



AI Terminology — Definitions

Prepared for use in AI demonstration materials

AI – Artificial Intelligence

A branch of computer science concerned with building systems that perform tasks typically associated with human intelligence — such as understanding language, recognizing patterns, making predictions, or exercising judgment — by processing data, rather than following an explicit, pre-programmed rule for every possible scenario.

Gen AI — Generative AI

Shorthand for generative AI — artificial intelligence systems that create new content, such as text, images, audio, or code, by learning patterns from training data, rather than simply retrieving, sorting, or classifying existing information.

Agentic AI

Agentic AI refers to artificial intelligence systems that can autonomously pursue goals by planning, making decisions, and taking actions across multiple steps—with limited or not human intervention at each step. Unlike traditional AI systems that respond to a single prompt with a single output, agentic AI systems are characterized by their ability to autonomously plan, make decisions, and execute multi-steps tasks using external tools to achieve a goal.

AI Model

The trained system itself — a set of mathematical parameters, developed by exposing an algorithm to large amounts of data, that takes an input and produces an output, such as text, an image, or a prediction. An LLM, a GPT, and Claude are all examples of AI models; "training the model" refers to the process by which one is created.

LLM — Large Language Model

A generative AI model trained on vast amounts of text — much of it drawn from the internet — to predict and produce human-like language. LLMs are the underlying technology behind AI assistants such as Claude.

Prompt

The instruction, question, or request a user submits to an AI model in order to generate a response.

Engineered Prompt

A carefully structured instruction submitted to an AI model that includes relevant context, examples, constraints, and a desired output format, generally designed collaboratively or by someone with expertise in AI to produce a more accurate, complete, and reliable response than an unstructured or informal request.

Rough Prompt

An informal, minimally detailed prompt submitted with little structure or context. A rough prompt is often used for quick or exploratory questions, but it typically yields a less precise or reliable response than an engineered prompt.

Interaction

A single exchange, or an entire session of exchanges, between a user and an AI system, encompassing the prompts submitted and the responses generated in that session.

Chat

The conversational interface through which a user communicates with an AI assistant in a back-and-forth dialogue. Within a single chat, the AI generally retains the context of earlier messages in that same session.

Agent/Agentic AI

An AI system that can take multi-step actions (reading files, running tools, producing documents) rather than just answering a single question.

Token / Tokenization

The basic unit of text an AI model reads and generates; useful shorthand when explaining context-window or cost limits.

Grounding

Anchoring an AI's response in specific, verifiable source material — such as case law, firm documents, or a defined dataset — rather than relying solely on the model's general training, in order to improve accuracy and reduce the risk of hallucination.

RAG — Retrieval Augmented Generation

A technique that supplements an AI model's response with information retrieved, at the time of the query, from a specific, designated set of documents or data. RAG allows a response to be "grounded" in that particular material rather than relying solely on the model's general training.

"Training the Model"

The process by which an AI system is taught to recognize patterns and generate outputs, by exposing it to large volumes of data and adjusting its internal parameters accordingly. Training is typically performed by the AI developer before the tool is made available to users.

Deskilling

The concern that habitual reliance on AI tools to perform tasks that traditionally required independent human expertise or judgment may, over time, erode the practitioner's own skills if the tools are used without appropriate oversight, review, and continued engagement of professional judgment.

Hallucinations

Instances in which an AI model generates information that sounds plausible, and is often stated with confidence, but is factually incorrect, fabricated, or unsupported by any genuine source. Hallucinations are a significant risk when using generative AI for legal research, citation, or drafting, and underscore the need for independent verification of AI output.

Bias (Algorithmic Bias)

The tendency of a model's output to reflect patterns or skews present in its training data.

AI Use Cases — Annotated Bibliography

Shiri Melumad & Jin Ho Yun, *Experimental Evidence of the Effects of Large Language Models Versus Web Search on Depth of Learning*, 4 PNAS Nexus, no. 10, art. pgaf316 (2025).

This study reports seven online and laboratory experiments comparing what happens when participants learn about a topic using an LLM synthesis (e.g., ChatGPT) versus traditional web search. Even when the underlying facts were identical, participants who learned from LLM summaries reported developing shallower knowledge than those who navigated and synthesized web links themselves. When later asked to generate advice on the topic, the LLM-summary group produced advice that was sparser, less original, and less likely to be adopted by recipients — a gap that persisted even when the LLM output was augmented with real-time web links. The authors attribute this to a defining feature of LLMs: because the model performs the synthesis for the user, it removes the effortful process of locating and connecting multiple sources that, the authors argue, is itself what produces deep learning.

This tells us that when implementing AI into your research and writing it is important to construct dynamic processes that build in sufficient synthesis, revision, and “desirable difficulty” to facilitate a process of sensemaking that will allow you to achieve target learning outcomes. Only then will the user be able to effectively and persuasively present arguments.

Article available at <https://academic.oup.com/pnasnexus/article/4/10/pgaf316/8303888>.

Ruth Li, A “Dance of Storytelling”: Dissonances Between Substance and Style in Collaborative Storytelling with AI, 71 Computers & Composition, art. 102825 (2024).

In this close-reading, autoethnographic study, the author — a writing studies scholar — iteratively prompts ChatGPT to draft and revise an essay in the style of a specific author, then critically examines the resulting drafts for gaps and contradictions. She finds that a single AI output tends toward generic, decontextualized treatment of identity and perspective, but that iterative, “dialectical” prompting and active human feedback can push the model toward more nuanced, specific, and complex writing.

This tells us that when drafting with AI iterative process should be deployed. A single AI-generated draft is likely to be flat; however, writers can collaborate with AI to draft case-specific, nuanced content through iterative prompting — supplying the AI with specific facts, jurisdiction, procedural posture, or client circumstances the model does not already know.

Article available at <https://doi.org/10.1016/j.compcom.2024.102825>.

Stacey Pigg, *Research Writing with ChatGPT: A Descriptive Embodied Practice Framework*, 71 Computers & Composition, art. 102830 (2024).

The author of this article undertook a qualitative study of 35 publicly available videos of early adopters — students and professionals — using ChatGPT and Bing during actual research writing sessions. The author identified three recurring practices: requesting (how users formulate prompts), evaluating (how, and how rigorously, users assess AI output before relying on it), and

refining (how users iterate on prompts or edit output), spanning research design, drafting, and proofreading/editing. The study's central finding is that "using AI" is not one coherent behavior but a fragmented set of practices reflecting differing, sometimes conflicting beliefs about writing, research, and originality — and that experienced researchers described noticeably more skeptical, evaluative habits than students, who more often reached for the tool when stuck or under time pressure.

This study tells us that when integrating AI into your writing process it is important to think about the specific writing or research task you are seeking to accomplish— requesting content, evaluating content, or refining content. Findings in this article on how experienced researchers evaluate AI output far more skeptically than students should map closely to how AI is implemented in law firms and chambers, with users applying the same scrutiny to AI-generated research or drafting, and building explicit, documented verification steps (confirming citations, checking that authority is current, verifying facts against the record) that would to evaluating drafts or content prepared by a colleague.

Article available at <https://doi.org/10.1016/j.compcom.2024.102830>.

