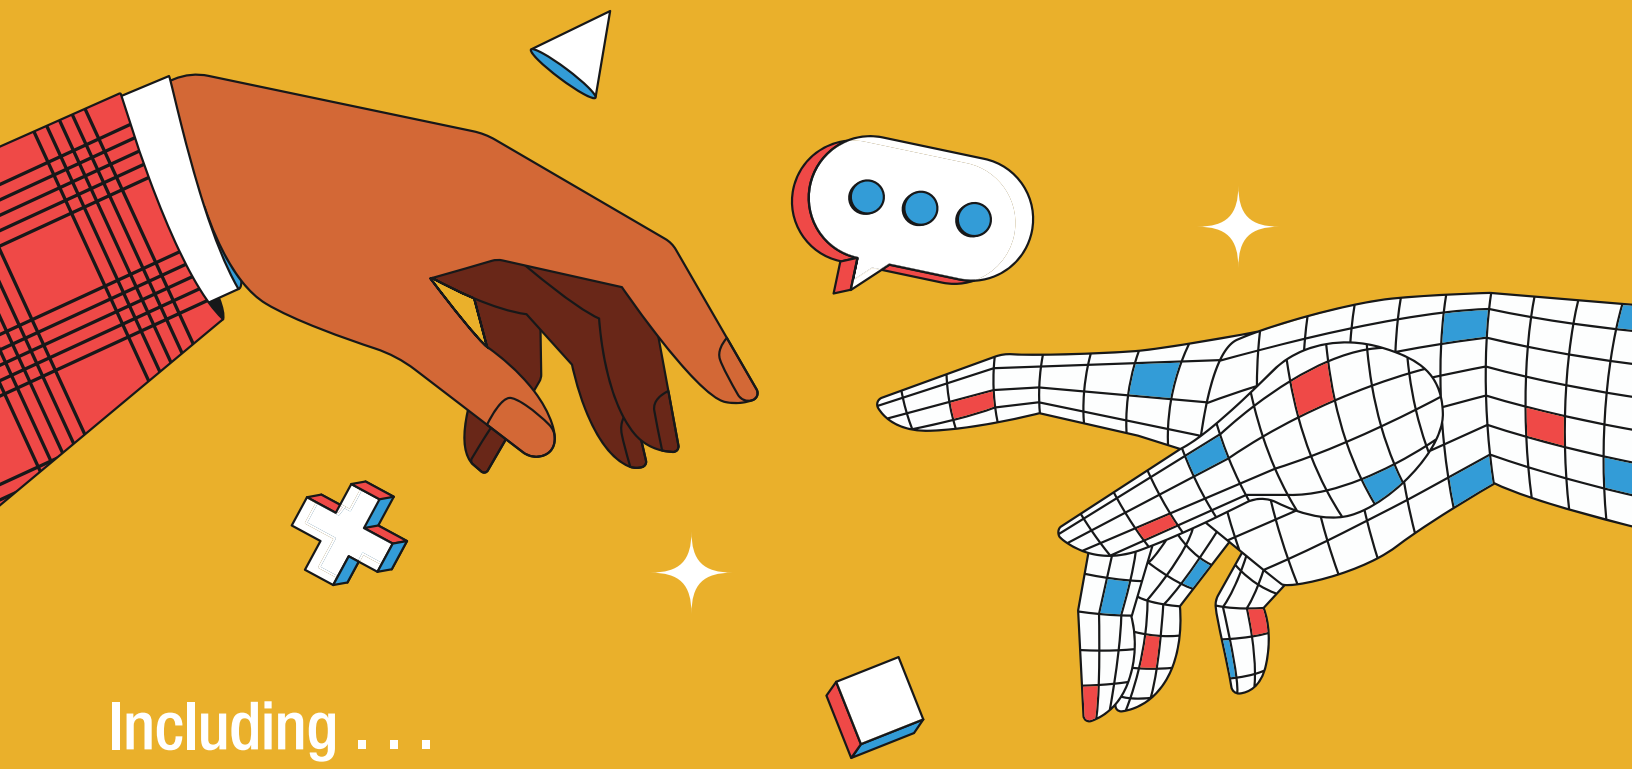
For The Defense™

dri
The magazine for defense, insurance and corporate counsel

February 2026

Artificial Intelligence



Including . . .

Defending the Algorithm™: The AI Daily Routine

Also in This Issue . . .

Lace Up Your Sneakers + Litigate: The Marathon as a Metaphor for the Life of a Lawsuit

And More!

THE CENTER
The Voice of the Civil Defense Bar.

Defending the Algorithm™*

By Henry M. Sneath, Acacia B. Perko, and Christopher M. Jacobs

The question isn't whether to develop AI expertise, but how quickly you can do so while fulfilling professional and ethical obligations to clients.

