



2026 Business Litigation Practice Group Seminar

June 3-5, 2026

AI in the Legal World

A Practical Guide for In-House and Outside Counsel

Jesse Conan

BECKER, GLYNN, MUFFLY, CHASSIN & HOSINKSI LLP

New York, New York

jconan@beckerglynn.com

D. Jay Davis, Jr.

CLEMENT RIVERS, LLP

Charleston, South Carolina

jdavis@ycrlaw.com

Part I: AI in the Legal Department

A Practical Guide for In-House and Outside Counsel

I. Laying the Groundwork: How to Organize for AI Success

Generative AI tools are only as useful as the underlying materials, instructions, and controls surrounding them. Before an organization expands AI use, it should consider the following:

- **Organize source materials**
 - Store contracts, leases, correspondence, and litigation files in a single, searchable location (SharePoint, iManage, NetDocuments, etc.).
 - Move documents systematically into approved repositories or AI-enabled workspaces to support higher-value analysis.
 - Use consistent naming conventions and folder structures across matter types.
- **Select an enterprise-grade AI tool**
 - Avoid consumer-grade tools such as free ChatGPT, consumer-grade Claude, or non-enterprise Copilot. These tools are not appropriate for confidential legal work — queries may be used to train public models.
 - Enterprise solutions, such as Harvey, CoCounsel, Microsoft Copilot for Microsoft 365, and Claude for Enterprise, can provide contractual and technical controls that are more appropriate for legal work, but the specific terms should be reviewed before use.
 - Always verify the platform's data handling terms, privacy policy, and whether your data is used for model training.
- **Establish an AI Use Policy**
 - Establish who within your organization can use AI tools, for what purposes, and with what categories of information.
 - Require attorney review of all AI-generated work product before use or filing.
 - Establish data handling and usage requirements.
 - Address privilege and confidentiality obligations in the policy.
- **Train your team**
 - Technology competence is an ethical obligation under ABA Model Rule 1.1, Comment 8.
 - Train attorneys and staff on effective prompting, verification requirements, and what AI cannot do.

II. High-Value Legal Use Cases

A. Litigation Support: Mining the Document Universe

Business litigation typically involves a large number of documents — leases and amendments, years of correspondence, internal communications — that must be rapidly analyzed for key provisions, factual timelines, and exposure assessments.

EXAMPLE: Tenant files a co-tenancy dispute claiming years of overpaid rent. AI can help counsel: (1) locate the operative co-tenancy language and amendments; (2) build a chronology from notices, correspondence, and payment records; (3) identify potential notice defects; and (4) compare the provision against other leases at the property. Counsel must still verify each cited source before relying on the output.

Potential time savings: AI can turn a full-day first pass into a focused attorney review, but the lawyer remains responsible for verifying the source documents, legal analysis, and final work product.

B. Deposition Transcript Analysis

After depositions, AI can be used to rapidly identify and cite key testimony across hundreds of pages of transcripts, dramatically reducing the time required to prepare deposition summaries, identify inconsistencies, and support summary judgment briefing.

EXAMPLE: Upload transcripts to an approved AI tool and ask it to: (1) identify testimony about a specific event or fact with page-line citations; (2) flag potential inconsistencies; and (3) summarize key admissions by witness. Treat the output as a draft work-product aid, not a filing-ready summary.

C. Lease and Contract Summaries

For organizations managing large portfolios of similar contracts, AI can generate standardized summaries identifying critical terms, non-standard provisions, and deadline obligations — enabling faster intake, easier comparison, and better institutional knowledge.

STEPS: (1) Organize leases/contracts in a structured folder. (2) Build a workflow or agent with a standard extraction template (rent terms, notice requirements, termination rights, renewal options, etc.). (3) Run the query across all documents. (4) Attorney reviews and approves output.

Implementation Guidance:

- Start with lower-risk, repetitive tasks, such as summarization, first-pass document review, and basic contract analysis.
- Build workflows or agents that standardize recurring outputs, such as lease summaries, contract summaries, or issue-spotting checklists tailored to the organization's needs.
- As the organization's AI use matures, move to more advanced workflows that require more design, testing, and attorney supervision. Examples: drafting responses and objections to document requests; preparing tabular reviews of contracts, pleadings, transcripts, or production sets.
 - Examples:
 - Drafting responses and objections to document requests.