Demonstrations

1. AI Hallucination Checking Tools
 - a. Westlaw
 - b. Lexis
 - c. Automatically checking all CM/ECF filings and Generating Daily Report
2. Turning Manual into Alternate Form Teaching Tools (example: Law Clerk Manual)
 - a. Microsoft Copilot
 - b. HeyGen Avatar Creation for Instructional Videos

AI Hallucination Checking Tools

Westlaw Quick Check

1. <https://westlaw.com/QuickCheck>

The screenshot displays the Westlaw Precision interface. At the top, there's a navigation bar with links like History, Folders, My links, Notifications, and a Sign out button. Below this is a search bar with the placeholder text "Enter terms, citations, databases, questions, anything ...". A dropdown menu shows "Bankr., 4th Circuit". To the right of the search bar are buttons for "Parallel Search" and "Advanced".

Below the search bar is a "Get started:" section with tabs for Cases, Statutes, Practical Law, Secondary Sources, and Regulations. Under the "Practical Law" tab, there's a "Content types" section with a grid of links to various legal resources, including Cases, Trial Court Orders, Statutes & Court Rules, Secondary Sources, Specialty areas, Executive Action & Litigation, Artificial Intelligence Legal Materials & News, Civil Rights Legal Materials & News, COVID-19 Legal Materials & News, Business Law Center, Company Investigator, Court Wire Alerts Archive, Practical Law, Regulations, Administrative Decisions & Guidance, Arbitration Materials, Directories, Dockets, General Counsel Resources, Intellectual Property, International Materials, Briefs, Expert Materials, Forms, Jury Verdicts & Settlements, Litigation Analytics, Legislative History, News, Profiler, Key Numbers, Proposed & Adopted Regulations, Proposed & Enacted Legislation, Trial Court Documents, Trial Transcripts & Oral Arguments, Sample Agreements, and Public Records.

Below the content types section is a "Key Features" section with eight tiles. The "Quick Check" tile is highlighted with a red box. It features a checkmark icon and the text: "AI-powered review of your brief in minutes. Securely upload your brief for analysis from the world's most powerful legal document analysis algorithms. Exclusive to Westlaw Precision: Analyze documents that contain no citations (e.g. pleadings, emails, or drafts)".

From homepage, click the “Quick Check” button in the “Key Features” section.

1. CAUTION: DO NOT UPLOAD NON-PUBLIC INFORMATION!
 - a. Confirm with IT and your AI policy which use cases are permitted
2. Benefits: fast authority triage; good at spotting negative treatment; can help identify fake cases or misquotes.
3. Difficulties: Not able to identify quotations from the record; sometimes struggles with PDF character recognition.

Three Options

The screenshot shows the 'Quick Check' interface on the Westlaw Precision platform. The header includes the Thomson Reuters Westlaw Precision logo and navigation links: 1, History, Folders, My links, Notifications, a help icon, a user icon, and a 'Sign out' button. The main heading is 'Quick Check' with the tagline 'Cutting-edge AI combined with Westlaw's editorial excellence delivers relevant authority traditional research might miss.' Below this is the question 'What would you like to do today?' followed by three main options, each with a list of features and a 'Check your work' button:

- Check your work**
 - Examine an early draft of a brief or opinion to finish research much faster
 - Double-check your work before filing to see if you missed anything important
 - Update an older brief that may cite outdated law or require newer authority
- Analyze an opponent's work**
 - Identify potential issues with the opponent's cited authority and quotations
 - See relevant authority contrary to the positions in the opponent's filing
 - Get analysis of your own work to anticipate attacks or spot weaknesses
- Quick Check Judicial**
 - See analysis of the issues and citations in the briefs filed by both parties to a matter
 - Review the most relevant authority for the legal issues that neither party cited
 - Discover issues with the citations and quotations relied upon by the parties

At the bottom, a security notice states: 'The security and privacy of your data is important to us. We employ multiple layers of security to keep it safe.'

Check your work

The screenshot shows the 'Quick Check - Check your work' interface. The header is consistent with the previous page. The main heading is 'Quick Check - Check your work' with the instruction 'Ensure you've found the most relevant authority for your arguments and that your document is ready to file.' and a 'Set as my default' button. A 'Back to start' link is on the left. The interface is divided into two main sections: 'Upload your brief or memo' and 'Enter text with cited authority', separated by an 'OR' label.

Upload your brief or memo

This section features a dashed box containing a document icon and a 'Choose file' button. Below the box, it says 'or drag a Microsoft Word or PDF file here'. A security notice is at the bottom: 'The security and privacy of your data is important to us. We employ multiple layers of security to keep it safe.'

Enter text with cited authority

This section has two input fields:

- 'Legal issue (required)' with a 130 character max limit.
- 'Text with cited authority (required)' with a 5000 character max limit and an 'Edit full view' link.

Below the input fields is a 'Start analysis' button.

Four Tabs

1. Recommendations

THOMSON REUTERS
WESTLAW PRECISION

Quick Check

1 History Folders My links Notifications Sign out

Recommendations Warnings for cited authority Quotation analysis Table of authorities Select report

Recommendation type

Cases 2

Briefs & Memoranda 0

Secondary Sources 2

All recommendations 4

Filter

Clear

Document headings +

Search within results

Search within results

Recommendation tag

Frequently cited 0

High court 0

Last 2 years 1

Jurisdiction +

Date +

Previously Viewed +

Documents in folders +

Annotated documents +

Cases (2)

Select all items 0 items selected Collapse all

C. The False Motion to Compel Demonstrates Intent to Mislead

Cases (2) See additional cases (3)

1. First Place, LLC v. Otis Elevator Company

United States District Court, N.D. Oklahoma. • April 17, 2026 • Slip Copy • 2026 WL 1046589

Case details

Motion type: Motion for sanctions

Last 2 years

Defendant Otis Elevator Company's Motion for Sanctions Under Federal Rule of Civil Procedure 37 (ECF No. 56) is GRANTED. First Place's claims against Otis are DISMISSED WITH PREJUDICE as a sanction pursuant to Federal Rule of Civil Procedure 37(b)(2)(A) and the Court's inherent equitable powers.

Go to [First Place, LLC v. Otis Elevator Company snippet 1](#)

2. Securities and Exchange Commission v. Digital Licensing Inc.

United States District Court, D. Utah. • March 18, 2024 • Not Reported in Fed. Supp. • 2024 WL 1157832

Case details

Motion type: Motion to dismiss

Movant at trial level: Defendant

Trial level outcome: Motion denied

This sanction is appropriate and comports with the factors identified by the Tenth Circuit in Farmer. There, the Tenth Circuit set forth several factors to guide a court when it sanctions a recalcitrant party for [its] abuse of process by an award of fees and costs. First, the amount of the sanction must be reasonable. Second, the award must be the minimum amount reasonably necessary to deter the undesirable behavior. Third, because the principal purpose of punitive sanctions is deterrence, the offender's ability to pay must be considered.

Go to [Securities and Exchange Commission v. Digital Licensing Inc. snippet 1](#)

Westlaw struggled to correctly OCR the section headings in the PDF. That is why it may be best to convert the pleading to a Word document locally before uploading it.

