to start steeping.

- Legal practitioners experienced with agentic AI recommended exploring YouTube and other resources for educational content from experts like Andrej Karpathy.
- Watch for forthcoming materials from Sedona Conference Working Group 13.
- Leverage your network to talk with AI experts in your field; it’s an incomparable way to build understanding in the context of what matters most to you and your team. A simple way to start? Attend an upcoming conference and seek out all the AI content you can.

Tool: A Visual Guide on AI Safety

AI isn't one-size-fits-all. This guide shows where common tools fall along the autonomy spectrum—and where legal teams should aim to operate in the interest of safety, defensibility, and client protection.

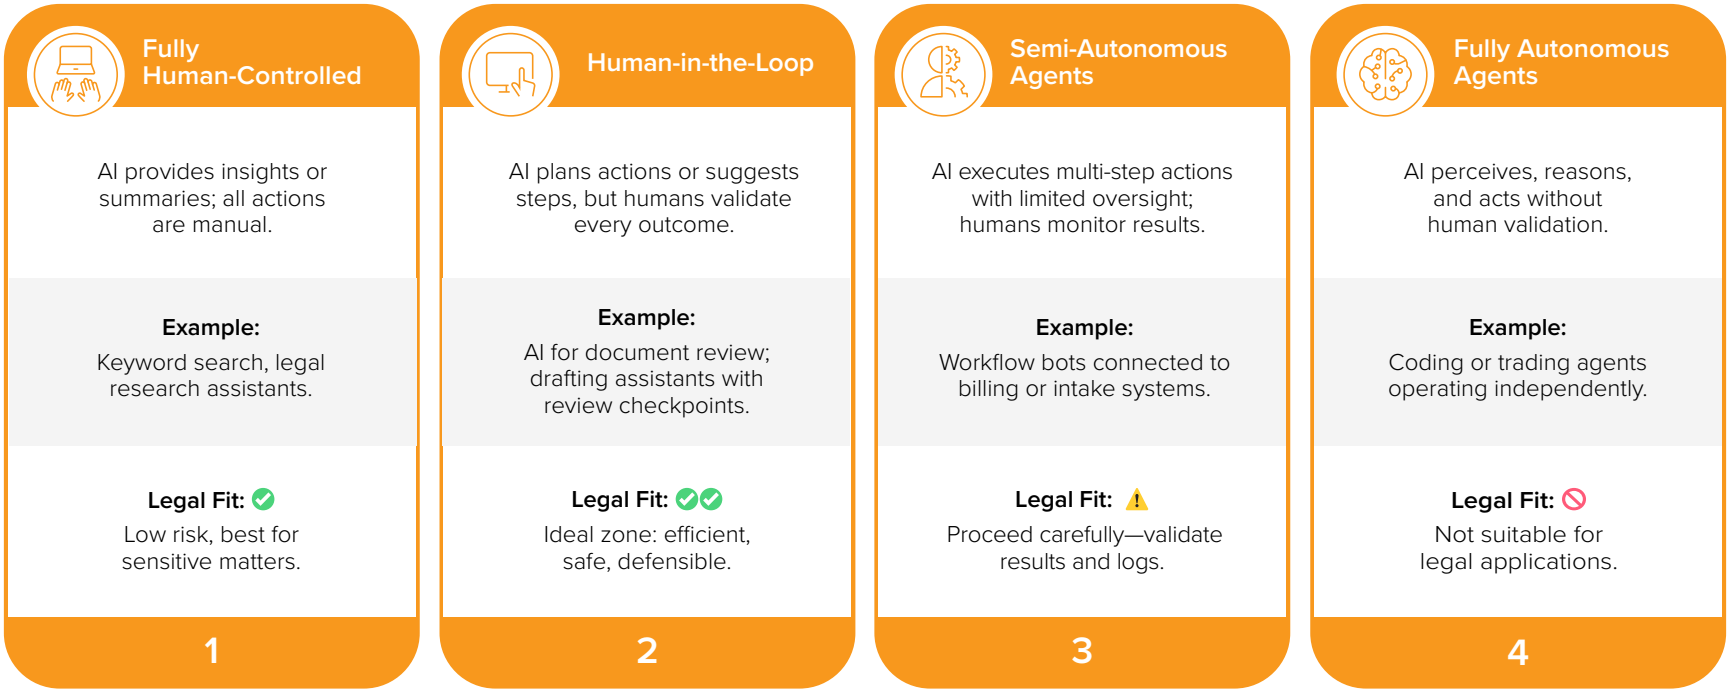

How to Use This Tool

- Consider the spectrum of autonomy and try to discern where a tool you have in mind sits.
- Questions to ask yourself to find that answer might include:
  - Are the outputs of this AI reviewable and reproducible?
  - Does the AI have access to multiple systems? Does it start, advance, and complete multi-step tasks on its own, or does it require prompting at each stage?
  - Are QC and validation checkpoints built into its recommended workflows?
- Align your practice by keeping automation within defensible boundaries. Keep a copy of this sheet in your team's AI playbook. When evaluating a new product, mark where it fits and note which safeguards you'd need to move it leftward on the scale.