- Preparing tabular reviews of contracts, pleadings, transcripts, or production sets.

III. Common Concerns and How to Address Them

1. Data Protection and Confidentiality

A principal risk in using AI for legal work is exposing confidential client, company, or privileged information through tools that lack appropriate contractual, technical, or ethical safeguards.

- **KEY** — Use approved enterprise-level tools with contractual commitments addressing confidentiality, data use, retention, access controls, and training on customer data.
- Review the vendor's privacy policy, Terms of Use, data retention terms, human-review practices, and model-training rights before using any tool with confidential information (ABA Formal Op. 512, 2024).
- Obtain client consent before inputting client information into AI tools that may not adequately protect confidentiality.

2. Hallucinations: AI Cannot Be Trusted for Legal Research

General rule: be highly skeptical of any AI-generated legal research output unless independently verified against reliable primary or secondary sources. Generative AI tools hallucinate. They generate plausible but fictitious citations, case holdings, and quotations. Studies have found leading AI legal research tools hallucinate 17–33% of the time.

WARNING: In *Mata v. Avianca* (S.D.N.Y. 2023), attorneys submitted a brief containing AI-generated fake case citations to a federal court. The attorneys were sanctioned \$5,000. The court held that there is nothing inherently wrong with using AI, but attorneys bear the gatekeeping responsibility to verify accuracy before filing.

- Always verify AI-generated legal research in Westlaw or Lexis before use.
- Do not cite AI-generated authority in advice, filings, or presentations without independent source verification.
- AI is most reliable when working with documents you have provided — your own leases, contracts, correspondence, or caselaw — rather than when generating legal conclusions from general knowledge.

3. Preserving Privilege and Work Product Protection

Privilege and work product protection depend on the purpose of the AI use, the confidentiality settings of the platform, the terms governing access to prompts and outputs, and whether counsel directs and supervises the work.

For Outside Counsel:

- Use approved enterprise-level tools with confidentiality, access-control, and data-use restrictions appropriate for client information.
- Obtain client informed consent before inputting client information into any AI tool (ABA Formal Op. 512).

- Boilerplate engagement letter language authorizing AI use is insufficient — consent must be informed and specific.

For In-House Counsel:

- Queries to AI platforms are not privileged on their own — AI is not an attorney.
- Queries should be made at the direction of counsel and documented as part of the legal work product process to maximize protection.
- Do not put privileged attorney communications or legal strategy into consumer AI platforms.

IV. Key Cases Every Legal AI User Should Know

Mata v. Avianca, Inc. / *678 F.Supp.3d 443 (S.D.N.Y. 2023)*

Attorneys sanctioned \$5,000 for submitting AI-generated fictitious case citations without verification. Establishes that using AI is permissible, but attorneys bear personal responsibility to verify the accuracy of all AI-generated content before filing.

United States v. Heppner / *No. 1:25-cr-00503 (S.D.N.Y. Feb. 2026)*

Defendant's Claude prompts and outputs held not protected by attorney-client privilege or work product doctrine. AI is not an attorney; communications with AI are not confidential because Anthropic's privacy policy permits data disclosure; and queries not made at counsel's direction are not work product.

Warner v. Gilbarco, Inc. / *No. 25-cv-10283 (E.D. Mich. 2026)*

Court held that attorney's use of AI tools in preparing work product was not discoverable. AI use in litigation preparation, when directed by counsel, may be protected work product — but the attorney must remain in control of the process.

In re Kellogg Brown & Root, Inc. / *756 F.3d 754 (D.C. Cir. 2014)*

In-house counsel's status does not dilute attorney-client privilege. Internal investigations conducted under the direction of an in-house legal department may be protected. The case is a useful analogy for documenting AI-assisted internal investigations under the direction of counsel, although it does not address AI directly.

Upjohn Co. v. United States / *449 U.S. 383 (1981)*

Foundational case confirming attorney-client privilege applies to corporations. The privilege protects communications made to facilitate legal advice — the organizational framework for understanding when in-house AI use fits within the privilege.