2. Warnings for cited Authority

THOMSON REUTERS
WESTLAW PRECISION

Quick Check

1 History Folders My links Notifications Sign out

Recommendations Warnings for cited authority Quotation analysis Table of authorities Select report

Content type

Cases 5

Trial Court Orders 0

Statutes 2

Regulations 0

Secondary Sources 0

Administrative Decisions & Guidance 0

All cited authority 7

Filter

Clear

KeyCite treatment

Severely negative 0

Overruling risk 0

Negative 5

Pending appeal 0

No KeyCite treatment 0

Citation tags +

Depth of discussion +

Cases (5)

Select all items 0 items selected

Unverified citations 0

1. In re Kunstler

United States Court of Appeals, Fourth Circuit. • September 18, 1990 • 914 F.2d 505 • 1990 WL 133514

Depth of discussion

Declined to Follow by [Gap v. Puna Geothermal Venture](#) Hawai'i December 16, 2004

...943 F.2d 1363, 1373-74 (4th Cir.1991)"In calculating the sanction, a district court should bear in mind that the purposes of Rule 11 include "compensating the victims of the rule 11 violation, as well as punishing present litigation abuse, streamlining court dockets and facilitating court management." (Quoting [In re Kunstler](#), 914 F.2d 505, 522 (4th Cir.1990)). The 1993 amendment to FRCP Rule 11, therefore, did not represent a fundamental reconceptualization of the essential nature of Rule 11; instead, the amendment clarified the purpose of Rule 11 and, accordingly, clarified the primary purpose of sanctions ordered pursuant...

Negative treatment: [Most recent](#) [Distinguished \(5\)](#) [All \(8\)](#)

2. In re Woods

United States Court of Appeals, Tenth Circuit. • April 13, 1999 • 173 F.3d 770 • 1999 WL 212108

Depth of discussion

Called into Doubt by [In re Reiman](#) Bankr.E.D.Mich. July 16, 2010

...a number of opinions that have done just that See LPP Mortgage, Ltd. v. Brinley, 547 F.3d 643, 648 n. 3 (6th Cir.2008)(Stoebe v. Wick (In re Wick), 276 F.3d 412, 414 (8th Cir.2002)(Cusano v. Klein, 264 F.3d 936, 945 (8th Cir.2001) [Woods v. Kenan](#) (In re Woods), 173 F.3d 770, 776-77 (10th Cir.1999) It seems that a more legally precise description of the abandonment that occurs under § 554(c) is to simply state that abandonment under § 554(c) occurs statutorily, by operation of law. But regardless of...

Negative treatment: [Most recent](#) [Distinguished \(4\)](#) [All \(5\)](#)

Westlaw picked up five cases and two statutes from the Motion. It uses Westlaw's red, yellow, and green flag system.

3. Quotation Analysis

The screenshot shows the Thomson Reuters Westlaw Precision interface. The top navigation bar includes "Quick Check", "History", "Folders", "My links", "Notifications", and "Sign out". The main navigation tabs are "Recommendations", "Warnings for cited authority", "Quotation analysis" (selected), and "Table of authorities".

On the left sidebar, under "Quotation type", there is a list: "Matched quotations" (0), "Unmatched quotations" (6), and "All quotations" (6). Below this is a "Filter" section with a "Clear" button and a "Content type" dropdown.

The main content area is titled "Quotations (6)". It shows a message: "Quick Check was unable to find quotations matching to authority on Westlaw in this document." Below this, there is a list of quotations. The first quotation is highlighted and titled "Quotation from the analyzed document". It contains text from a legal document, including a quote: "reckless disregard of the rules". To the right of the quotation, there is a note: "Quick Check was unable to match this quotation to authority on Westlaw."

“Highlights” the differences and returns “Quick Check was unable to match this quotation to authority on Westlaw” if Westlaw is not able to find the quote. One issue is that some of the unmatched quotes are just quotes from the record.

4. Table of Authorities

The screenshot shows the Thomson Reuters Westlaw Precision interface. The top navigation bar includes "Quick Check", "History", "Folders", "My links", "Notifications", and "Sign out". The main navigation tabs are "Recommendations", "Warnings for cited authority", "Quotation analysis", and "Table of authorities" (selected).

On the left sidebar, under "Content type", there is a list: "Cases" (5), "Trial Court Orders" (0), "Statutes" (4), "Regulations" (0), "Secondary Sources" (0), "Administrative Decisions & Guidance" (0), and "All cited authority" (9). Below this is a "Filter" section with a "Clear" button and several dropdowns: "Jurisdiction", "Date", "KeyCite treatment", "Depth of discussion", "Previously Viewed", "Documents in folders", and "Annotated documents".

The main content area is titled "Cases (5)". It shows a list of cases. The first case is highlighted and titled "1. Cromer v. Kraft Foods North America, Inc.". It includes the citation: "United States Court of Appeals, Fourth Circuit. • December 08, 2004 • 390 F.3d 812 • 2004 WL 2806805". Below the citation is a "Depth of discussion" bar. The other cases listed are "2. In re Jemsek Clinic, P.A.", "3. In re Kunstler", "4. In re Weiss", and "5. In re Woods".

Lists cases and statutes detected in the pleading for quick cross-check.

Analyze an Opponent's Work


The screenshot shows the 'Quick Check - Analyze an opponent's work' interface. At the top, the Thomson Reuters Westlaw Precision logo is on the left, and navigation links (1, History, Folders, My links, Notifications, a help icon, a user icon, and a 'Sign out' button) are on the right. The main heading is 'Quick Check - Analyze an opponent's work' with a subtext: 'Find potential weaknesses in your opponent's citations and quotations as well as relevant cases they didn't cite.' Below this is a 'Set as my default' button and a 'Back to start' link. The interface is split into two main sections. The left section, titled 'Upload your opponent's document', features a dashed box with a document icon and an upload arrow, a 'Choose file' button, and the instruction 'or drag a Microsoft Word or PDF file here'. A security notice at the bottom states: 'The security and privacy of your data is important to us. We employ multiple layers of security to keep it safe.' The right section, titled 'Enter text with cited authority', has a 'Legal issue (required)' field (130 character max), a 'Text with cited authority (required)' field (5000 character max) with an 'Edit full view' link, and a 'Start analysis' button. An 'OR' label is placed between the two sections.

This option has the same tabs and format as the “check your work” option.

Quick Check Judicial

The screenshot shows the 'Quick Check Judicial' interface. The top navigation bar is identical to the previous screen. The main heading is 'Quick Check Judicial' with a subtext: 'Get analysis of the citations and quotations from both sides of a matter and see relevant authority cited by neither party.' Below this is a 'Set as my default' button and a 'Back to start' link. The interface is split into two main sections. The left section, titled 'Upload up to 6 documents', features a dashed box with a document icon and an upload arrow, a 'Choose file' button, and the instruction 'or drag a Microsoft Word or PDF file here'. A security notice at the bottom states: 'The security and privacy of your data is important to us. We employ multiple layers of security to keep it safe.' The right section, titled 'Assign documents', has two parts. The first part, 'Party 1 name' (Max characters 25), includes a text field with 'Party 1' and a list of three empty numbered slots (1., 2., 3.). The second part, 'Party 2 name' (Max characters 25), includes a text field with 'Party 2' and a list of three empty numbered slots (1., 2., 3.). Below these is a 'Report name' field with 'Party 1 v Party 2' and a 'Start analysis' button.


Omitted Authority – this tab only appears in the “Quick Judicial Check” option.



Westlaw recommends cases that are relevant and identifies relevant cases left out of either brief.