Considerations and Ethics for Legal Practitioners

02

tl;dr

What You Should Know

- AI aids, but doesn't replace, legal judgment. Lawyers are still beholden to fulfilling competence requirements and making client-centered decisions under ABA Formal Op. 512 and Model Rule 1.1.
- What "competence" means always evolves. Lawyers must stay current on new tech, including generative and agentic AI, as part of ongoing professional educational requirements.
- Generally speaking, courts care about defensible results, not specific tools—so you must verify accuracy and reliability when leveraging AI in discovery (see FRCP 26(g) and TAR precedents).
- One tool or workflow won't suit all matters. Legal teams can use varied methods according to their needs and preferences as long as they're reasonable, documented, and technically sound (see Sedona Principle 6).

What You Should Do

- Stay informed by learning which AI tools are relevant to your practice and evaluate them for responsible use. CLEs, conferences, and online resources like webinars can help.
- When using AI for discovery, keep humans in the loop and validate the tool's outputs using metrics like precision, recall, and elusion—as well as citation checking and thoughtful QC workflows—before relying on them.
- Use transparent, well-supported tools—and talk about them! You should aim to share AI usage guidelines across your firm.

The Responsible Practice of AI in Law

According to Formal Opinion 512 from the American Bar Association's Standing Committee on Ethics and Professional Responsibility, "while GAI [generative AI] tools may be able to significantly assist lawyers in serving clients, they cannot replace the judgment and experience necessary for lawyers to competently advise clients about their legal matters or to craft the legal documents or arguments required to carry out representations."

The opinion goes on to say (emphasis added):

> **Emerging technologies** may provide an output that is of distinctively higher quality than current GAI tools produce, or may enable lawyers to perform work markedly faster and more economically, eventually becoming ubiquitous in legal practice and establishing conventional expectations regarding lawyers' duty of competence. Over time, other new technologies have become integrated into conventional legal practice in this manner. For example, "a lawyer would have difficulty providing competent legal services in today's environment without knowing how to use email or create an electronic document".[1]
>
> Similar claims might be made about other tools such as computerized legal research or internet searches. As GAI tools continue to develop and become more widely available, it is conceivable that lawyers will eventually have to use them to competently complete certain tasks for clients. But even in the absence of an expectation for lawyers to use GAI tools as a matter of course, **lawyers should become aware of the GAI tools relevant to their work so that they can make an informed decision, as a matter of professional judgment, whether to avail themselves of these tools** or to conduct their work by other means. As previously noted regarding the possibility of outsourcing certain work, "there is no unique blueprint for the provision of competent legal services. Different lawyers may perform the same tasks through different means, all with the necessary 'legal knowledge, skill, thoroughness and preparation'".[2] Ultimately, any informed decision about whether to employ a GAI tool must consider the client's interests and objectives.

Arguably, agentic AI falls into that realm of "emerging technologies"—and, as Ahmadia noted, it's erupted onto the market with remarkable speed. Lawyers are not required to use generative or agentic AI to complete their legal work, but building awareness of the tools available seems like a good strategy when they may have an impact on lawyers' ability to deliver competent legal services to their clients—especially in circumstances such as e-discovery, which frequently involve very large amounts of data and very short deadlines.

There is already a broad spectrum of AI usage among legal teams, particularly in e-discovery.

Automated contextual analysis and topic clustering of documents has been available for almost two decades. Early machine learning tools underwent rigorous scrutiny over the last decade, ultimately securing judicial approval for technology-assisted review (TAR) workflows on e-discovery projects provided technical competence and proper validation exercises are established.

Many more teams already successfully apply these same standards to the use of the latest flavors of legal AI, including generative and agentic tools, within their legal data-related workflows.

This diversity of workflows and strategies makes sense when every case, client, and legal team is so unique. It's also well-established and, in fact, protected, by industry best practices.

1 ABA Formal Opinion 477R: Securing communication of protected client information

2 Formal Ethics Opinion 08-451

Take Sedona Principle 6, which states that "responding parties are best situated to evaluate the procedures, methodologies, and technologies appropriate for preserving and producing their own electronically stored information" and, in Comment 6.b, notes that "responding parties should be permitted to fulfill their preservation and discovery obligations without preemptive restraint."

The Federal Rules of Civil Procedure—particularly Rule 26(g)—state that attorneys must, after conducting a "reasonable inquiry," certify that any discovery request, response, or objection:

> **(A)** with respect to a disclosure, is complete and correct as of the time it is made; and
>
> **(B)** with respect to a discovery request, response, or objection, is:
>
> > **(i)** consistent with these rules and warranted by existing law or by a nonfrivolous argument for extending, modifying, or reversing existing law, or for establishing new law;
> >
> > **(ii)** not interposed for any improper purpose, such as to harass, cause unnecessary delay, or needlessly increase the cost of litigation; and
> >
> > **(iii)** neither unreasonable nor unduly burdensome or expensive, considering the needs of the case, prior discovery in the case, the amount in controversy, and the importance of the issues at stake in the action.

"I think I can clearly satisfy that duty by using whatever tools I want if I build in some explicit steps to validate my results. It is not a process that needs to be sufficient to meet the bar set by 26(g); it is the outcome," argues Manfred Gabriel, a partner at Holland & Knight.

"This is the one great thing that came out of our long, sad struggle over TAR. It's not the process, beyond that lawyers have to be accountable for it—as long as you have a process for validating results at the end," he continues. "So I don't actually care if you use agentic AI, or non-agentic AI, or machine learning, or humans, or dowsing rods, as long as there is a process at the end by which we can calculate recall and precision."

Validation is essential, but competency is at the heart of the responsible practice of law, with and without the help of shiny new technology. ABA Model Rule 1.1 is competence: "A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation."