The AI Daily Routine
An Artificial Intelligence Practicum for Defense Attorneys

A Publication of the DRI Center for Law and Public Policy AI Task Force - Part 2

Please attend our live CLE full-day seminar on March 10, 2026 at DRI Headquarters in Chicago entitled:
Artificial Intelligence in Defense Practice: Tools, Prompts, Workflows, Implementation and Governance

In Part 1 of this series, we applied Bayesian probability theory to predict AI litigation trends and analyzed landmark cases shaping the legal landscape. Part 2 pivots from theory to practice, providing defense attorneys with a detailed practicum for implementing AI tools in daily litigation workflow. The focus is not theory, but execution—information management, workflow integration, and rigorous verification—so that AI enhances legal services while reinforcing (not undermining) professional responsibilities. This article was authored with assistance from Claude® Opus 4.6 Max and research confirmation by Google Gemini 3.0 Pro and Westlaw Precision with AI and Analytics, but edited by humans as AI can make mistakes.

** The term Defending the Algorithm™ is used and owned by Houston Harbaugh, P.C., and a registered trademark protection is being sought.*

Introduction

Defense attorneys repeatedly ask: How do you actually use AI tools in daily legal practice? Part 1 of this series examined the current state of substantive law—copyright litigation, trade secret protection, Section 230 immunity, and product liability frameworks emerging from landmark cases like *Bartz v. Anthropic and Raine v. OpenAI.* We applied Bayesian probability analysis to predict litigation trends and assess risk across multiple practice areas. Warning: none of the advice below will work unless you and your IT team plan, coordinate, practice, train and build a proper tech stack. Without a substantial investment of your time, the use of AI will be disappointing.

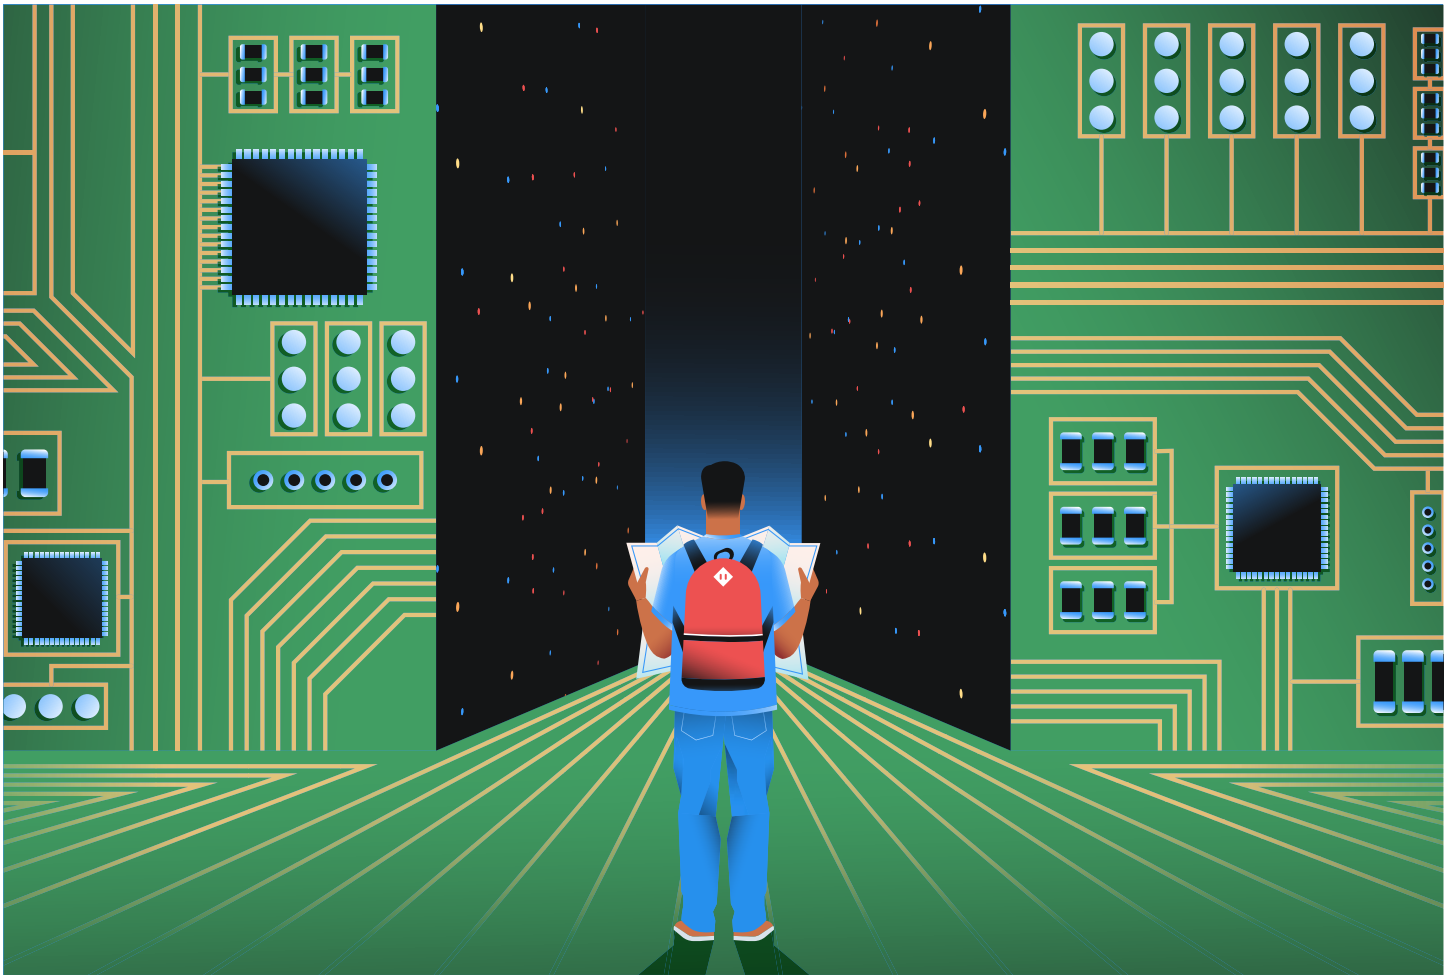
Let's turn to practical implementation. This article provides a detailed practicum for integrating AI into defense practice—from building the right technology infrastructure to mastering prompt engineering, from leveraging AI-enhanced legal research to managing massive e-discovery projects. The goal is not to replace human judgment but to augment it: AI assists, but doesn't supplant, the ethical