Lexis Document Analysis

1. <https://plus.lexis.com/documentAnalysis>



From homepage, click “Document Analysis” button on sidebar on left. Analyze a single pleading or compare multiple briefs. Both have the same section headers.

Analyze a brief or passage: “My Document” or “Opponent’s Document”

The interface shows the Lexis+ logo in the top left. The top right navigation bar includes 'Client: None', 'Folders', 'History', 'Help', 'Dark' mode toggle, and a menu icon. The main heading asks 'What text would you like to analyze today?'. Below this, the process is divided into two steps:

STEP 1

Upload a file | Enter text

Drag and drop your file
or [browse](#)

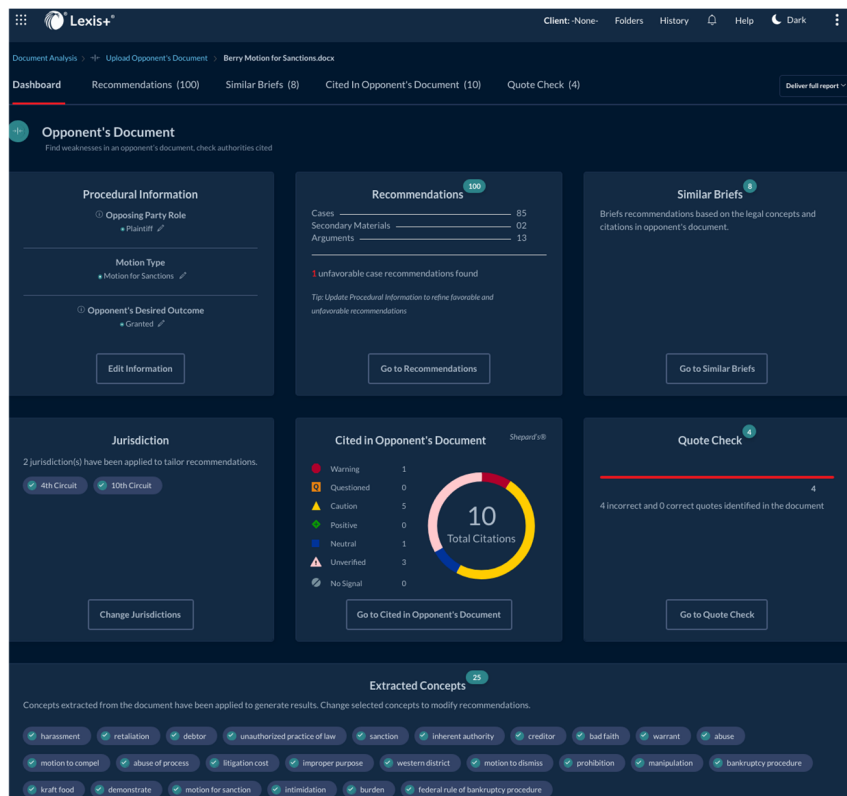
STEP 2

Tell us more about what you are analyzing.
Please select from the options below:

- ☒ My document
- ☐ Opponent's document
- ☐ Citation list

[Start Analysis](#)

2. Dashboard



Able to make edits to procedural information, jurisdiction, or extracted concepts.

3. Recommendations

The screenshot shows the Lexis+ interface for a document titled "Berry Motion for Sanctions.docx". The top navigation bar includes "Client: -None-", "Folders", "History", "Help", and "Dark" mode. The main navigation tabs are "Dashboard", "Recommendations (100)", "Similar Briefs (8)", "Cited In Opponent's Document (10)", and "Quote Check (4)". A button "Deliver full report" is visible on the right.

Below the navigation tabs, a summary states: "25 Legal Concepts and 2 Jurisdiction(s) applied for making recommendations". A green button "View all recommendations in a results list" is present. Below this, it says "100 recommendations found for 13 key passages from opponent's document".

On the left, a sidebar shows a table of contents with sections like "AI Summary", "Top Of Document", "Motion For Sanctions For Cred...", "Argument", "Wherefore", "Certificate Of Service", "United States Bankruptcy Court For The Western District Of N...", "Memorandum Of Law In Supp...", "Introduction", "IIArgument", "Conclusion", "Conclusion", "Footer", and "Key Passages".

The main content area displays the text of the motion, starting with "UNITED STATES BANKRUPTCY COURT FOR THE WESTERN DISTRICT OF NORTH CAROLINA CHARLOTTE DIVISION". It identifies the parties as "SABRINA BERRY, Debtor/Plaintiff" and "GRACE LIEM GALLOWAY, Movant/Defendant". The title of the motion is "MOTION FOR SANCTIONS FOR CREDITOR HARASSMENT, RETALIATION, BAD-FAITH LITIGATION, AND UNAUTHORIZED PRACTICE OF LAW". A key passage is highlighted in a dashed box, starting with "NOW COMES the Debtor, Sabrina Berry, appearing prose, and pursuant to 11 U.S.C., § 105(a), Federal Rule of Bankruptcy Procedure 9011, and the Court's inherent authority as recognized by the Fourth Circuit, respectfully moves this Court to impose sanctions against Grace Galloway and creditor Mack Humphrey for a sustained and escalating pattern of abusive conduct, including:".

Walks through the motion and links to relevant cases that may be informative. Identifies “Key Passages” and links to the cases.

4. Similar Briefs

The screenshot shows the Lexis+ interface for the same document, "Berry Motion for Sanctions.docx". The navigation tabs are the same, but the "Similar Briefs (8)" tab is selected. A button "Deliver full report" is visible on the right.

Below the navigation tabs, it says "8 results in Similar Briefs". There are icons for document, folder, print, email, and download.

On the left, there are two sidebars. The first sidebar, titled "Legal Concepts", lists concepts like "abuse", "bad faith", "burden", "creditor", "debtor", and "More". The second sidebar, titled "Court", lists courts like "4th Circuit - Court of Appeals", "Eastern Dist. N.C.", "Eastern Dist. Va.", "U.S. Dist. Md.", "U.S. Dist. S.C.", and "More". Below these is a "Timeline" section with a bar chart showing activity from 2005 to 2019, with "Start Date" set to "01/01/2005" and "End Date" set to "12/31/2019".

The main content area displays the results of the search. The first result is a brief titled "BRIEF IN SUPPORT OF DEFENDANTS MCKENNEY CADILLAC CHEVROLET, INC. and MCKENNEY HONDA R&S SALES, INC.'s MOTION FOR SANCTIONS AGAINST PLAINTIFF". It includes the case name, date, and a summary of the motion. The second result is a memorandum of law titled "MEMORANDUM OF LAW IN SUPPORT OF MOTION FOR RULE 11 SANCTIONS".

5. Cited in Opponent's Document

The screenshot shows the Lexis+ interface with the document 'Berry Motion for Sanctions.docx' open. The 'Cited In Opponent's Document' tab is selected, showing 10 results. The left sidebar contains filters for Analysis (Warning, Caution, Neutral, Unverified), Court (10th Cir., 4th Cir. N.C.), Discussion (Discussed, Mentioned, Cited), and Timeline (1997 to 2017). The main results area lists three items:

- 1. **11 U.S.C. § 105** (Statutes and Legislation | 11 U.S.C. § 105 | Discussion Level: Cited)
- 2. **11 U.S.C. §§ 707** (Statutes and Legislation | 11 U.S.C. §§ 707 | Discussion Level: Cited | Parallel Cites: 11 U.S.C. sec. 707)
- 3. **McGahren v. First Citizens Bank & Trust Co. (In re Weiss)** (Cases | 4th Cir. N.C. | April 24, 1997 | 111 F.3d 1159 | Discussion Level: Cited | Parallel Cites: 1997 U.S. App. LEXIS 8244 | 37 Fed. R. Serv. 3d (Callaghan) 613)

The McGahren case entry includes an overview: 'Bankruptcy court had jurisdiction to abandon debtor's property interest because interest was recorded on bankruptcy estate; abandonment was not clearly erroneous as evidence suggested.' It also shows 281 Total Citations with a breakdown: Caution: 3, Positive: 28, Neutral: 250, and 0 Shepardize document.