Further, competence isn't a one-time exercise, like passing an exam or achieving a permanent certification. It's an ongoing duty to "maintain the requisite knowledge and skill," which means "a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject."

Like their US counterparts, Australian lawyers must stay current with technological developments relevant to practice, a duty reinforced by regulators and increasingly tested in disciplinary proceedings. The Legal Profession Uniform Law (applicable in New South Wales and Victoria) requires lawyers to deliver legal services "competently and diligently," while other jurisdictions maintain equivalent standards.

The Law Council of Australia has acknowledged that artificial intelligence presents both opportunities and professional responsibilities, noting that while AI tools may enhance efficiency, they cannot substitute for a lawyer's professional judgment, ethical obligations, or duty to clients. The Victorian Legal Services Board & Commissioner echoes this position, emphasizing that lawyers must understand AI's capabilities and limitations to provide accurate advice and maintain ethical obligations, while warning that improper use may carry "ethical, professional, and legal risks."

Justice Bell's Practice Note SC GEN 23 (NSW Supreme Court) reflects a cautious approach, requiring practitioners who use AI for submissions to verify sources and legislative authority. Bell stresses that "the task of judging is a human one" and warns against abdication of judicial reasoning to machines.The Australian Government's Policy for the Responsible Use of AI in Government (v2.0, effective December 2025) sets out mandatory principles for transparency, accountability, and risk-based governance.

The EU's Artificial Intelligence Act (AI Act), meanwhile, takes a firmly risk-based approach, banning certain "unacceptable" uses of AI and imposing extensive obligations on high-risk systems in areas such as critical infrastructure, employment, credit, public services, law enforcement, and judicial or administrative decision-making. Lawyers advising EU clients must help determine whether AI systems fall into these categories and ensure compliance with governance, documentation, and procedural requirements. In practice, this often means embedding AI risk assessments and compliance checklists into legal advice and maintaining a strong focus on fundamental rights, safety, fairness, bias mitigation, data protection, and explainability. The Act also explicitly promotes AI literacy, indicating that lawyers need sufficient technical understanding to assess risk, ask informed questions, and ensure ethical AI use.

The UK, for its part, has opted for a principles-based approach that relies on existing regulators and laws to apply core principles such as safety, fairness, transparency, accountability, and explainability to the use of AI. Legal guidance is clear: AI may support legal work, but can never replace humanjudgment or professional responsibility. Judicial guidance, updated in 2025, warns strongly against uncritical reliance on AI-based tools—noting risks like hallucinated case law, bias, deepfakes, and privacy breaches—and reinforces that responsibility for submissions always rests with the human lawyer. The Law Society and Bar Council echo this stance, stressing supervision, verification against primary sources, confidentiality, and explainability.

Across all of these jurisdictions, AI fluency has become a core professional competency: lawyers who understand AI's limits and apply rigorous oversight are best positioned to use it responsibly while upholding ethical and legal standards.

It should be noted that this isn't an undertaking you must go alone. Bringing in the right help to support your ongoing education, advise on which technology may be right for your team, and the best way to leverage fit-for-purpose agentic AI in your practice is an excellent strategy—whether it's fellow practitioners with more technology experience, consultants, service providers, or vendors who can help fill in any skills gaps within your team.

Build Responsible AI into Your Practice

Agentic applications of AI for legal are still nascent, and many new tools are sure to splash into the legal community in the coming months. As you—with guidance and input from colleagues, clients, peers, and/or experts—begin to explore them, sticking to these basic ethics will help you navigate a shifting landscape with grace (and emerge with some serious skills to define your future career path).

The key is to lean on *your* skills, and keep building them—not to rely too heavily on those shiny new tools.

"I like to give the tool direction, like a prompt, and ask it to present to me a plan or two for what it wants to do. And then we go back and forth talking about the plans—*we* meaning me and the machine, ha—and collaborate on how we want it to act," says Lenora Gray, associate director and data scientist at Elevate.

"After it does the work, I want to test the output myself. I don't like the idea of an LLM testing itself because I've seen some crazy behavior around that. One time, I had it write its own unit test and run the test itself to make sure the code it had produced was working—and then it deleted the unit test in order to pass it," she continues.

So yes: keeping a human in the loop is crucial for maintaining the quality, clarity, and control that matter so much in every legal data intelligence context.

But what does this *mean* in practice?

Whether you're using Claude Code to develop in-house review innovations like Gray or leveraging purpose-built tools like Relativtity aiR for Case Strategy, it means that AI—even agentic AI—is no "easy button." While agentic systems can significantly accelerate and even transform your work, there is no set-it-and-forget-it solution to legal tasks; they will always require the wisdom and strategic, contextual thinking at which humans excel.

So, when evaluating agentic tools for legal use cases, look for the following green flags:

- **The ability to test and iterate prompts** to ensure you're getting the results you want before running it on colossal data sets
- **Built-in validation workflows,** such as calculations for well-established benchmarking metrics like precision, recall, and elusion rate
- **Transparent reasoning,** perhaps delivered in the form of rationales and citations that you can read, verify, and interrogate further with the tool
- **Honest and principled developers;** you can only trust a tool as far as you can trust the team building it
- **Hands-on support teams** who will help you use the tool in the way it was designed, optimized for your needs and workflows

Finally, be sure to socialize these standards across your team—in every department. Prioritizing the responsible use of *any* type of AI is an essential tactic for success in every kind of twenty-first-century organization.