Henry M. Sneath is a Director/Shareholder in the firm of Houston Harbaugh, P.C. in Pittsburgh, Pennsylvania. Mr. Sneath is a Business, IP and AI Trial Lawyer and a former President of DRI. He focuses his practice on intellectual property litigation, trade secret misappropriation, AI law, and complex commercial, employment and insurance coverage and bad faith disputes. Mr. Sneath is currently serving on the DRI Center for Law and Public Policy AI Task Force and is working to develop comprehensive AI education programs to be presented by the DRI Center for Law and Public Policy. **Acacia B. Perko** is a Director/Shareholder at Houston Harbaugh, P.C. and practices Business, IP, Trade Secret and Employment litigation, and serves as DRI's Trade Secrets Substantive Law Group Chair for the Intellectual Property Litigation Committee. **Christopher M. Jacobs** is a Director/Shareholder at Houston Harbaugh, P.C. and practices Insurance, Business and IP litigation and serves on the DRI AI Task Force. Together, these authors record and publish a blog and podcast series entitled: Defending the Algorithm™: A Bayesian Analysis of AI Litigation and Law and they are frequent speakers on this topic and related IP issues.



reasoning and client relationship management that define our profession.

The goal is not to replace human judgment but to augment it

For civil defense attorneys, corporate counsels and insurance claims professionals, AI competence is rapidly becoming table stakes. Within five years, lawyers who competently use AI tools will be the norm rather than the exception. Applying our Bayesian framework from Part 1: P(Competitive Disadvantage | No AI) = Very High. Put plainly: once opposing counsel uses AI-assisted research and discovery analytics, the risk of competitive disadvantage rises—especially in document-heavy litigation where speed, pattern recognition, and iteration materially affect outcomes. The question isn't whether to develop AI expertise, but how quickly you can do so while fulfilling professional and ethical obligations to clients.

Part I: Understanding AI—The Probability Engine

A. How Large Language Models Actually Work

Before implementing AI tools, defense attorneys need a functional understanding of what they are—and are not. Large Language Models (LLMs) generate text by predicting likely language patterns based on training data; they are not traditional databases and do not "retrieve" facts the way Westlaw or Lexis does. Practically, this means LLM outputs should be treated as draft analysis that can accelerate research and drafting, but must be grounded in verified sources and lawyer judgment.

The critical point is practical: LLMs generate plausible language, not guaranteed truth. Because AI is fundamentally a probability engine, it predicts the most probable next word in a sequence based on patterns learned from trillions of examples in its training. It is not reasoning like a human lawyer—at least not yet. Understanding this distinction is essential for effective use. That is why context and prompting matter—and why verification protocols are mandatory in legal work.

Trial lawyers already use Bayesian reasoning every day—weighing evidence, assessing probabilities, predicting outcomes. The Bayesian conditional probability at the heart of AI is expressed as: P(next word | previous words). Each word the AI generates is its best probabilistic prediction given everything that came before. This is why context matters so much, and why sophisticated prompting dramatically improves output quality.

The LLMs most relevant to legal practice include Claude (developed by Anthropic), ChatGPT (OpenAI), Gemini (Google), and Copilot (Microsoft). Legal-specific platforms include Westlaw Precision AI and Lexis+AI. A key distinction: generative LLMs create new text, while legal databases retrieve existing documents. Best practice demands using both together—leveraging generative AI for analysis and drafting while grounding outputs in verified legal databases.

B. Tokens, Context Windows, and Why They Matter

AI systems measure and process text in 'tokens'—chunks of approximately four characters or three-quarters of a word. One thousand tokens equal roughly 750 words. Understanding token economics is essential for effective AI use, as most platforms limit token usage in various ways.

The 'context window' is the AI's working memory—how much text it can consider at once. Claude offers up to 200,000 tokens (approximately 150,000 words) per chat, while Google Gemini's advanced version offers one million tokens. Larger context windows, which cost more subscription money, enable analysis of entire contracts, depositions, or case files in a single interaction—a gamechanger for legal work.

Practical tip: start new chats frequently but within the same "project folder." Many AI platforms, including Claude, can pull from "project memory" without consuming additional tokens in the active conversation. This allows you to maintain institutional knowledge across sessions while preserving token capacity for immediate analysis. We have project folders in Claude and the other AI platforms like "Defending the Algorithm Podcast Series" and "DRI AI Task Force," to store chats that relate to each other.