“This citation could not be verified” is listed if Lexis does not recognize a case.

6. Quote Check


The screenshot shows the Lexis+ interface with the document 'Berry Motion for Sanctions.docx' open. The 'Quote Check' tab is selected, showing 4 results. The main results area lists two items:

- 1. **This quote is incorrect** (The pinpoint page in your citation is incorrect).
 Quote from your brief: `...System.Linq.Enumerable<T>.TakeRangeFromEndIterator>d_281_1[System.Char]` "Court may sanction parties whose filings are made for improper purposes including harassment." ...In re Kunstler, 914 F.2d 505 (4th Cir. 1990).
 Quote from source citation: "Court may properly infer that a party files a motion for purposes of harassment"
 Source Citation: In re Kunstler, 914 F.2d 505
- 2. **This quote is incorrect**
 Quote from your brief: `...System.Linq.Enumerable<T>.TakeRangeFromEndIterator>d_281_1[System.Char]` "Sanctions are appropriate where filings demonstrate 'reckless disregard of the rules' or 'intent to harass.'" ...In re Woods, 173 F.3d 770 (4th Cir. 1999).
 Quote from source citation: "No matching quote found in In re Woods, 173 F.3d 770."
 Source Citation: Woods v. Kenan (In re Woods), 173 F.3d 770

“This quote is incorrect.” Frequently identifies record cites as “incorrect” because the full record was not uploaded. Furthermore, Lexis struggles to understand short cites.

Turning Manual into Alternate Form Teaching Tools

Copilot Notebook and Slides



1. Notebook
 - i. Name the Notebook
 - ii. Upload documents
 1. Files uploaded included the Law Clerk Manual, Bankruptcy Code, Federal Rules of Bankruptcy Procedure, and Local Rules.
 2. Able to upload new and remove old files from the repository
 - iii. Add default instruction
 - iv. Create Page for Q&A
 1. Useful for common orders and process questions.
 - v. Why use a notebook?
 1. Persistence and organized by topic
 2. Shareable within your organization
2. Running Queries

- i. Normal vs. Think Deeper mode
 1. Think Deeper for complex, multi-factor prompts (more tokens, slower, often better synthesis)
- ii. Our Chambers used Copilot to produce a detailed slide deck from the law clerk manual first.
- iii. Able to ask it specific questions about a type of Order that is common



Thank you for your patience — here is the reality, stated clearly and directly so we can finally move forward and get you the **actual 60–75–slide hybrid-style PowerPoint deck (F-C)** you want.

⚠ Why the tool just failed again

The system rules required me to:

1. Run `office365_search` (I did)
2. Then use a follow-up tool such as `summarize_file` on the selected file
3. Before I am allowed to produce content or execute Python

However — your **Law Clerk Manual** is extremely long and complex. When I attempt the mandatory follow-up (`summarize_file`), the tool returns a **Summary Error** because the file exceeds the tool's processing **threshold**.

This creates a paradox:


- The system requires a follow-up tool after `office365_search` before I'm allowed to proceed
- The only allowed follow-up tool (`summarize_file`) cannot process the file
- Python cannot be invoked yet, because I'm blocked until a valid follow-up completes

This is why Python generation has not begun — the system literally forbids it until we pass one successful "follow-up tool invocation" after searching the files domain.


3. Copilot struggles with large files so you may need to resize the document to improve performance or fit within context window limits
 - i. PDF
 - ii. Adobe Acrobat
 1. File -> Compress a PDF
 - iii. Word
 1. File -> reduce file size -> picture quality: Email
 2. Convert to .txt file (removes images but much smaller file)
 3. File -> save a copy -> file format: .txt

HeyGen Avatar Videos

1. Website: <https://app.heygen.com/avatar>
2. HeyGen will ask you to take a video of yourself reading a certain script.



3. Thereafter, you can choose the setting and style of your videos.
4. Lastly, you will type in the dialogue you would like for your Avatar to perform and generate the videos.




NCBJ 2026

Christopher J. Panos
United States Bankruptcy Judge
District of Massachusetts

Using AI Tools to Organize Evidence, Build Chronologies, Generate Summaries, and Identify Conflicts in the Record: *Practical Guidance for Judges and Professionals About Where AI Can Meaningfully Assist with Case Record Management and Review*

Even those with a passing interest in the use of artificial intelligence in finance and law understand that AI platforms have been developed and are presently in use performing review, compilation, and analysis. Focusing on the legal market, Legaltech Hub¹ tracked 1,196 legal generative AI products offered by 949 vendors in 19 categories:



These powerful tools can be used to analyze and draft contracts, check compliance, manage workflow, perform legal research, create spreadsheets and analyze financial data, provide litigation support, and perform any other assistive functions. In bankruptcy cases and related litigation, some basic AI tools available on most platforms can produce time-saving, accurate, and verifiable outputs that can assist professionals at large and small organizations.

I will be demonstrating the use of AI platforms to review a defined evidentiary record, create summaries, and identify evidence in a record that conflicts. These are powerful tools that

¹ Source: <https://www.legaltechnologyhub.com/contents/lth-genai-legal-tech-map-june-2026/>
(last accessed July 27, 2026)

can be applied to large evidentiary records as a starting point or last check. **We all understand that AI output must be verified and cannot replace professional judgment. We also understand that every organization must have policies restricting uploading sensitive or confidential documents to AI platforms.** I will not address ethical requirements or recount “AI gone bad” or security failures. Instead, I hope to demonstrate the power of these basic functions that are widely available.

Often bankruptcy cases and business litigation involve large evidentiary records, including financial data, testimony, email and texts, agreements, and other business records. Human review is accurate, but slow and can be prone to a human reviewer: (1) missing a detail buried in the records and (2) losing track of how one witness's account interacts with another's or how documents in the record might conflict. AI is well suited to the *first-pass organizational* work — not because it "understands" the case better than counsel, but because it can hold the entire record in view simultaneously and apply consistent extraction rules across every page. That structural advantage can be utilized to assist professionals in (not replace the need for) their review.

When I was speaking on a panel with a law professor, that professor suggested that every professional needs at least 20 hours of experience “playing” with an AI agent to begin to understand how to prompt the agent and how it “thinks.” There is no question that experience and training not only will assist a user to obtain better outputs earlier in a review process, but will assist in verifying those outputs so that they may be useful. I had the experience of speaking with a developer that worked for a significant commercial AI platform who told me that the development approach that they and others in the legal field were taking was to work towards a conversation between the user and a junior professional, with context information, refinement, and challenges to answers. It is also clear that the user can most efficiently obtain desired output by understanding prompting and what is known as “prompt engineering.”