Recommended Resources

Opinions on this topic abound and conversations are plentiful. Here are a few references to get you started:

- Read this excellent article on the ethics of agentic AI, written by Tara Emory and Maura Grossman for the *National Law Review*.
- For a closer look at Formal Opinion 512 and its implications for e-discovery practitioners, read this blog post.
- If you'd like a deeper exploration on the ethics of using generative AI for e-discovery, here's an informative, on-demand webinar you'll love.

Tool: Green-Flag Checklist for Evaluating Legal AI Tools

When researching a new AI tool—especially an agentic one—aim for the high bar. No product is perfect, but some more clearly reflect ethical design, thoughtful engineering, and defensible use in legal practice than others. Use this checklist to spot genuinely fit-for-purpose innovation.

Category	What to Look For	Why It Matters
Transparency	Outputs include *rationales, citations, or reasoning* traces that users can review and interrogate.	Builds explainability and trust; supports defensible outcomes.
Ethical Development	Developers publish *ethical use principles*, disclose limitations, and define appropriate use cases.	Signals accountability and alignment with professional standards.
Validation & Accuracy	The tool measures *precision, recall, and elusion rate* or equivalent performance metrics.	Enables outcome validation—critical for FRCP 26(g) compliance.
User Control	Users can *test, customize, and iterate prompts* before running large jobs or client data.	Preserves human oversight and control over outputs.
Support & Training	Vendor provides *hands-on onboarding, responsive support, and iterative improvement loops.*	Drives adoption success and informed tool use.
Security & Compliance	Meets recognized standards: *SOC 2 Type II, ISO 27001, HIPAA, FedRAMP,* etc.	Protects sensitive client data and ensures compliance readiness.

How to Use This Tool

- Evaluate new tools before adoption or pilot testing using this list. Create a checklist featuring each of these green flags. You can also use a grading system, if a yes/no checkbox isn't enough. For example, give 2 points to tools with documented proof they meet these standards; 1 point to partially compliant tools; and 0 points for tools with no supporting evidence or unknown status.
- Reassess existing tools—especially free or freemium ones—to spot hidden risks. Do they hit most of these boxes, or are there any big misses? If the latter is true, act quickly to minimize your risk exposure. You can share your evaluations internally to illustrate what responsible AI readiness looks like in practice.
- Keep on that governance documentation train! Use this checklist as a first-step, repeatable exercise in your organization's standard procurement and AI ethics policies.

Green-Flag Checklist for Evaluating Legal AI Tools

	Category	What to Look For	Why It Matters
☐	Transparency	Outputs include *rationales, citations, or reasoning* traces that users can review and interrogate.	Builds explainability and trust; supports defensible outcomes.
☐	Ethical Development	Developers publish *ethical use principles*, disclose limitations, and define appropriate use cases.	Signals accountability and alignment with professional standards.
☐	Validation & Accuracy	The tool measures *precision, recall, and elusion rate* or equivalent performance metrics.	Enables outcome validation—critical for FRCP 26(g) compliance.
☐	User Control	Users can *test, customize, and iterate prompts* before running large jobs or client data.	Preserves human oversight and control over outputs.
☐	Support & Training	Vendor provides *hands-on onboarding, responsive support, and iterative improvement loops.*	Drives adoption success and informed tool use.
☐	Security & Compliance	Meets recognized standards: *SOC 2 Type II, ISO 27001, HIPAA, FedRAMP,* etc.	Protects sensitive client data and ensures compliance readiness.

Practical Use Cases

03

tl;dr

What You Should Know

- Agentic AI is already in use in the real world, from supply chain optimization to warehouse automation. Legal teams are joining the movement. Agentic AI now powers document review, privilege review, drafting, intake processes, contract analysis, and other data-heavy legal tasks.
- AI isn't (and can't be) a lawyer. Even highly capable systems can't meet ethical or character standards required for legal practice.
- Agentic AI can be built or customized in-house. Tech-savvy lawyers and teams are coding bespoke tools to fit their workflows and clients.

What You Should Do

- Experiment safely by starting with controlled pilots of trustworthy tools on data-heavy or repetitive tasks—and with strong human oversight. Work alongside service providers or software developers to ensure you're making the most of these experiments and learning from the experts.
- Keep learning, because the tech landscape is going to keep changing. Explore resources like LegaltechHub and trusted vendor blogs to stay current on AI tools and trends.

AI Agents in the Broader World

Now that we're grounded in agentic definitions and theory, let's see what AI agents look like in practice.

Supply Chain Directors

Bennett Borden, founder and CEO of Clarion AI Partners, staffs a team of legal technologists and data scientists to build bespoke AI tools for their clients. One of his favorites, he tells us, is an agentic system that directs and redirects their client's supply chain components, like shipping routes, in real time.

"This was for a giant retailer, and we were able to take their immense stores of data about all the supply chain and related issues they had on hand and enrich that data with other types of information: weather patterns, political developments, safety concerns," he explains. "Then, an ensemble of agents—a multi-agent system in which each agent does its job according to its own constitution—sets AI against itself to discuss potential solutions until they come to a consensus."

These agents **perceive** developments or risks along potential routes based on real-time data updates, **think** about the best ways to avoid or minimize those risks, and **act** by sending out instructions based on that analysis.

Warehouse Robots

Autonomous warehouse bots are another perfect example of agentic AI in action.

These bots use a suite of sensors to **perceive** their environment, map out the space of the warehouse, and examine the boxes of items. Tasked with organizing these boxes—according to their contents, order status, fragility, age, or other qualities—the bots must then **think** and analyze where each box belongs. Finally, they **act** by labeling, sorting, and moving the boxes as needed, as well as updating central inventory records.