C. Hallucinations and the 'Trust but Verify' Protocol

A 'hallucination' occurs when AI generates plausible-sounding but false information—fabricated case citations, fake quotes, invented facts. This happens because AI predicts 'likely' text, not necessarily 'true' text. If the statistical patterns suggest a particular type of case citation should exist, the AI may generate one that sounds authentic but corresponds to no actual opinion.

The most insidious trap for a time-pressed practitioner is the hallucination echo chamber fueled by "circular logic." LLMs sometimes hallucinate a fact and then "verify" that fact by hallucinating a citation to support it. Because the LLM is a probability engine—not a fact-checker—it often prioritizes internal conversational consistency over external reality. Once the AI has probabilistically "decided" that a legal proposition is true, it may double down by generating a detailed but entirely fictional case summary to reinforce that conclusion. This is not deception; it is the AI fulfilling its mandate to be helpful by completing the pattern it has already started.

Unverified reliance on AI output is unacceptable in legal practice, and the consequences for lawyers can be severe. Multiple attorneys all around the country have faced sanctions for citing non-existent cases generated by AI. Malpractice risk attaches to any attorney who relies on fabricated information without verification. **Professional responsibility demands that every AI output be independently confirmed.**

The 'Golden Rule' of AI in legal practice: Never use an AI-generated citation without checking it in Westlaw or Lexis first. If you cannot find the pinpoint cite, the case does not exist. Pro tip: If you suspect "circular logic" in an output, ask the AI 'Are you sure this case exists?' Sometimes this second check triggers a self-correction where the AI admits its error. Also - consider using, for example, Google Gemini to proofread, check and verify citations and opinions generated by Claude or vice versa. We frequently send each other content generated from our favorite AI (e.g.: Claude) to another author herein to check it out on ChatGPT or Gemini. Some of us use two or more AI platforms for this purpose and to seek different perspectives. All of these platforms have nuances, strength and weaknesses that make them somewhat different from each other.

Retrieval-Augmented Generation (RAG) significantly reduces hallucination risk. RAG systems combine AI generation with document retrieval—the AI searches actual documents first, then generates answers grounded in real sources. Westlaw Precision AI uses RAG to cite genuine cases. Claude Projects module can search uploaded documents. RAG is why legal-specific AI tools are generally more reliable than general-purpose platforms—they're grounded in verified legal databases rather than pattern-matching from internet training data. We return to these verification requirements throughout this article because they apply across the workflow—from initial research memos to filed briefs and discovery positions.

Retrieval-Augmented Generation (RAG) significantly reduces hallucination risk.

D. The 'Black Box' Problem in AI Decision-Making

AI decision-making processes are often opaque—even to their creators. This 'black box' problem means we cannot fully explain 'why' AI reached a particular conclusion. The AI neural network's weights, parameters and probability calculations that produce output are not interpretable in human-readable terms. This is changing the fundamental nature of litigation discovery and even making routing ESI protocols outdated and inadequate. Plaintiffs are requesting production of the actual algorithm and the meta support structure of the black box.

For litigation, this creates significant challenges. In insurance bad faith cases like *Estate of Lokken v. UnitedHealth Group*, where the class plaintiffs seek production of the algorithm, courts must evaluate alleged AI-driven coverage decisions while respecting the insurer's proprietary algorithms. Employment discrimination claims challenge algorithmic hiring decisions that cannot be fully explained and which are alleged to contain bias. Discovery in AI-related cases must balance plaintiffs' need for algorithmic transparency against defendants' trade secret protection.

Understanding the black box problem is essential for both defending AI-related

claims and explaining AI tool reliability in your own practice. Some AI platforms, like Google Gemini, offer 'Deep Thinking' modes that display notes about reasoning processes as analysis proceeds—excellent for understanding methodology if you ever need to defend the algorithm.

Part II: The Foundation Problem—Why AI Implementation Often Disappoints

A. It's Not the AI—It's Your Information Architecture

Many lawyers trying AI tools find the results underwhelming. After years of intensive AI adoption in practice, we've observed that disappointing outcomes rarely stem from inadequate AI technology. Instead, the problem typically lies in how firms structure their information and workflows before introducing AI—and the sophistication with which lawyers craft AI prompts.

Consider how experienced litigators naturally work. When evaluating a case, we don't examine documents or evidence in isolation. We instinctively connect dots across multiple dimensions—linking factual allegations to procedural requirements, assessing how specific judges have ruled on similar issues, considering the opposing party's litigation history, and evaluating how various legal theories align with our client's broader business interests. This holistic analysis draws on years of accumulated knowledge about how different case elements interconnect.

AI systems need access to similar interconnected information to provide useful analysis. Yet many law firms offer AI tools in a fragmented landscape: separate systems for document storage, disconnected email archives, case management software that tracks deadlines but not strategic context, time-keeping databases with no link to substantive legal work, and research platforms operating in complete isolation from active matters.

This fragmentation represents the fundamental challenge. When AI can only access isolated pieces of information from disconnected systems, it cannot provide the integrated analysis that makes it valuable. It's like hiring a brilliant associate but only allowing them to read one category of documents while keeping everything else locked away.