V. Ethical Rules Governing AI Use

Rule	AI Implication
Rule 1.1 — Competence	Attorneys must maintain competence in technology, including AI tools relevant to their practice.
Rule 1.4 — Communication	Consider whether and how to disclose AI use to clients, especially when AI processes client information.
Rule 1.5 — Fees	Ensure fees remain reasonable and accurately reflect how AI affected the time, staffing, and value of the work.
Rule 1.6 — Confidentiality	Use only AI tools that adequately protect client confidential information from disclosure.
Rule 3.1 — Meritorious Claims	Verify all AI-generated legal research and citations before relying on them in any filing or argument.
Rule 3.3 — Candor to Tribunal	AI-generated fake citations submitted to a court may constitute a sanctionable violation.
Rule 5.1 / 5.3 — Supervision	Supervising attorneys are responsible for AI outputs generated by subordinates and staff.

VI. The Bottom Line

AI is not a replacement for attorney judgment — it is a force multiplier for attorney time. The organizations that will benefit most are those that:

- Invest in document organization before deploying AI tools.
- Select enterprise-grade platforms with appropriate data protection.
- Establish clear policies governing how, when, and by whom AI is used.
- Train their attorneys to verify AI output and maintain professional responsibility.
- Use AI for high-volume, document-intensive tasks where the time savings are greatest.

REMEMBER: AI can accelerate reading, sorting, and drafting, but lawyers remain responsible for judgment, strategy, and verification. Use AI to create more time for the legal analysis, counseling, and advocacy that require a lawyer.

Source Materials

The following materials were referenced in preparing this presentation:

- ABA Formal Opinion 512 (July 29, 2024) — Generative Artificial Intelligence Tools.
- Mata v. Avianca, Inc., 678 F.Supp.3d 443 (S.D.N.Y. 2023).
- United States v. Heppner, No. 1:25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026).
- Warner v. Gilbarco, Inc., No. 2:25-cv-10283 (E.D. Mich. 2026).
- In re Kellogg Brown & Root, Inc., 756 F.3d 754 (D.C. Cir. 2014).
- Upjohn Co. v. United States, 449 U.S. 383 (1981).
- ISO/IEC 42001:2023 — Artificial Intelligence Management Systems.
- ISO/IEC 27001:2022 — Information Security Management Systems.

Part II: EFFECTIVE LEGAL TOOLS TO CONSIDER AND USE CASE IDEAS

MICROSOFT COPILOT: AN ASSISTANT THAT NEVER SLEEPS

Introduction

Microsoft Copilot represents a meaningful shift in how AI can be used in the practice of law. Unlike standalone AI tools, Copilot is built directly into the Microsoft 365 ecosystem that most firms already rely on every day: Word, Outlook, Excel, PowerPoint, SharePoint, and Teams. Because it sits within that environment, it operates on a foundation of existing permissions, matter-level access restrictions, and enterprise-grade security.

This portion explains what Copilot is, how it works, and how lawyers can use it to summarize information, draft documents, manage inboxes, prepare for meetings, and automate routine workflows. It also describes advanced capabilities — Prompt Coach, scheduled prompts, Copilot Notebooks, and Copilot Agents — that go beyond basic chat and help attorneys systematize work at scale.

Note— though this presentation is focused on Copilot, the content and approach will work with most enterprise-level AI platforms, as highlighted below. Our notes in this regard will focus on two such platforms, Harvey and Claude for legal.

Harvey is a law-specific AI-platform built for the legal and professional service industry. Harvey is not itself an LLM. Harvey, rather, is a platform designed to facilitate the use of AI to run and complete legal tasks. When a prompt is submitted to Harvey, it will run that prompt on one of several LLMs depending on the nature of the task. (*i.e.*,

OpenAI, Claude, Gemini). Claude for legal is predicated on Anthropic's AI-platform and comes with a suite of law-focused AI workflows and tools.

What Copilot Is — and Why It Matters

Microsoft Copilot is an AI assistant embedded directly in core Microsoft 365 applications. Within these apps users can call Copilot to summarize, draft, rewrite, analyze, or organize information. For lawyers, this matters because these tools are where nearly all legal work already takes place. Copilot does not replace legal judgment. Instead, it accelerates labor-intensive parts of practice that are not uniquely legal: sifting through your enterprise workspace (including calendar and emails), drafting standard communications, creating outlines, summarizing case law, and organizing documents. It can synthesize information from a firm's existing content — case files, memos, correspondence, contracts, and pleadings — subject to the same permissions that govern access to those materials.

Importantly, Copilot respects Microsoft 365 permissions, Purview controls, and auditing mechanisms. That means it cannot access matters or folders the user cannot access. It also means that prompts and outputs remain within the organization's security framework — an essential consideration for privacy, confidentiality, and ethical obligations.