Examples of Agentic AI for Legal Teams

Now, imagine those warehouse bots are analyzing, labeling, and sorting documents instead of boxes—but digitally, in data stores rather than between shelves on an industrial floor.

That's agentic AI for document review, which is just one legal data intelligence use case for this technology.

Relativity aiR for Review, built for this purpose, combines agentic AI workflows with human guidance. It can review hundreds of thousands of documents per hour while users remain in control, directing and confirming its output—with a goal of eliminating repetitive, time-consuming cognitive tasks to help review teams produce better results, faster.

Other agentic AI use cases include case strategy, privilege review, and issue coding, as well as other legal data intelligence categories like contract review, data breach responses, and internal investigations.

Legal professionals also use agentic AI to manage tasks like case intake and triage, helping them sort new clients and matters according to practice area and priority level, offer round-the-clock answering services, and ensure timely and accurate communications when and where they're needed.

Noteworthy, of course, is that none of these tasks involve agentic AI acting like an independent lawyer—and it shouldn't. The unauthorized practice of law by non-lawyers *does* apply to technology, and while AI may be capable of passing the bar exam, as one practitioner pointed out, it can't meet the requirements of character and fitness that only a human person can.

It's also worth noting that the scope of agentic AI's potential for legal teams doesn't end with out-of-the-box tools.

Like Lenora Gray's team at Elevate, tech-savvy legal practitioners also leverage agentic tools to code their own applications, which help them customize the software they've already invested in and methodize unique workflows into repeatable best practices.

And at Bennett Borden's firm, every lawyer is also a technologist—and they leverage AI agents to build bespoke AI solutions for their clients, as well as their own team.

Recommended Resources

- Visit LegaltechHub for scores of information about available legal technology, from startups to top players.
- Follow relevant industry and vendor blogs to stay up to date on the latest developments in tools that have potential for you and your team.

Tool: Spot the Agent Reflection Exercise

Agentic AI is everywhere—from retail logistics to document review—but recognizing its *pattern* (perceive → think → act) helps you imagine new applications for your own practice. Use this quick exercise to turn everyday observation into professional insight.

Example	How It Perceives	How It Thinks	How It Acts	What Could This Look Like in Legal
Warehouse bots sorting packages	Uses cameras and sensors to "see" boxes	Categorizes based on content, order status, etc.	Moves boxes and updates inventory	AI system categorizes documents, flags relevance, updates doc status
Supply chain agent rerouting shipments	Monitors weather and route data	Analyzes alternate paths and risks	Redirects trucks to safer routes	AI triages discovery priorities or reallocates doc review resources
AI triages discovery priorities or reallocates doc review resources	Reads incoming queries	Determines intent and next best step	Responds or escalates	Intake agent classifies client requests and routes them internally

How to Use This Worksheet

1. Identify three agentic systems you've encountered in the world.
2. Break down how each one perceives, thinks, and acts.
3. Translate what you've learned into ideas for your own workflows.

Spot the Agent Reflection Exercise

Example	How It Perceives	How It Thinks	How It Acts	What Could This Look Like in Legal
Look for automated systems that you encounter every day, while you're driving, shopping, browsing for shows, etc.	How does this system take in data from its environment, digital or physical?	What does the system do with that data, once collected? How and why does it analyze it?	Once that analysis is complete, what does the system do? Is it autonomous?	Can you think of a parallel legal exercise that calls this thought process to mind?

How to Get Started

04

tl;dr

What You Should Know

- We're in an "opt-in" moment. Using AI isn't yet mandatory, but developing fluency now will pay off as it becomes standard practice. The biggest barrier to using agentic AI is psychological: overcoming hesitation and building comfort with experimentation. Those who start now will come out ahead.
- Agentic AI is a great opportunity to learn by playing. Safe, hands-on experimentation builds intuition, creativity, and understanding faster than theory ever could.
- Data and design matter. Understanding where your data lives and how it flows is key to making AI useful—and safe—in your workflows.

What You Should Do

- Begin by testing AI on low-stakes, repetitive tasks—personal or professional—and iterate. And play *safely*. Use approved tools, never disclose sensitive data, and practice in "sandboxes" to explore agentic frameworks.
- Map your firm's ecosystem by identifying automatable workflows, data touchpoints, and where human judgment adds the most value. Build buy-in by sharing your enthusiasm as you learn more and develop that agentic mindset.
- Make fostering an AI-enabled culture within your firm, with appropriate guardrails, a top priority. Encourage team discussions about AI policies, risk tolerance, and pilot opportunities.
- When using AI, proactively measure and refine. Track validation metrics (e.g., accuracy, bias, impact) and bake those touchpoints into your process from the start.

Develop an Agentic Mindset

For many, the biggest hurdle to getting started with agentic AI—or any new technology, for that matter—is simply taking the first step. Where do you begin? What tool do you try? Which problem do you aim to solve first?

Consider how most of us consume new shows (and movies, and docuseries…) in the 2020s: no TV guides, no scheduling life around a premiere or finale, no waiting for DVDs (or VHS tapes) months after a season concludes. We can choose to watch almost whatever we want, whenever we want to, from a vast digital library of fresh and old-favorite options.

Those who want to watch new programs as they air on live TV, and watch them only then, certainly can. But it isn't the norm anymore, is it?