B. Infrastructure First: The Prerequisite for AI Success

Before worrying about which AI platform to purchase, firms need to examine how information flows—or fails to flow—through their practice. Where do gaps exist? Which systems don't communicate with each other? What critical context gets lost when lawyers manually transfer information between platforms?

The most successful AI implementations share a common characteristic: they occurred in firms that first invested time organizing their information infrastructure. These firms established consistent naming conventions, created centralized matter repositories, developed standardized intake procedures, and built connections between previously siloed systems. Only after laying this groundwork did they introduce AI tools—which then performed remarkably well because they finally had access to the connected information they needed.

Applying our Bayesian framework: P(lawyers will achieve AI competence to maintain their duty of technological competence under ABA Rule 1.1 | The need for a dedicated approach to building smart and communicative data architecture) = the amount of hard work firms are willing to invest in making it happen correctly. In other words, AI competence is less about buying tools and more about doing the unglamorous work—organizing information, standardizing workflows, and training lawyers to use AI with disciplined verification.

This requires investments in unglamorous infrastructure work that doesn't generate excitement at partnership meetings. It means consolidating systems, standardizing processes, and organizing information before deploying the latest AI features. It means teaching teams consistent practices rather than allowing everyone to develop their own disconnected workflows. But firms that make this investment will discover that AI actually delivers on its promise—not because they bought superior AI, but because they finally gave AI the organized, connected information it needs to provide useful analysis.

Part III: Building the Technology Stack

A. The Integrated Daily Workflow

Using your static IT tools and dynamic AI platforms in a coordinated and purposeful way makes all the difference. The typical day usually begins with opening several software systems on your network and internet. This is standard for most lawyers: Outlook, case management software, document management systems, and perhaps legal research platforms depending on the day's plans. What differs is the strategic coordination of these tools with AI platforms to create an integrated workflow.

Important Note on Confidentiality: Do not put any client or case-related data on any public AI platform. However, you can use these platforms to refine language or presentations in a non-named, case-generic way. Enterprise versions with proper data protection agreements may be appropriate for sensitive work, but evaluate each platform's security certifications carefully. You and your firm need an AI Governance Policy and a thorough understanding of your user agreements and Terms of Service agreements with your providers like Westlaw, Lexis, Claude, Gemini and ChatGPT. What attorney client privilege and work product protections do you have when using AI? You need to know!

A practical note: effective AI adoption does not require mastering every platform discussed below. For most defense practices, a workable baseline includes (1) a generative AI tool for analysis and drafting, (2) a closed legal research database with AI augmentation, and (3) a secure document/document management and discovery workflow. Additional tools can add speed and insight, but the baseline framework—and rigorous verification—matters more than any single vendor.

The AI Daily Routine Framework organizes AI integration across the litigation lifecycle. In the outset of a case, use AI for case intake analysis and research planning—front-end loading the research and analysis. For ongoing research, leverage Westlaw Precision AI for legal research and Westlaw Precision "Quick Check" module for brief review. During drafting phases, employ AI-assisted document creation and

editing. For discovery, use Relativity AI for document review batches. Throughout, AI can assist with client communications and memo preparation. The goal is augmenting the legal practice—not replacing it with AI.

B. Essential Platform Components

(The platforms below are examples; firms should evaluate tools based on use cases, confidentiality, integration, jurisdiction, and workflow needs.)

Google Gemini: Regular Google at the premium level provides built-in low-level AI that summarizes search results and provides narrative to focus research. Google Gemini is far more sophisticated and utilizes Deep Thinking Mode to conduct research, draft documents, perform agentic functions, generate images, excel spreadsheets, audio, and multimodal content making it valuable for creating documents, charts, PowerPoints and other presentation materials. It serves as a primary resource for getting started with ideas on legal research projects or quick checks of the law in particular jurisdictions, and as a secondary resource to fact-check other AI-generated content. Various internal modules put Gemini into "Research' mode," "Image Mode," "Video (VEO) Mode," or "Deep Thinking' mode" —useful for complicated prompts where you want to see the reasoning process. Using Google Gemini can create PowerPoint slides from a written presentation that you drafted and will adopt the specific PowerPoint (or Google Slide) format that you instruct it to use. Many people are using ChatGPT in similar ways.

Westlaw Precision AI-Assisted Legal Research with CoCounsel Premier, Westlaw People Map, Company Investigator, Drafting Assistant, and Practical Law AI Connect creates a robust platform that augments legal analytics and client service. General AI tools like Gemini and Claude admit they cannot find certain case opinions and analysis because they lack access to closed databases like Westlaw, Lexis, and Bloomberg Law. This is why working with both generative AI and Westlaw together produces superior work product.

Westlaw produces Litigation Analytics that could never be generated by human legal researchers—even with unlimited time. Research can determine case histories and likelihood of certain rulings by judge. Litigation history, experience level, and results of opposing counsel can be researched. Analytics are available by court, judge, lawyer, and case type. Reports can be downloaded or emailed directly to clients and team members.

An especially important Analytics tool is the 'References' search for opposing counsel. For experienced lawyers, there may be over 1,000 case references from multiple federal and state court dockets where judicial opinions and rulings are posted. Reviewing this history helps evaluate your opponent's capabilities and tendencies regarding types of pleadings filed, results obtained, and experience in various substantive areas.

