



2026 Professional Liability Practice Group Seminar

February 19-20, 2026

Exploring the Practical Applications and Future Trajectory of AI in
Legal Practice: From Cutting-Edge Arbitration and Classifier Models
to Hallucinations and Pro Se Ai Pleadings

Brian Sableman
Moderator
BROWN & JAMES, P.C.
St. Louis, MO
bsableman@bjpc.com

D. Jay Davis, Jr.,
Partner
CLEMENT RIVERS, LLP
Charleston, SC

AI Arbitration

The Vision

Judge.ai represents what has been called “the platonic end result of the textualist project”—the theoretical ideal of algorithmic adjudication where legal rules are applied with perfect consistency, speed, and precision. The insight is simple: on almost every conception of law, consistency is all a judiciary really needs. Not wisdom, not intelligence, not a pulse. Consistency. A judge doesn’t have to be wise; the lawmaker can be wise for him. But he must be consistent—without that baseline, there is no rubric at all.

This is why large language models like ChatGPT fail as judges. Jonathan Choi’s landmark experiment demonstrated that LLMs are hopelessly prompt-sensitive: ask the same legal question 2,000 different ways and you get wildly different answers. LLMs are designed to predict the next likely string of characters, not to connect factual patterns to legal outcomes. When there’s no consensus in the training data—as is typical in adversarial legal disputes—the LLM becomes a grab bag of inconsistent results.

Why Classifiers Work

Arbitrus, the system underlying IDRE-235, is not an LLM. It is a classifier—a fundamentally different architecture. Where an LLM’s purpose is to guess the next string of letters, a classifier’s purpose is to group factual inputs and connect them to legal outcomes based on patterns learned from actual case law. The result is analogical reasoning reduced to its purest form: pattern recognition.

We replicated Choi’s experiment with Arbitrus on the bench. Across 150 question variants, Arbitrus delivered perfect consistency—100% decision uniformity with zero hallucinations. Every time it saw X fact pattern, it issued Y answer. It never saw X and issued anything other than Y. This makes Arbitrus the only judiciary in the world that passes the Choi test.

IDRE-235: Theory Made Real

IDRE-235 is Fortuna’s credential-pending application to CMS to become a certified Independent Dispute Resolution Entity under the No Surprises Act. It is the first real-world deployment of Judge.ai—a supervised automated arbitration system for healthcare billing disputes between providers and insurers.

The system employs a voting ensemble of classifier models trained on federal case law to weigh provider qualifications, market characteristics, patient complexity, institutional factors, and negotiation patterns. These models produce an equity multiplier that adjusts the Qualifying Payment Amount baseline. The final award follows baseball-style arbitration: whichever party’s offer comes closest wins.

The No Surprises Act created a perfect proving ground. The Fifth Circuit’s Texas Medical Association decisions locked CMS into a six-factor test—but with over 3 million cases filed and 363,000+ still backlogged, the system is drowning. IDRE-235 offers one definitive version of the statutory test applied the same way every time. Our Texas federal validation demonstrates exceptional performance: 87.1% accuracy, 0.957 AUC, 92.2% precision.

Human-in-the-Loop

IDRE-235 maintains human supervision. Qualified reviewers certify every AI-generated ruling before issuance, verifying evidentiary accuracy, jurisdictional compliance, and mathematical precision. The human arbitrator signs each final determination.

Add Session Title Here (Page 2 Title Style)

This is textualism’s fulfillment through innovation. The Choi test isn’t just a gotcha for chatbots—it’s the definitive entrance exam for the American judge. Arbitrus passes. And it does it in 30 seconds flat.

Practical AI tools for lawyers

AI has enormous potential in the legal field. Lawyers and clients alike are embracing AI. Of course, some AI tools are better than others. Microsoft CoPilot stands out as a highly capable enterprise solution. AI notetakers and deposition analytics stand out as stand-alone tools.

- CoPilot
- Stand-alone tools: AI notetakers and deposition summaries/analytics

How CoPilot Can Safely and Meaningfully Assist You

Introduction

Over the past couple of years, lawyers have heard promises that artificial intelligence would revolutionize the practice of law. Much of that early hype felt disconnected from the realities of a law firm—confidentiality concerns, inconsistent outputs, limited adoption, and a lack of tools built specifically for legal workflows. Microsoft Copilot represents a meaningful shift. 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 paper explores how Microsoft Copilot can be used responsibly and effectively inside a law firm. It explains what Copilot is, how it works, and the ways lawyers can use it to summarize information, draft documents, manage inboxes, prepare for meetings, and automate routine workflows. It also describes advanced capabilities such as Prompt Coach, scheduled prompts, Copilot Notebooks, and Copilot Agents—tools that move beyond basic chat functions and help attorneys systematize work at scale. The goal of this paper is straightforward: to demonstrate that AI really can help the practice of law—if lawyers understand what it does, how it works, and how to control it.