Odds are the version of you who watched *Grey's Anatomy* in 2005 would be pretty confounded by the way you watch it today. When the show's complete seasons started to hit Netflix in 2009, and then new episodes became available the day after airing on Hulu in 2024, it felt different at first—but you adapted to each change sooner or later. And before you knew it, that was just how TV was done.

Similarly, AI offers an entirely new way of tackling so many tasks—at work, at school, at home. At some point, it will simply be the way things are done in many areas of life. But right now, today, picking up an agent and getting going with it feels like a big step.

That's because this is an opt-in moment, particularly for generative and agentic AI.

Here, you face choices that feel a bit confusing, begin building new skills, and change well-established habits. It's a lot of transformation, and while it'll someday feel normal, you probably aren't there yet.

In legal, especially, the act of opting in is big. Learning and using the tools themselves isn't the colossal part; building the confidence and buy-in to defend and advocate for them is.

You'll find that many AI experts, including AI Visionaries, share similar advice here: don't overthink it.

Start small. Begin by asking yourself a small question whenever you find yourself doing a repetitive or research-intensive task: can AI help me automate this?

Consider starting with tasks outside of work (or at least outside of your performance-reviewable tasks) to soothe your hesitation, and make exploring agentic AI feel less risky and more fun.

If you periodically go to Google to check for news about a topic of interest, pause—AI can do it for you. Go to ChatGPT and explain that you'd like to be notified whenever a favorite author announces a new book on the horizon. Ask it to set up a watcher, which will periodically scan available sources for such an announcement and notify you if it finds one. You'll only need to think about it when there's actual news to be read, making for one less distraction the rest of the time.

Another example that falls into a more professional realm, if your firm is a Microsoft 365 shop, would be to create a Monday Brief tool using Microsoft Copilot's agent builder. In Copilot Studio, describe an agent that will review your calendar at the start of every week. Instruct it to send you an email at 6:00 a.m. every Monday, with a summary of your agenda for the week: meetings (and how heavy your meeting load will be), who you'll be talking to, what you're required to prepare. See how you like it when the first email come through; you can jump back into Copilot Studio to iterate on your

instructions as needed.

These are pretty basic agentic workflows, but that's where you should begin: the basics. Again, don't overthink it.

The more you engage in exercises like this, the faster your mindset will evolve. Eventually, reaching for a little productivity boost from AI will come as naturally as typing in a Google search.

A Word on Safe Sandboxes

Ray Mangum puts it simply: "Playing with it is the best way to learn. Go out and do it and don't be intimidated—just start using it. Come up with a real problem and try to solve that problem. Even if you don't make it all the way through, you're going to learn a tremendous amount from doing this."

The best way to build intuition around agentic AI is to play—but to play safely.

Here are a few ethical, low-risk ways to do that:

- **Try general-purpose tools** like **ChatGPT, Perplexity**, or **Copilot** for low-risk, individual tasks to practice giving structured, multi-step instructions. Try out the examples described in the previous section or experiment with your own ideas.
- **Learn about agentic frameworks** like **Claude Code, GitHub Copilot,** or open-source sandboxes such as **CrewAI** or **OpenDevin** if you're technically inclined.
- **Use non-sensitive data only.** Do not create problems for your firm by downloading or accessing unauthorized tools at work (a practice referred to as "shadow AI"). Keep things light and productivity- or creativity-focused, within approved applications and without disclosing any information you and your firm are entrusted to keep safe.
- **Compare reasoning styles.** Ask the same question across tools and observe how each plans and explains its response. You'll get a feel for the styles and features of differing tools, which will lend some fascinating insight into the spectrum of AI innovation and development.

Think of this as building muscle memory for working with generative and agentic systems. The goal isn't to master a specific tool—it's to understand how these systems perceive tasks, plan their approach, and act toward completion.

It's also to build new habits for yourself, training your brain to innovate and solve problems in new ways.

Leveraging Agentic AI to Evolve Your Practice

As you build this knowledge, you can also begin advocating for the safe use of fit-for-purpose AI within your team. Start fostering that AI-enabled culture across your firm by asking the right questions:

- Do we have any internal AI usage policies or other guidance frameworks? If not, can we start building some of these materials?
- Are your colleagues using AI—sanctioned or otherwise? How can you help deter unsanctioned use and direct that curiosity in better directions?
- What workflow problems do we consistently experience as a team? As a firm? How might agentic AI help automate some of those rote tasks and help us start tackling those problems?

Once these conversations begin, even if you're not yet cleared to use agentic tools in production, conduct some internal mapping:

- Identify which tasks are repeatable, rule-based, or ripe for automation.
- Document where your data lives and how it moves between systems.
- Think about where human judgment adds the most value, and where automation could relieve cognitive load.
- Assess your firm's risk tolerance, preferred success measurements, and creative appetite. This will help you distinguish the earliest-available opportunities for experimentation from pipedreams.

These insights will prepare your team to run thoughtful, compliant pilots

once governance frameworks are in place.

And then? Manfred Gabriel reminds us that technical experimentation is only half the battle.

“Data you have in a tool like RelativityOne already gives you a very mature repository function. But all the other data that is floating around your enterprise or in people’s heads, is not in a usable form in that state,” he says. “Connecting systems is a pain. But thinking about not just the sexy agents, but where your data is and how it flows, is a key to making this work.”

That’s a powerful cue for legal professionals. Understanding your data ecosystem—the repositories, processes, and bottlenecks that shape your daily work—is foundational to adopting agentic AI responsibly.