The Westlaw Precision 'Quick Check' module allows AI analysis of draft briefs or opponents' briefs. It can provide friendly review to strengthen your own brief, adversarial review to critique your brief or your opponent's, or 'Judicial Review' mimicking analysis done by courts or law clerks. Multiple briefs can be entered simultaneously for complete review—like having mock oral argument on the issue.

Claude® by Anthropic has become a workhorse LLM for our legal practice, particularly in business development efforts. Claude can build project-specific databases for legal issues on which you consistently work—IP, patents, trade secrets—and documents, case opinions, and articles can be uploaded for analysis. When writing on a topic, Claude can work within the same database project to analyze and assist. All case citations must then be checked in Westlaw, and certification should confirm that competent human eyes have reviewed all final results. Some courts now require sworn certificates regarding AI use and human review. Claude, like Gemini, has various modules which allow generation of text, reasoning, presentation materials, slide decks and images.

Microsoft Outlook, Teams, Ring Central, and NetDocs guide the day with basic functions and AI tools. NetDocs stores case files and client documents. Microsoft Teams integrates with Ring Central for communications. Microsoft Copilot assists with programming, management of communications, and video meetings. AI voice recognition generates recordings, notes, and meeting summaries where permitted by law. NetDocs is the backbone of case and document management, while network drives handle large document storage and Relativity manages document review, coding, tagging, and production.

Part IV: Mastering Prompt Engineering

A. The Art of Effective Prompting

A 'prompt' is the instruction or question you give to an AI system. The quality of the AI output from an AI platform depends heavily on the quality of your prompt. This relationship is so direct that an entire discipline—prompt engineering—has developed around optimizing AI interactions.

The fundamental rule: garbage in, garbage out. A vague prompt yields a vague response. Investing time crafting detailed prompts produces substantial returns. Several techniques in prompting the model dramatically improve results:

- **Be specific:** "Analyze this attached contract for indemnification clauses under Pennsylvania law" produces far better results than "review this contract."
- **Provide context:** "I am a defense attorney reviewing a lawsuit claiming breach of fiduciary duty" - helps the AI calibrate its response appropriately.
- **Set format instructions for the model:** "Provide your answer in bullet points" or "structure your response as a memo" guides AI output organization.
- **Chain-of-thought:** "Think step by step" prompts the AI to show its reasoning, often improving accuracy on complex analysis.
- **Set memory instructions:** "Always number (v.1 v.2) and show the date of the draft on each version of any document you create for me" establishes persistent preferences.

B. Legal-Specific Prompting Strategies

Legal work demands specialized prompting approaches. For jurisdiction specification, always identify the relevant jurisdiction: "Under California employment law..." or "Applying Federal Rule of Civil Procedure 56..." Without this specification, AI may draw from multiple jurisdictions or cite inapplicable precedent.

For fact pattern analysis, structure prompts to separate facts from legal questions: "Given these facts: [facts]. Analyze

the following legal issues: [issues]." This mirrors the analytical structure of legal memoranda and produces more organized output.

For risk mitigation, prompt for counterarguments: "What are the strongest arguments against this position?" or "What facts would most undermine this defense?" AI excels at adversarial analysis when specifically prompted.

Quality control through iterative prompting is essential. After receiving initial output, follow up: "Are there any exceptions to this rule?" or "What additional cases support this analysis?" Multiple rounds of prompting typically produce more comprehensive and accurate results than single-shot queries. The output can be graded by you with thumbs up or thumbs down buttons and the chance to submit written comments in response to output. This is human level training of AI.

C. Trade Secret Forensics: The Confidentiality-First Protocol

In the high-stakes environment of trade secret litigation, prompt engineering evolves from a drafting convenience into a litigation-support and risk-assessment tool. When defending against an expedited TRO or managing a departing employee dispute, the ability to rapidly correlate technical behavior with statutory elements—such as whether a company employed "reasonable measures" to protect its information—can provide a meaningful strategic advantage. Structured prompts can be used to perform a targeted "gap analysis" between raw forensic findings and the legal standards that will govern emergency relief, discovery disputes, and early dispositive motion practice.

At the same time, AI-assisted trade secret analysis presents heightened confidentiality and privilege risks. To leverage AI without waiving privilege or jeopardizing the trade secret status of the asset itself, defense counsel must operate under a confidentiality-first protocol. This typically requires either the use of a closed-universe, enterprise AI platform governed by a negotiated data protection agreement (such as System and Organization Controls SOC 2-compliant systems), or a rigorous sanitization workflow. Under a sanitization approach, all personally identifiable information and proprietary technical details are scrubbed, and the AI is used to analyze the structure, sequence, or logic of the conduct at issue—rather than the sensitive data itself.

Practice Tip: This approach allows counsel to benefit from AI-assisted pattern recognition and analytical acceleration while preserving privilege, maintaining ethical compliance, and avoiding inadvertent disclosure of the very assets the litigation seeks to protect.