Make sure to closely review the terms and conditions of your firm's contract with Microsoft. Not all Enterprise 365 accounts are equal in terms of the security settings and permissions. At a minimum, be certain before activating Co-Pilot that Co-pilot will not train on your firm's documents.

Writing Better Prompts: Copilot's Built-In Prompt Coach

Copilot's usefulness depends heavily on the clarity of the instructions it receives. To help write effective prompts, Microsoft created Prompt Coach, a built-in agent and template available in Copilot Studio. Prompt Coach analyzes an attorney's instruction and provides feedback to improve structure, clarity, specificity, and compliance.

For example, many lawyers inartfully start with vague prompts like: "Summarize this case." Prompt Coach will help write more purposeful instructions such as: "Summarize the arguments asserted by the movant and draft a neutral explanation of the grounds for the motion, relevant facts, and requested relief for inclusion in an internal case memo." In addition, Prompt Coach checks for Responsible AI considerations — flagging ambiguous instructions, unsupported assumptions, or overly broad requests that may inadvertently pull in irrelevant data.

Note— Other platforms have also adopted prompt assistants. For instance, within Harvey, the chat window has an "improve prompt" function that functions in a similar way.

Automating Recurring Work with Scheduled Prompts

Many legal tasks follow predictable patterns — case updates, email digests, deadline reviews, or industry news monitoring. Copilot allows users to schedule prompts to run automatically at set intervals. Once configured, results appear directly in Copilot chat without requiring manual execution. Some examples you can try include:

- **Weekly Case Law Digest:** Every Monday, Copilot summarizes the top five appellate decisions across selected practice areas, complete with citations and short practical implications.
- **Morning Matter Status Reports:** Copilot compiles upcoming deadlines from Outlook and Planner, highlights recent client communications, and flags items requiring partner review.
- **AI and Industry News Digest:** Useful for firm leadership or marketing, Copilot summarizes developments in AI law, ethics guidance, or even online mentions of the firm.

Copilot Notebooks: Controlling Context and Scope

Copilot Notebooks are one of the most helpful features in the product. A Notebook is a controlled workspace in which the lawyer explicitly defines the universe of documents and references that Copilot is permitted to consider when answering a question. Within that bounded scope, Copilot's answers are grounded, traceable, and reusable.

NOTE – Harvey's file-storage system allows users to set up different Vaults within Harvey and upload materials into those Vaults.

Why Scope Matters in Legal Work

A general-purpose AI chat is a wide-open conversation. It will pull from any document it has access to, weight relevance by its own logic, and reach for adjacent files when it thinks they might help. That behavior is fine for casual research. It is not fine for a privileged matter, a deposition outline, or a damages memo where the wrong file pulled into the wrong answer can become a problem. A Notebook reverses the default. Instead of asking Copilot "what do you know," you are telling Copilot "here is where to look for what to know." That shift produces four practical benefits for legal practice:

- **Grounded answers.** Responses are tied to the documents you placed in the Notebook, not to whatever else lives in the firm's SharePoint or your OneDrive.
- **Reduced cross-matter risk.** Because the scope is explicit, there is far less chance that materials from an unrelated client matter will bleed into the analysis.
- **Consistency across the team.** A partner, an associate, and a paralegal asking similar questions of the same Notebook get answers that draw from the same record.
- **Defensible review.** You know, at any time, exactly which documents were available to the AI when it produced a given draft, summary, or chronology.

What Can Live Inside a Notebook

A Notebook can contain a wide range of reference material. The supported items typically include:

- Word documents (pleadings, memos, briefs, correspondence).
- Excel spreadsheets (damages models, exhibit indexes, billing detail).
- PowerPoint slides (mediation presentations, internal CLEs).
- OneNote pages (case notes, witness outlines, strategy memos).
- Meeting notes and transcripts from Teams.
- Web links (statutes, rules, public filings, court orders).
- PDFs that have been uploaded into the workspace.

The practical point is that the Notebook becomes a curated case file inside Copilot, designed to be queried. It is not a substitute for the matter file itself — it is a working layer on top of it.

NOTE – Vault's serve the same function within Harvey.

Six Notebook Worth Building

Notebooks are most useful when their scope is intentional. Below are six patterns I have found genuinely useful in my own litigation practice. Each can be created in minutes once the discipline of curating documents is established.