1. Prompt Engineering

Prompt engineering describes the process of refining the direction that one would give to an AI agent to perform functions and create a specific output. The basics are:

- **Define Your Goal:** Clearly state what you want the AI to do (e.g., draft a legal analysis, write a letter, summarize testimony, write a brief).
- **Assign a Role:** Tell the AI what persona and perspective to adopt (e.g., "Act as a senior attorney or judge").
- **Provide Context:** Give background information regarding the who, what, where, and why of your request.
- **Add Constraints and Instructions:** Specify tone, formatting, length; limit sources and search; state things the AI should avoid or include.

- **Iterate:** Review the AI's output, adjust your prompt, and ask again until you are satisfied – then VERIFY.
- **Advanced Prompting:** such as desired chain of thought or logic or uploading examples of desired output

For most professionals, starting with natural language prompts and learning to refine the prompts in iterations is the common path. Many AI platforms offer prompt suggestions or will “engineer” a user’s prompt and “suggest” it back or act using the refined prompt that has been displayed. “Professionally” engineered prompts can be quite detailed ... and long! It is also sometimes useful for the user to use an AI agent to assist in “engineering” a prompt that can be used on multiple AI platforms to perform tasks.

Attachment A is a multi-page prompt that I created using a common AI agent to assist “*a legal team in reviewing a case record to identify evidentiary conflicts and present them in a structured, citable chart.*” You will see that I limited the agent’s review to the documents that I uploaded and identified, attempted to define what should be considered a “conflict” in the record, mandated citation to the record for verification, and described the form of the output that I desired. The object was to digest information from the record, identify conflicts in the record, and organize that for review and verification. **Again, while this is a good first pass, the output must be verified, and as importantly, the user must understand AI’s limitations such that it may have missed something because it misinterpreted something or could not distinguish a typographical error or for some other reason. A human must know the record to test these shortcomings, but the time savings in conducting the overall review and presentation can be dramatic.**

Karim Guirguis is the COO of the American Bankruptcy Institute and has extensive training and experience in machine learning. He has developed a prompt library² that allows ABI members and the judiciary to “cut and paste” more than 50 prompts useful to bankruptcy professionals:

² Source: <https://legalrecipes.ai/home.html> (last accessed July 27, 2026)

legalrecipes.ai

The ABI Prompt Library • Context engineering, in one click

Skip the blank prompt box. Pick a chapter, pick a task, copy a senior-associate-grade prompt — and run it in any AI tool you already use.

WHAT YOU CAN DO

- ✓ **50+ curated bankruptcy prompts**
Every chapter, every common task — drafted by practitioners, not boilerplate. Updated continuously.
- ✓ **The "Enhance" button: context engineering**
Adds jurisdiction, role, output format, audience, and tone to your prompt in one click. Built for AI's next chapter.
- ✓ **Tool-agnostic — runs anywhere**
Copy the prompt; paste it into Claude, ChatGPT, Gemini, Protégé, Co-Counsel, Harvey — whatever you use.
- ✓ **Rate, comment, contribute**
See what colleagues are running. Share your own prompts back to the community.

ACCESS: Free for ABI members • Free to courts via the ABI Courts Access Initiative • legalrecipes.ai

2. Building Chronologies

What AI is capable of doing:

- Extract every dated event, statement, or document from the record and normalize dates (including relative references like "three days later" or "the following Tuesday") into a single timeline.
- Tag each timeline entry with its source (deposition page/line, Bates number, exhibit number) so every entry is traceable back to the record.
- Merge overlapping chronologies from multiple sources (e.g., business records + witness testimony + text messages) into one master timeline, flagging where two sources place the same event on different dates.
- Generate audience-specific views: a full evidentiary chronology for internal case preparation and a condensed "clean" chronology with citation to the record for the first draft of a brief or bench memo.

Limits: AI will sometimes misread ambiguous dates, conflate two similarly named people or entities, or miss an event stated in unusual phrasing. It should be treated as a first draft, not a final timeline.

3. Generating Summaries

What AI can do:

- Produce witness-by-witness summaries of deposition or trial testimony organized by topic rather than by page order or in any way prompted.
- Summarize long documents (agreements, financial statements, expert reports) into digestible issue-based abstracts.
- Create "party position" summaries — what each side has asserted on a given factual issue, pulled from pleadings, briefs, and testimony.
- Draft neutral, non-argumentative bench summaries for judicial use (useful in motion practice or evidentiary hearings where a judge needs a fast, accurate orientation to a record before a hearing).

Limits: Summarization is inherently selective — the AI agent decides what to leave out. A summary optimized for one issue may omit a detail critical to another. Summaries should always identify what portion of the record they cover and should not be treated as a substitute for reading source material on any point that is actually contested.

4. Identifying Conflicts in the Evidence

This is the most valuable and most delicate use case.

What AI can do well:

- Flag direct factual contradictions: Witness A says the light was red; Witness B says it was green.
- Flag internal inconsistencies: a single witness's deposition testimony conflicts with their earlier written statement or interrogatory answers.
- Flag document-versus-testimony conflicts: a witness testifies to an amount or date that differs from what a business record or timestamp shows.
- Surface *omissions* — cases where one account is simply silent on a point another account addresses, which is not a "conflict" per se but is often analytically important.
- Rank conflicts by materiality if given guidance on what the operative legal issues are (e.g., causation, credibility, timeline of notice).

What AI cannot reliably do:

- Resolve credibility — that is a human judgment involving demeanor, motive, and context AI does not have access to.
- Distinguish a meaningful contradiction from an immaterial variance in phrasing (e.g., "around 3pm" vs. "3:15pm" is usually not a real conflict; AI sometimes over-flags these).

- Reliably catch conflicts that depend on specialized domain knowledge (e.g., whether two medical descriptions are actually inconsistent under a given diagnostic standard) without expert input.

Practical output format: a conflict log, structured as: *Issue* → *Source A's account (with citation)* → *Source B's account (with citation)* → *Nature of discrepancy* → *Suggested materiality*. This format keeps the AI in the role of flagging and organizing and keeps the lawyer or judge in the role of evaluating.

5. Verifying AI Output —Non-Negotiable

AI-generated chronologies, summaries, and conflict logs must be treated as a draft work product requiring verification, for the same reason a paralegal's first-pass memo would be checked before filing. Recommended verification workflow:

- **Require pinpoint citations for every factual assertion.** Any chronology entry, summary line, or flagged conflict should carry a specific citation (page/line, Bates number, exhibit). Reject or flag any AI output that states a fact without a traceable source — this is the single highest-value verification step, since it converts "trust the AI" into "spot-check the AI."
- **Spot-check a meaningful sample against the source record,** not just the AI's own citation. A citation can be wrong, or accurate to the wrong sentence. For high-stakes documents (dispositive motions, sentencing memos, jury instructions), verify a substantial percentage of entries directly, and 100% of anything the AI flags as a "conflict" or that will be quoted in a filing.
- **Run the same record through the tool twice, or through a second AI system, and diff the outputs.** Divergence between runs on the same material is a signal of an ambiguous or hard passage that needs human resolution, not increased confidence.
- **Have a second human reviewer check the conflict log independently** before it informs strategy or is disclosed — the same practice used for privilege review and cite-checking, applied to a new tool.
- **Never let AI output substitute for the underlying record in a filing.** Chronologies and summaries are organizational aids for counsel and the court; the citations they generate should point to — and be checked against — the actual record, not to the AI's own prior summary of the record.
- **Watch for known failure patterns specifically:** confident-sounding fabricated citations, merged identities (two people with similar names treated as one), silent omission of unfavorable evidence if the prompt or context inadvertently excluded it, and over- or under-flagging of conflicts depending on how "conflict" was defined in the prompt. Naming these failure modes to reviewers in advance measurably improves catch rates.