What Copilot Is—and Why It Matters to the Legal Profession

Microsoft Copilot is an AI assistant embedded directly in core Microsoft 365 applications. Within Word, PowerPoint, Outlook, Teams, Excel, and SharePoint, 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 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 operates in a secure, enterprise environment. It respects Microsoft 365 permissions, Purview controls, and auditing mechanisms. That means it cannot access matters or folders the user cannot access. It also means prompts and outputs remain within the organization’s security framework—an essential consideration for privilege, confidentiality, and ethical duties. For law firms, Copilot is the first widely deployed AI system that works with—not around—the structures lawyers need to protect client information.

Writing Better Prompts: Copilot's Built In Prompt Coach

Copilot's usefulness depends heavily on the clarity of the instructions it receives. To help legal professionals 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 start with vague prompts like: *"Summarize this case."*

Prompt Coach will help and even 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 to optimizing clarity, Prompt Coach checks for Responsible AI considerations—flagging ambiguous instructions, unsupported assumptions, or overly broad requests that may inadvertently pull in irrelevant data. Firms can extend Prompt Coach with custom prompt libraries and allowing standardized prompting across practice groups. The result is not only better AI output but also a more consistent approach to how associates and staff direct AI tools in legal work.

Automating Recurring Work with Scheduled Prompts

Many legal tasks follow predictable weekly or monthly 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 legal examples might include:

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

Copilot Agents: Automating Multi-Step Legal Workflows

In my opinion, Copilot Agents are among the most advanced and useful capabilities available. Built in Copilot Studio, agents are autonomous assistants capable of reasoning through multi-step instructions, accessing knowledge sources, and interacting with applications. To build an agent, lawyers or staff follow five steps:

1. **Choose a Template or Build from Scratch** Copilot Studio offers templates or a blank canvas to create your own.
2. **Connect Knowledge Sources** Agents can be tied to SharePoint libraries, OneDrive folders, or uploaded documents such as case law or procedural rules.
3. **Define Instructions and Actions** Lawyers specify how the agent should behave—what it should look for, when it should respond, and which tools it may use.
4. **Test and Refine** The agent is run in a sandbox environment to ensure accuracy, safety, and compliance.
5. **Publish and Monitor** Once deployed, usage analytics and user feedback help the firm continuously improve the agent.

Agents allow firms to automate recurring or complex workflows that were previously impractical without custom software development. I have used them to assist with repetitive workflows in my daily practice.

Narrative: How Specialized Agents Elevate a Litigation Workflow

In a typical litigation cycle, a lawyer or client may review dozens of documents a day— medical records, client reports, deposition transcripts, expert opinions, email threads, and mediation statements. What once required hours of parsing, summarizing, and reorganizing can now be handled by a series of tightly focused AI agents, each trained to handle a specific segment of the work. *These agents don't replace judgment; instead, they return the time needed to exercise that judgment effectively.* Just a couple of agents I have created.

Medical Record Analyzer: The agent reads the entire record set, extracts dates of service, diagnoses, treatments, and any contradictory entries, and then assembles them into a coherent chronology. While not a substitute for actual review, it can enhance and organize the review of expansive medical records, allowing the partner to focus on medical significance rather than clerical labor.

Deposition Analysis Assistant: When a fact witness or plaintiff is deposed, this agent extracts not only high-level themes but also key admissions, points of impeachment, and opportunities for future discovery. It highlights inconsistencies with prior statements, medical records, and interrogatory answers. By the time counsel sits down to prepare a motion or mediation memo, the most relevant testimony is already curated and ready to for evaluation, enhancement, and confirmation by human review.

Client Report Follow-Up Agent: This agent streamlines follow-up for my defense team. It produces a clean, structured follow-up task list that the partner can circulate and monitor. The result is clear, timely, and consistent communication without consuming billable hours that could be better spent elsewhere.

Together, these 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.

Using Copilot Notebooks to Control Context and Scope

Copilot Notebooks provide a controlled workspace for legal projects, ensuring that Copilot's answers are grounded solely in specified documents and resources. Within a notebook, lawyers can add references—including Word files, PowerPoint slides, Excel spreadsheets, OneNote pages, meeting notes, and links—and then pose questions or tasks. This is especially useful for complex matters where scope must remain tightly confined. For instance:

- A litigation team may build a notebook containing only the pleadings and discovery responses for a particular case.
- A transactional attorney may assemble a notebook of prior deal documents to guide drafting.
- A partner may use a notebook to prepare a damages outline, deposition sequence, or case chronology.

Notebooks also allow personalization of tone and format and offer audio summaries to help attorneys quickly refresh on materials—especially useful when returning to a file after some time. In essence, it is a secure way to have an “AI conversation” with all the documents in your notebook or case.