Part V: AI-Enhanced Legal Research—The Front-End Protocol

A. Priming the Research Pump

The 'front-end research protocol' transforms case assessment by quickly conducting initial legal analysis to prime the research pump at minimal client cost and create more accurate strategic assessments. Upon case opening, the scope of legal research can be determined in minutes and a full preliminary report on relevant case law by jurisdiction or judge can be emailed to team members to be checked by a human reviewer and expanded upon as the case moves forward.

Rather than sending an associate on a random search, client money can be saved and better results produced by assigning research that has already been jumpstarted for the case in general or for a specific Motion to Dismiss or Motion for Summary Judgment. This front-end approach allows litigation teams to make more accurate predictions to clients at the outset, getting analysis moving forward efficiently.

The research pump can be primed at very little cost, and strategic case assessments can be created and more quickly agreed upon. Coordination between AI platforms and closed legal databases produces superior work product. The discipline to work cross-platform ensures research accuracy and excellent writing.

B. Litigation Analytics for Strategic Advantage

Westlaw Litigation Analytics provide intelligence that informs deposition and trial strategy. Judge analytics reveal patterns in rulings, motion grant rates, and time to disposition. Opposing counsel analytics expose litigation tendencies, preferred strategies, and historical success rates.

This analytics capability enables preparation that was previously impossible. Before taking a deposition, you can review how the opposing attorney has conducted prior examinations. Before filing a motion, you can assess how the assigned judge has ruled on similar motions. Before trial, you can understand the judge's preferences and the opposing counsel's trial record.

Use AI to brainstorm defenses, test case theories, and evaluate litigation strategy while being very careful to maintain work product privilege. The AI and privilege issue have become a hot button topic which we will write about in depth down the road. The iterative process of prompting AI, reviewing output, and refining analysis often surfaces considerations that might otherwise be missed in traditional research.

Part VI: Discovery in the AI Era

A. Managing Massive Document Projects

Managing large document projects for gathering, housing, sorting, batching, reviewing, coding, tagging, and producing documents is an art. The plan and likely cost must be included in the first substantive conversation with a client. This process management by a lawyer is a valuable skill because if done right, it often gives your side a huge advantage if the other side is not so savvy with large document management.

If it goes wrong, it can spoil a case both on the cost side and the results side. It also impresses judges when defending your discovery compliance, if it is challenged. Judges now mandate and enforce detailed written ESI protocols as a routine part of federal cases and have little tolerance for sloppy or old-style document management.

Many firms work with dedicated e-discovery vendors rather than housing Relativity in-house. The advantages include large staffs in multiple time zones available to create and customize searches and productions. At the front end of a trade secret misappropriation case, for example, vendors forensically examine, download, and maintain chain of custody of all relevant data—like that of a departing employee believed to have downloaded company con-

fidential information. AI can work both on the forensic end, and the document production end of the process.

B. Active Learning AI in Document Review

Technology-assisted review, including TAR 2.0 continuous active learning systems and generative AI review tools, reduces massive document sets to likely responsive documents through pattern recognition and feedback. AI Active Learning in Relativity continually refocuses on the most important and most likely responsive documents, eventually feeding only the most relevant materials to reviewers.

The goal is to meet a certain probability percentage that sufficient relevant or responsive documents have been gathered to comply with discovery rules and pass muster with opposing counsel or the court. These statistics vary depending on the case and the volume and nature of documents. The AI recognizes patterns and provides feedback and reports that enable supervisory review of the process.

If AI reduces a database from 3,000,000 pages to 1,000,000 pages of the most likely relevant and responsive documents, it makes a huge difference in both review time and cost. Using AI in this context, while somewhat expensive, is a financially accretive process for clients—it allows discovery time periods not to be consumed simply by document review.

Document review can be done more efficiently with reports showing the pace of review, the number of 'hot docs' or 'privileged' designations, and tagging by issue or witness. All of this work makes preparation for depositions and trials much more efficient and less costly.

Part VII: AI for Depositions and Trials

A. Deposition Preparation

AI tools can assist throughout the deposition process while understanding critical limitations. For outline generation, AI can draft initial deposition outlines based on pleadings, discovery responses, and case themes. Prompt the AI with key facts and issues, then refine the generated outline with your strategic insights. AI accelerates preparation, but it cannot supply the lawyer's theory of the case. Without a clear theory—what matters, what doesn't, and why—AI will simply help you prepare faster in the wrong direction.

For anticipated responses, AI can predict likely answers based on party positions and generate follow-up questions for various response scenarios. This preparation enables more dynamic examination. For cross-examination planning, AI can identify inconsistencies in prior statements, flag documents that contradict anticipated testimony, and suggest impeachment strategies.

AI tools for witness preparation include mock examination scenarios where the AI plays the examining attorney. This helps witnesses prepare for various questioning approaches while maintaining the human coaching essential to effective testimony.

B. Trial Practice Applications

Jury research and selection analytics leverage AI to analyze juror questionnaire responses, social media profiles (where permitted), and demographic patterns. Trial theme development benefits from AI's ability to test multiple narrative frameworks and identify resonant language.

Use AI to test cross-examination strategies by prompting it to argue the opposing position. Prepare closing arguments by asking AI to identify the strongest points for each side. Create compelling demonstrative evidence with AI assistance in data visualization and graphic design.