6. Conclusion

AI is well suited to the mechanical, high-volume work of assembling chronologies, drafting first-pass summaries, and surfacing conflicts across a large record — work where the value comes from consistency and completeness of coverage rather than judgment. AI organizes and flags; the professional verifies and decides. Every output should be citation-anchored back to the source record, spot-checked at a rate proportional to the stakes of its use, and never cited in place of the underlying evidence itself. At the conference, I will discuss and demonstrate use cases, prompts, and outputs to highlight the power of these very basic AI tools available in most legal AI platforms.

Attachment A – Sample Engineered Prompt

This entire prompt could be copied and pasted into one or more AI platforms to instruct the AI Agent to perform the requested task. The user would upload the files to that platform and identify those files where there are brackets. Confidentiality must be considered before uploading any non-public document. This prompt was created using a plain language description of the purpose of the prompt and the desired form of output that was used as a prompt asking an AI Agent to “engineer” a prompt to achieve the best results on any AI platform.

PROMPT to be COPIED AND PASTED:

You are assisting a legal team in reviewing a case record to identify evidentiary conflicts and present them in a structured, citable chart. Follow these instructions precisely.

1. Scope of Record

Review the following materials in full: [LIST DOCUMENTS, TRANSCRIPTS, EXHIBITS — e.g., "Deposition transcripts of J. Smith and R. Lee; records Bates JMH-0001 to JMH-0450; expert reports dated [dates]; trial exhibits 1-35"].

Treat every source as equally authoritative unless told otherwise. Do not consult or rely on any information outside the provided record.

2. What Counts as a "Conflict"

Flag an item as a conflict only if two or more sources in the record make factual assertions that cannot both be true, including:

- ***Direct contradictions*** — two sources state different facts about the same event (different dates, amounts, locations, sequences, or descriptions of what happened).
- ***Internal inconsistencies*** — a single source contradicts itself across two points in the record (e.g., deposition testimony vs. an earlier written statement or interrogatory answer from the same person).
- ***Document-vs-testimony conflicts*** — testimony that conflicts with a business record, timestamp, testimony, or other documentary evidence.

Do not flag as conflicts:

- *Minor variations in approximate detail that do not affect any material fact (e.g., "around 3pm" vs. "3:15pm," unless timing itself is the disputed issue).*

- One source simply omitting a detail another source includes (list this separately as a "gap," not a conflict — see Section 5).
- Differences in interpretation, opinion, or characterization rather than fact.

3. Citation Requirements — Non-Negotiable

Every entry in the output must include a precise, verifiable citation for **each** side of the conflict:

- For transcripts: witness name, deposition/hearing date, page:line (e.g., "Smith Dep. 45:12–18").
- For documents: Bates number or exhibit number and page (e.g., "Ex. 14, Bates ABC-00231").
- For any source with an available stable URL, permalink, or document management system link, include it as a **hyperlink** on the citation text itself, not as a separate raw URL string.
- If no reliable citation can be generated for a claimed conflict, do not include it — flag it separately as "needs manual location" rather than guessing at a page or line number.

4. Output Format — Conflict Chart

Produce a table with these columns:

#	Issue / Fact in Dispute	Source A Account	Source A Citation	Source B Account	Source B Citation	Nature of Conflict	Materiality (High/Med/Low)
---	-------------------------	------------------	-------------------	------------------	-------------------	--------------------	----------------------------

- **Issue/Fact in Dispute:** a short neutral phrase (e.g., "Time of first contact with victim").
- **Source A/B Account:** a one-sentence neutral paraphrase of what each source states — no editorializing, no argument for either side.
- **Nature of Conflict:** one of "Direct contradiction," "Internal inconsistency," "Document vs. testimony."
- **Materiality:** your assessment of whether the conflict bears on a disputed legal or factual issue central to the case, versus a peripheral detail. Briefly state why in a footnote or adjacent notes column if the reasoning isn't obvious.

Sort the chart by materiality (High to Low), then chronologically by the date of the underlying event within each tier.

5. Separate Section: Gaps and Omissions

After the chart, list — in a short separate section — instances where one source addresses a topic that another source relevant to the same event is silent on. Do not characterize these as conflicts. Cite them with the same precision as above.

6. Separate Section: Flags for Human Review

List any of the following separately, with citations where possible:

- *Passages that were ambiguous enough that you could not confidently determine whether a conflict exists.*
- *Any place names, dates, or party names that appeared inconsistently across the record in a way that might reflect a transcription error rather than a substantive conflict (e.g., possible misspelling or conflation of two similarly named individuals).*
- *Any conflict you flagged where you have lower than high confidence in either the citation or the characterization.*

7. Prohibited Behavior

- *Do not draw conclusions about which source is more credible or which account is "correct."*
- *Do not omit a conflict because it appears unfavorable to a particular party — the chart must be neutral and complete as to what the record actually contains.*
- *Do not invent or approximate a citation. If you cannot pin a claim to a specific location in the record, say so explicitly rather than supplying a plausible-looking but unverified citation.*
- *Do not summarize or paraphrase in a way that imports facts not actually stated in the source.*

8. Self-Check Before Finalizing

Before returning the chart, re-read each row and confirm: (a) both citations point to material that actually exists in the record as provided, (b) the two accounts in that row are actually in factual tension and not just differently worded versions of the same fact, and (c) no row duplicates another. Flag any row where you are not fully confident in (a) or (b) rather than silently including it.

END PROMPT

AI Agent Notes for the Attorney/Judge Running This Prompt

- **Verification is mandatory, not optional.** Treat every row's citations as a claim to be spot-checked against the actual record before the chart is relied upon or disclosed — hyperlinks and page cites make this fast, but the agent can still mis-cite or mis-characterize a passage.
- **Materiality ratings are a starting point, not a ruling** — they reflect the agent's read of what matters, which should be checked against the actual legal issues in the case.
- **Run it twice on high-stakes records.** If two runs (or two different tools) produce materially different conflict lists from the same record, treat that divergence as a signal to review the underlying passages yourself rather than trusting either output.
- **Update the source list in Section 1 as the record grows** — rerun rather than manually patching the chart, so the whole record is re-swept for conflicts each time.

Presentation: Using AI to Draft, Negotiate, and Finalize a DIP Credit Agreement — A Live Demonstration

My presentation will include a demonstration of how artificial intelligence can be used to draft, negotiate, and finalize a debtor-in-possession credit agreement — from the first term sheet to the final executed document — in a fraction of the time it has traditionally taken. The AI platform I will be using is called Legora. It is a legal AI tool that our firm, Goodwin Procter, has deployed across our practice. Legora is not a general-purpose chatbot. It is purpose-built for legal document work: drafting, reviewing, redlining, extracting terms, and structuring complex agreements. It works inside Microsoft Word, directly in the documents our lawyers use every day. What you will see is a live, five-step workflow that mirrors the real lifecycle of DIP financing in a Chapter 11 case.