Email and Calendar Efficiency with Copilot in Outlook

Email management is one of the most time-consuming aspects of legal practice. Copilot in Outlook significantly reduces that burden. Copilot can:

- **Summarize email threads** and identify deadlines, client instructions, and open issues.

Add Session Title Here (Page 2 Title Style)

- **Draft replies** in a specified tone—formal, persuasive, neutral—and adjust the length based on the user’s preference.
- **Provide one-click thread summaries** directly within Outlook.

For example, an attorney might instruct: “Draft a polite follow-up requesting the signed engagement letter by Friday, in a formal tone.” Copilot can generate a draft reply for the attorney to review and edit. It is critical, however, to maintain human oversight. Lawyers must review all outputs for accuracy, privilege, and tone before sending—an ethical safeguard that cannot be outsourced to AI.

Best Practices and Security Considerations for Law Firms

To ensure responsible and effective use of Copilot, law firms should adopt the following practices:

1. **Be Specific in Every Prompt:** Include the audience, format, scope, and applicable sources. Vague prompts produce vague results.
2. **Use Grounding Discipline:** Notebooks and controlled libraries help keep Copilot’s scope narrow and prevent irrelevant or privileged documents from entering the context.
3. **Maintain Human Oversight:** Every AI-generated output—whether a draft email, brief, or summary—MUST be reviewed carefully for accuracy, privilege, tone, and factual grounding.
4. **Train Your Teams:** Develop prompt galleries, templates, and firm-wide standards. They will benefit enormously from examples of strong, legally structured prompts.
5. **Rely on Enterprise Security:** Copilot uses Microsoft’s built-in protections, including encryption, DLP policies, and compliance guardrails.

Conclusion

Microsoft Copilot brings practical AI capabilities directly to the tools lawyers use every day. Its value lies not in replacing attorneys or hard work, but in reducing time spent on routine, repetitive, or organizational work. By using Prompt Coach to write better prompts, schedule recurring tasks, create agents, organize matter-specific notebooks, and leverage Outlook’s drafting and summarization tools, lawyers can improve efficiency and increase their focus on high-value legal strategy.

As with any new technology, adoption requires education, oversight, and thoughtful implementation. But when used properly, Copilot can meaningfully enhance modern legal practice—fulfilling, at long last, the promise that AI can help.

Stand-alone AI tools

AI notetakers

AI notetakers automate the process of capturing, organizing, and summarizing information from meetings and calls. Summaries condense long conversations into concise highlights, key points, and action items. Action tracking identifies tasks, deadlines, and responsibilities mentioned during meetings. Microsoft CoPilot Notes is built into Teams and Outlook. Many third-party ai notetakers, like Read.ai, integrate with CoPilot.

AI notetakers have many applications for lawyers and law firms. Lawyers can utilize AI notetakers for new client meetings, in which topics like case strategy, next steps, and specific tasks are usually discussed. AI identifies the core information and tasks which can easily be assigned to team members (associates). Witness interviews can be efficiently analyzed for themes and contradictions. Users can ask CoPilot to perform tasks based on the meeting notes, like drafting a follow-up email.

AI deposition tools

AI deposition tools help lawyers by automating transcript review, surfacing key testimony, spotting contradictions, and generating summaries that save hours of manual work. Many stand-alone AI deposition platforms charge a flat fee per deposition or a flat fee depending on the number of pages. Accordingly, unlike legal research or enterprise tools like CoPilot, AI deposition analytics can be purchased on a case-by-case basis, which gives law firms the ability to pass on costs to clients.

Discussion

AI Arbitration

AI arbitration and Arbitrus.ai is here to stay. Would you use arbtirus.ai to resolve a dispute?

AI tools

The expanding ecosystem of AI tools available to legal professionals' presents a variety of questions: How much AI is too much in practice? What should clients reasonably expect, and pay for, when AI is involved? How do lawyers abide by their ethical obligations when utilizing AI. And does AI truly deliver cost savings, or simply shift the workload?

Implications for liability insurance

Major insurers have recently sought permission from US regulators to offer corporate policies which exclude liabilities tied to businesses deploying AI chatbots and agents. Should insurance companies be allowed to exclude AI-related liabilities? The concern is that an insurer may be able to pay a \$100M loss to one company that had one agent deliver one incorrect medical diagnosis for example. But the insurer cannot afford a mistake that is replicated by the agent 1,000 times or 10,000 times before it is discovered. Will AI be uninsurable or how can this issue be resolved?

Claims involving mistakes related to AI are growing. For example, employees have uploaded medical records into ChatGPT and asked it to print out summaries, despite receiving training on HIPAA privacy requirements. Most companies now have AI policies and are doing AI training. What additional training is needed? What should insurers require?

AI is increasing the amount of insurance fraud significantly. Fraud increases costs for insurers which is ultimately passed down to consumers. Can insurers use AI to combat this increase in fraud?