So is referring back to that risk tolerance and keeping success and validation in mind as you innovate.

“How do you identify the risks of your use of AI, reasonably mitigate those risks, and prove that mitigation? Build those metrics into your system—it’s a lot easier to use them if they’re inherent to the system than go back and add or calculate them later,” advises Bennett Borden.

“We are dealers in information and insight. Discovery is all about finding out what happened, why, and who did it. Lean into learning about how to build these agents on your own. Play with it—give it a constitution: ‘Here’s your job and the rules I want you to work by,’” he continues. “Set it loose and give it feedback on its output. You aren’t programming; you’re forming a psychology for one of the most powerful minds ever to exist, and you need to think about how you put controls around it to ensure the behavior is exactly what you want it to be.”

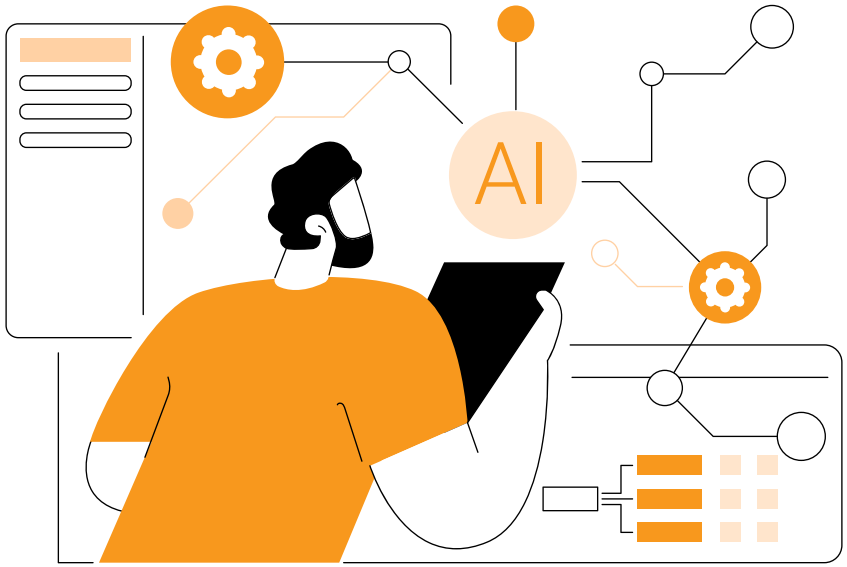

Recommended Resources

- Sign up for practical, hands-on AI newsletters, like TAAFT and The Rundown, for ideas on using AI within and beyond your day-to-day work and start thinking in new ways.
- Explore innovations backed by The LegalTech Fund to see what’s out there in legal tech and learn more about how your peers are innovating with agentic AI.
- Facing a resistant colleague or leader? Learn about Maurer’s Three Levels of Resistance and Change to better understand their mindset and build a well-suited approach to helping them get comfortable with AI. You can also learn some tips about advocating for AI internally—before your firm falls behind—in this on-demand webinar.

Tool: A Prompt for Using Agentic AI in Your Professional Development Planning

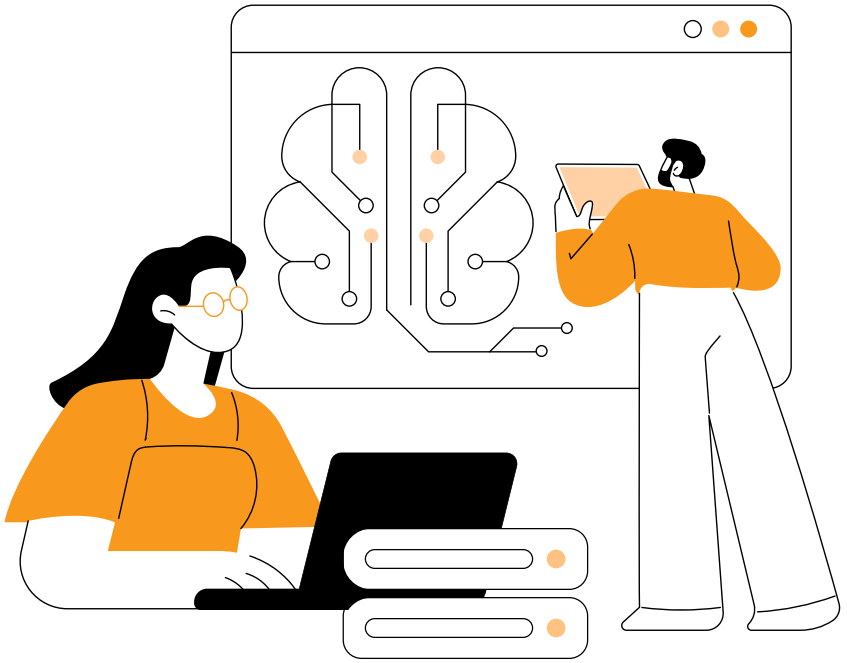

Many of the popular chatbots currently on the market offer agentic functionality to some degree. Using them in this capacity gives you an opportunity to practice your generative AI, prompting, and agentic AI skills simultaneously.

We recommend beginning with low-risk, but potentially high-impact, personal projects. This will help you hone those agentic instincts, spot more opportunities to use and learn about AI, and reduce the cognitive effort required to complete relatively simple but research-intensive tasks.

The bonus? Habitually collaborating with agentic AI on everyday tasks will lay the groundwork for new kinds of creativity and innovation at work, too.