1. The Matter Notebook

One Notebook per active matter, containing the operative pleadings, the scheduling order, the most recent discovery responses, and any expert reports. This is the workhorse — the Notebook you open to ask routine questions about the file (“what is plaintiff’s current theory on causation,” “what deadlines fall in the next thirty days,” “what economic damages are being claimed”).

2. The Deposition Preparation Notebook

A more focused workspace built for a single deposition. It contains the witness’s prior statements, relevant medical records or business documents, the operative complaint, and any prior testimony from related witnesses. From this Notebook, Copilot can draft outline topics, identify likely impeachment points, and flag inconsistencies between the witness’s statements and the documentary record. Because the Notebook is scoped, the answers are tied to actual evidence rather than the model’s general intuition.

3. The Damages Notebook

Designed for a single case, containing medical records, billing summaries, wage records, life-care plan documents, and any economist or vocational reports. From this Notebook, Copilot can help draft a damages outline, surface inconsistencies among providers, and prepare summary tables for mediation. The lawyer remains responsible for the numbers; Copilot is responsible for not losing track of where they came from.

4. The Expert Vetting Notebook

A Notebook built around a single opposing expert. Contents include the expert’s CV, prior reports produced in your case, transcripts of prior testimony, published articles, and any prior Daubert rulings. Copilot can compare the expert’s prior opinions to the current report, flag changes in methodology, and surface candidate cross-examination topics.

5. The Transactional Precedent Notebook

Although my own practice is litigation, the same pattern translates directly to transactional work. A Notebook is populated with prior deal documents of the same type — say, asset purchase agreements in a particular industry. The drafting lawyer can then ask Copilot to identify how a particular clause has been handled across precedent, draft a comparable provision, or surface unusual deviations.

6. The Appellate Brief Notebook

For an appeal, a Notebook can house the operative record on appeal, the trial court orders, the briefs filed below, and any controlling authority. Copilot can help build an issue outline, locate record citations, and ensure that arguments on appeal track arguments preserved below.

A Working Definition. A Copilot Notebook is a curated, scoped workspace that tells Copilot exactly which documents and references to consider, and implicitly tells it to ignore everything else. It is the single most important feature for keeping AI useful in legal work.

Building a Notebook: A Step-by-Step Workflow

Setting up a Notebook is straightforward, but doing it well requires the same discipline that goes into building a litigation binder. These same principles are applicable to other enterprise-level platforms such as Harvey or Claude. The steps below describe how I build a Notebook for a new matter:

- **Define the purpose first.** Decide what the Notebook is for — discovery, a specific motion, deposition prep, or the entire matter. A Notebook with no defined purpose collects everything and is useful for nothing.
- **Curate the source documents.** Add only the documents that belong inside the purpose. Operative pleadings, current discovery responses, key correspondence, controlling rules. Resist the urge to dump the entire matter file.
- **Add reference authority where helpful.** For litigation Notebooks, this may include your Rules of Civil Procedure, the local rules of the court, or any standing order specific to the judge. Links to public versions of these materials work well.
- **Set personalization preferences.** Notebooks allow you to specify a default tone (formal, neutral, persuasive), preferred output format (bullet, narrative, outline), and preferred level of detail for the notebook itself. Set these and the outputs will take less work or iteration to perfect.
- **Test with known questions.** Before relying on the Notebook, ask it questions to which you already know the answer. This is how you verify that the scope is correct and the sources are complete.
- **Use it, refine it, and prune it.** As the matter evolves, add new operative documents and remove materials that are no longer relevant. Stale Notebooks produce stale answers.

Notebook Discipline: What to Keep In, What to Keep Out

Curation is the discipline that distinguishes a useful Notebook from a digital junk drawer. The following table summarizes how I think about what belongs inside a litigation Notebook and what does not:

Keep In	Keep Out
Operative pleadings and the most recent amended versions.	Superseded drafts and outdated pleadings.
Current discovery responses and supplementations.	Earlier responses already replaced by supplementation.
Key correspondence that establishes positions or deadlines.	Routine scheduling emails that add noise without substance.
Expert reports disclosed in the case.	Draft reports from the expert that were not produced.
Court orders, scheduling orders, and any standing orders.	General research saved “just in case” from unrelated matters.