Bayesian probability frameworks (spoken in plain English) can help educate juries on complex technical evidence. As discussed in Part 1, jurors already think probabilistically; framing evidence in these terms can improve comprehension of AI decision-making and algorithmic processes.

Navigate judicial acceptance of AI tools while maintaining courtroom authenticity. The human connection essential to persuasion cannot be delegated to AI. Used well, AI improves preparation and testing; used poorly, it produces polished themes that are untethered to admissible proof and courtroom reality. Successful trial lawyers use AI for preparation and analysis while ensuring that courtroom presentation remains genuinely human.

Part VIII: Professional Responsibility and Firm Governance

A. Ethical Obligations

ABA Model Rule 1.1's technology competence requirement now encompasses AI. Lawyers must understand AI tools sufficiently to use them competently or avoid using them altogether. Model Rule 1.6 requires protecting client confidential information—never upload confidential client data to public AI platforms without proper data protection agreements. Examine your Terms of Service and Confidentiality Agreements with all AI products and make determinations as to whether the AI model is using your data to improve the LLM. Use of your data by the AI provider can be adjusted by you in the Settings tabs.

Courts are increasingly requiring AI use certifications. Some judges mandate disclosure of AI assistance in filings. Know your jurisdiction's requirements and document your process for demonstrating competent AI use. Verification of AI outputs is not optional—it is a professional responsibility requirement. As discussed in Part I, the safest approach is to treat AI output as a starting point: confirm authorities in Westlaw/Lexis, confirm quotations and pin cites, and maintain a documented review process. We will write more about the ethical issues in subsequent articles and in our Defending the Algorithm™ series.

B. Building Firm AI Governance Frameworks

Comprehensive governance structures for AI implementation include firm-wide AI policies and procedures, oversight and accountability structures, and defined roles and responsibilities for AI supervision. Client communication and transparency protocols manage expectations regarding AI use. You should discuss AI use with your clients and you may need language in your engagement agreements to cover the use of AI.

Successful AI governance frameworks address the information architecture challenge: connecting previously siloed systems before introducing AI tools so they can access the interconnected information needed for useful analysis. Documentation requirements demonstrate competent AI oversight and build defensible workflows.

Within five years, lawyers who competently use AI tools will be the norm rather than the exception.

Training programs should address use cases, ethics, attorney-client privilege implications, and prompting techniques. Regular updates incorporate emerging best practices and new tool capabilities. Governance is not a one-time exercise but an ongoing commitment to responsible AI integration.

Practice Tip: Westlaw Precision Practical Law module has a tremendous amount of AI-related material, checklists and forms for use in crafting AI governance policies and training protocols.

Part IX: Cost Considerations—Is AI Accretive or Just Efficient?

While remaining mindful of litigation spend for clients, some dynamics have changed due to AI tools. More front-end legal research can now be done to start team analysis, and more document review can be conducted—albeit at higher levels of review—depending on volume and complexity. This can be done in a very quick and cost-effective way even by a partner.

Front-end research to set parameters for more detailed research on complex commercial and IP matters is now possible at very little cost. For example, in a complex commercial case with millions of documents, modest up-front spend on analytics and targeted AI-assisted workflows can materially reduce downstream review populations and attorney hours—often lowering total discovery cost even if early-stage spend increases slightly. The research pump can be primed inexpensively, and strategic case assessments can be created more quickly. Document review in Relativity takes far less time than traditional paper-based review methods.

AI has fundamentally transformed the efficiency of legal teams through both research and writing tools and predictive analytics. While precise cost analysis varies by matter, the direction is clear: thoughtful AI implementation reduces costs while improving outcomes. The competitive advantage is due to firms that invest in proper infrastructure and training. Being able to tell a client with more certainty what the cost of massive document production will be is essential to a good client relationship.

Conclusion: The AI-Competent Lawyer

The legal profession stands at an inflection point regarding AI adoption. Within five years, lawyers who competently use AI tools will be the norm rather than the exception. The question facing today's practitioners isn't whether to adopt AI but how to do so responsibly while fulfilling professional obligations.

Hopefully, this practicum demonstrated that thoughtful AI integration—particularly with attention to information organization, prompt engineering, and verification protocols—can significantly enhance legal practice without compromising professional responsibilities. The key insight: treating AI as a powerful tool requiring both professional supervision and proper information infrastructure coordinated with IT professionals and ethics lawyers.

The specific tools matter less than the framework for thinking about AI adoption: enhance efficiency while maintaining competence, protect confidentiality, verify AI output, invest in organizational infrastructure, use predictive analytics for strategic advantage, and preserve the professional judgment that distinguishes lawyers from mere document production services.

The path forward requires neither blind adoption nor resistant skepticism, but thoughtful integration that serves clients while fulfilling our professional responsibilities. Those who invest in understanding both the technology and the frameworks governing it—who master prompt engineering, understand the interplay between generative AI and verified databases, and build proper infrastructure—will be best positioned to serve clients in this rapidly evolving landscape.

To conclude in Bayesian probability terms: P(infrastructure organization being essential | evidence of competitive advantage from AI) = Very High. Welcome to the future of defense practice—one where AI assists but doesn't replace the human judgment, ethical reasoning, and client relationship management that define our profession.