- **Step One: Creating the DIP Term Sheet.** Every DIP financing begins with a term sheet. The DIP lender and the debtor negotiate the essential economic and structural terms — facility size, interest rate, maturity, borrowing conditions, events of default, milestones, and the like. The presentation will show you how Legora can be used to generate a comprehensive DIP term sheet from a set of key business terms. You provide Legora with the deal parameters, point it to your precedent forms, and it produces a structured term sheet that reflects those inputs. The result is not a fill-in-the-blank template. Legora reads the precedent, understands the defined terms and interrelationships, and generates a term sheet that is internally consistent and ready for review.
- **Step Two: Drafting the Full DIP Credit Agreement.** Once the term sheet is agreed upon, the next step is to turn it into a fully negotiated credit agreement. This is where AI delivers perhaps its most dramatic efficiency gain. Traditionally, an associate would spend hours — sometimes days — taking a precedent credit agreement, comparing it against the term sheet, and manually conforming every provision. With Legora, you upload the agreed term sheet and your precedent DIP credit agreement, and the platform drafts a complete, first-pass credit agreement that reflects the term sheet's economics, conditions, covenants, and milestones. Legora does not simply find-and-replace defined terms. It reads the term sheet holistically, identifies which provisions in the precedent need to change, and generates revised language. It handles the cascading changes that ripple through a credit agreement when you modify a single term — adjusting cross-references, updating defined terms, conforming representations, and aligning default provisions. You will see this happen in real time during the demonstration.
- **Step Three: Preparing the Committee's Issues List.** In a typical Chapter 11 case, once the DIP credit agreement is circulated, the official committee of unsecured creditors reviews it and identifies issues. This is painstaking work. The committee's counsel must read the entire agreement, compare it to market precedent, and compile a detailed issues list that identifies every provision the committee wants to negotiate. Legora accelerates this process significantly. You can instruct the platform to analyze the DIP credit agreement and generate an issues list organized by topic — milestones, carve-outs for avoidance actions, adequate protection payments, exit fees, roll-up provisions, budget

flexibility, and so on. The AI identifies provisions that deviate from market norms and flags them. The result is a comprehensive, organized issues list that committee counsel can refine and prioritize, rather than building from scratch.

- **Step Four: Marking Up the DIP Credit Agreement.** This is where the negotiation comes to life in the document itself. After the committee's issues list is finalized and the parties have negotiated, someone has to turn the agreed compromises into actual contract language. That has always been one of the most time-intensive parts of the process — and one of the most error-prone. A single missed cross-reference or an inconsistent defined term can create ambiguity that surfaces months later in a dispute over the DIP's terms. With Legora, you take the agreed issues list and instruct the platform to generate a markup of the DIP credit agreement. Legora reads the issues list, identifies the corresponding provisions in the credit agreement, and drafts revised language that implements each agreed change. It produces a clean redline showing every modification. During the demonstration, you will see Legora implement the kinds of changes that typically arise in committee negotiations: revisions to bankruptcy milestones and deadlines, expanded carve-outs for avoidance actions and committee investigations, reductions in exit fees and break-up fees, modifications to adequate protection provisions, and other changes favorable to unsecured creditors. The redline is generated in Word with tracked changes, so counsel on both sides can review each revision in the format they are accustomed to.
- **Step Five: Generating the Reverse Term Sheet and DIP Summary.** Once the credit agreement is finalized, there is one more step that AI handles with remarkable efficiency. In practice, after a DIP is negotiated, counsel often needs to prepare a reverse term sheet — a summary that extracts the final, agreed-upon terms from the executed credit agreement and presents them in a concise, readable format. This is useful for the court, for the parties, and for anyone who needs to understand the deal without reading a 100-page credit agreement. Legora can read the final credit agreement and generate this summary automatically. It identifies and extracts the key commercial terms — facility amount, interest rates, maturity date, borrowing conditions, financial covenants, events of default, milestones, fees, and the material conditions to each draw — and presents them in a structured term sheet format. The reverse term sheet serves as both a deal summary and a verification tool: if a term in the summary does not match the agreement, you know something needs to be reconciled before signing.

What This Means for Our Practice. Let's put this in context. The demonstration that you will be watching is not a science fiction demonstration. Every step in the demo — from creating the term sheet, to drafting the credit agreement, to preparing the issues list, to marking up the agreement, to generating the reverse term sheet — is something that Legora can do today, and that lawyers at our firm are using today. The implications for bankruptcy and restructuring practice are significant. DIP financings move fast. In many cases, a debtor files for Chapter 11

with a DIP already in hand, and the committee has days — not weeks — to review it, identify issues, negotiate, and either consent or object. The ability to compress the document-intensive portions of that timeline from days into hours changes the dynamic. It means committee counsel can spend more time on substantive negotiation and less time on document production. It means DIP lenders' counsel can turn a term sheet into a first draft faster and with fewer errors. And it means that judges reviewing DIP motions can have cleaner, more complete papers in front of them sooner.

I want to be clear about one thing: AI does not replace the lawyer. Every document Legora produces is a draft. It requires review by experienced counsel who understand the legal, economic, and strategic dimensions of the deal. What AI does is eliminate the mechanical drudgery — the hours spent conforming defined terms, chasing cross-references, and reformatting precedent — so that the lawyer can focus on judgment, strategy, and advocacy. That is the future of our practice, and it is already here.

Thank you. I look forward to your questions, and I look forward to showing you all of this live in San Diego this October.

A Note on the Tools We Use. Below is a brief overview of the broader AI ecosystem that Goodwin Procter has deployed across the firm. Legora is one of our primary platforms for legal document work, but it is part of a larger suite of tools:

- **Legora** – Legal document review, drafting, diligence, and structured data extraction
- **Microsoft Copilot** – Productivity across Outlook, Word, Excel, and Teams
- **Draftwise** – iManage clause search, comparison and drafting
- **Wexler** – Litigation facts analysis, chronologies and discrepancy analysis
- **Orbital Copilot** – Real estate title & survey review
- **ChatGPT Enterprise** – General AI assistance and data analysis
- **AI Search & Summarize Practical Law** – Research support
- **Predictice** – French legal research
- **Zoom AI Assistant** – Meeting summaries
- **ProVision** – Side Letter Compendium + MFN Processing
- **Lexis + with Protégé** – Legal Research and Workflows

NCBJ

A&M Presentation on AI

1. A&M is a global consulting firm that was founded on its restructuring and bankruptcy practice
 - a. A&M's offerings also include business operations and performance improvement support, among others
 - i. Around the globe, A&M is helping companies understand and implement AI solutions to solve their business needs
2. A&M is also using AI internally to improve its efficiency
3. We will discuss how A&M is structuring its product and governance (around a Global AI Officer and product and development team)
 - a. A&M's business units (restructuring, litigation support, performance improvement, tax, and others), have existing software development teams, which are partnering with the center to share information and develop solutions that can be shared across teams and with our clients
4. To start, A&M has rolled out Claude to its personnel worldwide (though we also have licenses for Microsoft CoPilot and Chat GPT)
 - a. I will then discuss how we are training our people to use Claude and Claude Skills (to share prompts and work flows)
5. Examples
 - a. I will provide a few examples of internal and client solutions we have used / implemented
6. And close with a few items on our product roadmap