Audio Summaries and Other Notebook Skills

Notebooks also support audio summaries. The lawyer can ask Copilot to produce a spoken overview of the materials in the Notebook, which can be listened to in the car or between meetings. This is particularly useful when returning to a matter after a period of inactivity — such as a case that went dormant during discovery or a file picked up after

a colleague's departure. Rather than re-reading the entire file, the lawyer can listen to a concise summary and then drill into specifics in the Notebook itself. Copilot has also recently added additional Notebook skills, including the ability to generate: summaries of materials, PowerPoint presentations, mind maps, and other services. These are in testing mode but so far appear to be very helpful.

Common Pitfalls to Avoid

- **Overstuffing the Notebook.** A Notebook with five hundred documents is not a Notebook — it is a search engine. Keep the scope tight and intentional.
- **Mixing matters.** Never combine two clients' materials into a single Notebook, even for what seems like a related issue. The whole point is scope discipline.
- **Forgetting to update.** A Notebook frozen at the time of filing will quietly give wrong answers six months later when the case has moved on.
- **Treating answers as final.** A grounded answer is still a draft. The same human-review obligation that applies to every Copilot output applies here.

Copilot Agents: Automating Legal Workflows

An agent is a purpose-built assistant that lives inside Copilot Studio. It is given a defined job, a set of knowledge sources it may consult, and a set of instructions for how to behave. Once published, it can be invoked repeatedly — by the lawyer who created it or by anyone in the firm with access — to perform that same job consistently. Agents are running in my own litigation practice today and have meaningfully changed how work get done.

Note — Agents or workflows can be built in Harvey or Claude and most other enterprise-level AI platforms.

Three Things Agents Are — and Three Things They Are Not

It helps to be precise about what an agent is, because the word “agent” is overused in AI marketing. In Copilot Studio specifically:

- **An agent is a configured assistant** with a fixed purpose, a defined knowledge base, and explicit instructions about how to respond.
- **An agent is reusable** across matters and across users in the firm, producing consistent output every time it is invoked.
- **An agent is auditable.** You can see what it was told to do, what sources it relied on, and what it produced.

And just as importantly:

- **An agent is not an autonomous.** It does not exercise judgment, and it does not replace the human who supervises it.
- **An agent is not a free-form chatbot.** It is scoped, instructed, and intended to do a defined job — not to do whatever the user asks.

Agent, Scheduled Prompt, or Chat Session?

A practical question is when to use an agent versus a scheduled prompt versus an ordinary Copilot chat. The simplest way to think about it:

Mode	Best For / Example
Chat session	One-off questions, exploratory drafting, or quick summaries. Example: “Summarize this email thread” in Outlook.
Scheduled prompt	A recurring task that follows the same instruction each time. Example: Every Monday morning, summarize last week’s appellate decisions on legal malpractice.
Agent	A multi-step workflow tied to specific knowledge sources, used repeatedly. Example: A Medical Record Analyzer that reads a record set and produces a chronology.

Building an Agent: The Five-Step Process

Building an agent in Copilot Studio follows the same five steps every time, but the quality of the result depends heavily on how carefully each step is executed.

Step 1: Choose a Template or Build from Scratch

Copilot Studio offers a library of templates organized by category, and a blank canvas for custom builds. Templates are useful when an agent will perform a well-defined task that has been done before — a summarization agent, a research agent, a meeting-prep agent. The blank canvas is appropriate when the workflow is firm-specific and does not map cleanly to a template. Try starting from a template the first time and moving to the blank canvas only after you have built two or three agents successfully.

Note — Harvey and Claude each offer a set of pre-built agents and also provide a blank canvas for custom builds.

Step 2: Connect Knowledge Sources

An agent without good knowledge sources is a guessing machine. An agent with focused, current, well-organized knowledge sources is a force multiplier. Knowledge sources can include:

- SharePoint sites or specific libraries within them (firm-wide brief banks, form files, expert files).
- OneDrive folders (a specific case folder for a matter-specific agent).
- Uploaded documents (a single statute, a particular set of standing orders, a procedural rulebook).
- Public websites (court websites, statutes, regulatory pages) for agents that need to reference public authority.

Keep the knowledge sources tight. An agent pointed at “the entire SharePoint” is an agent pointed at nothing in particular.