For example, try using these widely available tools to:

- Research, propose options for, and plan an upcoming vacation
- Send a weekly briefing to your inbox highlighting recent news in your areas of interest and why it matters to you specifically
- Create a monthly meal plan that will generate and send weekly shopping lists to you, crafted according to your tastes, budget, available time, and cooking skills

To get you started—and get your professional wheels turning without risking any sensitive, client, or matter-specific data—here's a prompt you can try right now.

Begin by opening your favorite commercially available chatbot (ChatGPT is a familiar option, and this prompt was also tested on Perplexity and Google Gemini with decent results). Use the following information to build your prompt.

Step 1: Provide background on your career

The simplest way to do this is by uploading your resume as an attachment. If you prefer, you can also write a brief description of your role and career ambitions. It might look something like this:

> **I am a legal data intelligence practitioner focused on litigation and internal investigations. Currently, I live in Chicago, Illinois and work as a project manager on the legal team at a large company in the pharmaceutical industry. My work is mostly focused on legal technology support and implementation, e-discovery project management, and connecting our department to IT, outside counsel, and other internal stakeholders as needed for each project. I have a bachelor's degree in history and have been working in similar roles for about seven years.**

You can add any nuance or detail that feels important, but succinct and direct descriptions should work just fine as a starting point in the absence of a formal resume.

> ***Note:*** *If you don't have an updated resume on hand but you're an active LinkedIn user, you can visit your LinkedIn page, click Resources just below your name and current company, and select Save to PDF to create one based on your profile.*

Step 2: Describe your professional interests and career goals

Next, provide some insight into what you love about your job, what skills you'd like to hone further, and the path you see for your future career. Here's a continuation of our hypothetical example:

> **My favorite part of my job is the focus on technological innovation—I always enjoy experimenting with new tools to see how they can increase our efficiency and prove our team's value to the company. I also like my role as a connector between teams; helping to bridge the gap between our tech experts and legal practitioners is a dynamic, human, and interesting challenge that keeps me on my toes. I am a Relativity Certified Administrator and need to earn 12 RCE credits every year to maintain that certification. I'd love to challenge myself and become a Relativity Expert in the next couple of years. Long term, I'd like to find a job like "Head of Legal Data Intelligence" or something similar. I imagine I'll probably stay in-house, but can also see myself at a law firm if my skills transfer and the opportunity is right. For now, I am happy in my role but hope to climb the ladder internally to the best of my abilities.**

If you are an attorney and earning CLE credits is an important part of your professional development goals, be sure to include that information. Same goes for any ongoing certifications you hold and need to maintain.

> ***Note:*** *If you have a paid subscription to ChatGPT, add the command "***/agent***" to the very beginning of this prompt or select* **Agent mode** *from the dropdown by clicking the + on the left side of your input box. This will instruct ChatGPT to deliver an execution-focused, deeper-research output rather than one presenting light research in a conversational style.*

Step 3: Instruct the agent on the type of output you want

As the final component of your initial prompt, tell the model what you need from it. "Help me craft a professional development plan" is too vague. Be specific about things like the timeline you're planning for, what sort of resources you're interested in pursuing, and your budget. If you'd like, you can copy and paste the following directly (with a few tweaks in the bracketed portions, as needed):

> **Based on what you see there, please help me plan my professional development strategy for the next 12 months. I would be interested in attending [2-4 conferences (preferably no more than one per quarter)] throughout this period. I'd also like to see recommendations on online publications to read or follow, highly rated books that may be helpful to me, and free or low-cost online courses. And don't forget about my [certification/bar] requirements, either. Please break your plan down in detail (with links, discrete goals I should aim for, etc.) for each quarter and aim to balance [LDI, collaboration, and technology/innovation-focused topics].**

Don't be tempted to expand on this portion for many paragraphs. Leaving some room for interpretation, grounded in your basic needs and timeline, is a great way to allow the AI to do its best work: conducting the research, gathering the results, and presenting its findings in a logical, digestible format that serves you well and goes a few steps further than the structure you may have come up with on your own.

Note: *Feeling adventurous? Start this chat with your chosen LLM (or a new one) multiple times and see how changes in your linguistic style influence its outputs. Experiment with more aggressive or curt language (e.g., avoid "please," "thanks," and any extra fluff) to see if the model responds more concisely. Try starting one prompt with a role-defining statement—e.g., "You are a business coach hired to help mid-career legal practitioners launch into their professional ambitions; today, you're sitting down with me to help me plan my professional development strategy for optimal career advancement by 2030"—and compare the results to one without it. These experiments can be a fun way to see what type of tone yields the results you like best.*

Step 4: Engage in some conversation to adjust the plan as needed

Once you enter that initial prompt, the model should present you with a fairly meaty strategy based on your current status, future ambitions, and top interests. Give it a good, close read. See what stands out, what lands, what sounds a bit off.

And next? This is the part where AI really shines: in collaborating with you, rather than delivering an output and disappearing into the void. Ask it why it recommended one conference and not another. Tell it you've heard of a new certification and are wondering if being an early adopter is a smart idea or a risk of wasting your limited time. Engaging in this back-and-forth, and allowing it to update your strategy accordingly, will help you truly tailor the strategy to your needs and interests so that it's as useful as possible.

Note: *Prompting is an art as well as a science. It's also a practice that's evolving quickly as AI improves and engineers craft more effective and intuitive user experiences. If you find this exercise interesting and want to learn more about the nuances of AI prompts and how to finesse them, check out this e-book.*

Relativity