Step 3: Define Instructions and Actions for your Agent

Instructions tell the agent how to behave. They are written in plain English, and the clearer they are, the more useful the agent will be. A useful instruction set typically specifies:

- **Role.** Who the agent is acting as (“act as a litigation paralegal,” “act as a defense associate”).
- **Job.** What the agent is supposed to produce, in concrete terms.
- **Sources.** Which knowledge sources to use and which to ignore.
- **Format.** How the output should be structured (chronology, outline, table, memo).
- **Tone.** Formal, neutral, persuasive, or matter-of-fact.
- **Boundaries.** What the agent should not do — for example, never produce a legal conclusion, never invent citations, never speculate beyond the documents.

Step 4: Test and Refine

Every agent is run before it is used on live matters. Testing means running the agent against a known input and comparing the output to what an experienced lawyer or paralegal would produce. The first version of an agent is almost never the best version. Refinement typically involves tightening the instructions, narrowing the knowledge base, or adding examples of good and bad outputs.

Step 5: Publish, Monitor, and Improve

Once an agent is performing reliably, it is published so others in the firm can use it. Usage analytics and user feedback feed continuous improvement. Agents that are heavily used and rarely corrected are working well. Agents that are rarely used, or used and then heavily revised, deserve a second look — either the agent needs work, or the workflow itself needs to be redesigned.

Litigation Agent Ideas

What follows are some sample uses for a defense litigation practice. All of them return time to the lawyer that would otherwise be spent on clerical labor — time that is better spent on strategy, judgment, and client communication.

Medical Record Analyzer

Purpose. Reads a complete medical record set, extracts dates of service, providers, diagnoses, treatments, and any contradictory entries, and assembles them into a coherent chronology.

Inputs. Uploaded medical records (PDF, scanned images with OCR, or Word).

Outputs. A timeline of treatment with date, provider, and brief description; a list of inconsistencies and gaps; a separate list of any entries flagged as potentially significant.

Why it matters. It does not replace a lawyer’s review. It enhances and organizes the review so that the lawyer can focus on medical significance rather than clerical labor.

Deposition Analysis Assistant

Purpose. Reads a deposition transcript and extracts high-level themes, key admissions, points of impeachment, and opportunities for further discovery.

Inputs. Deposition transcript and (optionally) the deponent’s prior statements, key medical records, and answers to interrogatories.

Outputs. A theme summary, a list of admissions with line and page references, a list of inconsistencies

with prior statements, and a candidate list of follow-up discovery items.

Why it matters. By the time counsel sits down to prepare the most relevant testimony is already curated and ready for evaluation, enhancement, and confirmation by human review.

Jury Verdict Research Agent

Purpose. Searches the firm's internal verdict database (or a connected SharePoint library of verdict reports) for comparable verdicts based on injury type, venue, and other parameters provided in the prompt.

Inputs. Internal verdict reports stored in SharePoint, plus the parameters the lawyer specifies (jurisdiction, injury, plaintiff age, liability profile).

Outputs. A short comparable-verdict summary table with case names, venues, injuries, and amounts, along with citations to the underlying reports.

Why it matters. Verdict research is repeatable, slow, and high-value. An agent makes the search consistent and the citations auditable.

Daily Inbox Triage Agent

Purpose. Reviews the lawyer's Outlook inbox each morning, identifies new client communications, deadlines, and items requiring partner attention, and produces a one-page triage summary.

Inputs. Outlook inbox, calendar, and a list of active matters.

Outputs. A triage summary with three sections: items requiring response today, items requiring response this week, and informational items.

Why it matters. Email is the largest single drain on a litigator's time. A triage agent run at 7:30 a.m. converts an hour of sorting into ten minutes of decision-making.

Common Failure Modes

- **Instructions that are too vague.** "Summarize this deposition" is not an instruction; it is an invitation for the agent to guess what you wanted.
- **Knowledge sources that are too broad.** An agent pointed at the entire SharePoint will pull in unrelated documents and weaken every answer.
- **No human review step.** An agent's output is a draft. Treating it as a finished product is the fastest way to create a problem that should never have existed.
- **No iteration.** A first-version agent is rarely the best-version agent. Without a refinement loop, the agent's output drifts away from what the lawyer actually needs.

The Cumulative Effect. Together, agents form a kind of silent litigation support team — each stepping in when needed and assisting with the heavy lifting so the lawyer can focus on strategy, judgment, and

client communication. The work becomes not only faster but also more consistent, more organized, and ultimately more valuable. The number of agents to address tasks in the day-to-day practice of law is, in practical terms, unlimited.