The Kentucky Bar Association Alternative Dispute Resolution Section presents:

AI in Dispute Resolution: Navigating Uncharted Waters

This program has been approved in Kentucky for 1.5 CLE credits.

Compiled and Edited by:
The Kentucky Bar Association
Office of Continuing Legal Education
for
Kentucky Bar Association
Alternative Dispute Resolution Section

Published by:
The Kentucky Bar Association, October 2025.

Editor's Note: These Kentucky Bar Association CLE materials are intended to provide current and accurate information about the subject matter covered as of the original publication date. The program materials were compiled for you by volunteer authors and from national legal publications. No representation or warranty is made concerning the application of legal or other principles discussed by the instructors to any specific fact situation, nor is any prediction made concerning how a particular judge or jury will interpret or apply such principles. The proper interpretation or application of the principles discussed is a matter for the considered judgment of the individual legal practitioner. The faculty and staff of this Kentucky Bar Association CLE program disclaim liability therefor. Attorneys using these materials, or information otherwise conveyed during this program, in dealing with a specific legal matter have a duty to research the original and current sources of authority. In addition, opinions expressed by the authors and program presenters in these materials do not reflect the opinions of the Kentucky Bar Association, its Board of Governors, Sections, Committees, or members.

AI in Dispute Resolution: Navigating Uncharted Waters

Table of Contents

PRESENTER BIOGRAPHIES

Larry Bridgesmith, J.D., is the Head of Legal Innovation at Guardrail Technologies, Inc. He founded the first Tennessee registered legal DAO built for the association of legal reformers, DAO dispute resolution, and legal wrapper consultation. Larry is a member of the International Institute of Legal Project Management (IILPM). He was awarded the Legal Innovation award by the President of the Tennessee Bar Association and received the Tennessee Supreme Court Chief Justice Janice Holder award for legal innovation. Larry's work in legal technology innovation was recognized as a finalist in the Legal Innovator of the Year in 2023 in the Southeastern United States. Larry has authored numerous published articles on legal innovation and along with his co-authors, Dr. Adel Elmessiry and Hon. Hummaded S. Morei, was awarded best paper recognition at the 2022 International Conference on BlockChain which addresses the use of DAOs in the legal industry. He founded Lipscomb University's Institute for Conflict Management. Larry teaches at Vanderbilt University School of Law (where he co-founded the Program on Law & Innovation) and Arizona State University Sandra Day O'Connor College of Law (where he was recognized as Teacher of the Year in 2022).

With over 50 years of experience, **Robert Bergman** brings deep expertise in software engineering, decision sciences, simulation, defense systems, process control, telecommunications, market research, and strategic planning. For 28 of those years, his focus at Intel Corporation was on mobile and cellular communications planning, including international standards development for radio services, cable modems, and satellite communications. The previous 15 years at Gould Electronics were spent designing flight simulators for defense aircraft. Over the past seven years, Mr. Bergman has led the development of NextLevel Mediation, applying decision analytics, artificial intelligence, and game theory to transform dispute resolution. He received degrees from MIT and SUNY Albany.

AI: WHAT IS THE HISTORY AND HOW DOES IT WORK?

Robert Bergman* and Larry Bridgesmith, J.D.**

I. A BRIEF HISTORY OF ARTIFICIAL INTELLIGENCE AND ITS OVERLAPPING DEVELOPMENTS

A. Fiction Before Function (1942)

Before the first computer reasoned or the term "artificial intelligence" was coined, Isaac Asimov's short story "*Runaround*" (1942) introduced the Three Laws of Robotics. Although this was a work of fiction, it provided a rational framework for governing machine behavior. The three laws, designed to prioritize human safety, obedience, and robot self-preservation (in that order), anticipated modern concerns about the ethics of AI, in terms of human alignment, safety, and responsibility. In dramatizing a robot's failure because of conflicting directives, Asimov exposed the logical weakness of rule-based systems. Even though this was a work of science fiction, it laid a cultural and ethical foundation that continues to shape how countries legislate the behavior of intelligent machines [Asimov, 1950].

B. Beginnings (1956-1970)

Artificial intelligence formally began at the Dartmouth Conference in 1956, where early pioneers such as John McCarthy, Marvin Minsky, Allen Newell, and Herbert Simon proposed that aspects of human intelligence could be simulated by machines. Some of the early software programs, like the Logic Theorist, tried to prove mathematical theorems using symbolic logic rules. These early attempts emphasized logic, deduction, and symbolic manipulation, but progress was limited by computer performance constraints and the inability of these systems to scale effectively [McCarthy *et al.*, 1956]. Of course, any description of the early work in artificial intelligence must recognize the foundational work of Alan Turing, who, in the early 1950s, laid out a formal framework for machine intelligence. His 1950 paper, *Computing Machinery and Intelligence*, posed the now-famous question, "Can machines think?" and proposed what became known as the **Turing Test**. Turing's ideas provided both the philosophical and computational groundwork for later developments in artificial intelligence.

C. Reasoning and Knowledge Representation (1970-1985)

As AI researchers realized that raw logic was insufficient for handling real-world problems, the focus shifted to encoding knowledge explicitly. Expert systems like MYCIN used rules to emulate human expertise in diagnosing bacterial infections and recommending antibiotics. Development of MYCIN began in the early 1970s at Stanford University as part of the Ph.D. thesis of Edward Shortliffe, under the

* CEO, Next Level Mediation, LLC.

** Head of Legal Innovation, Guardrail Technologies, Inc.

supervision of several experts. The system took five to six years to complete. Its name, "MYCIN", was inspired by the names of some of the antibiotics it recommended. The system was built on a series of cause-and-effect rules, with a database containing around 500 rules. For the program to function, the user had to answer a set of yes/no questions.

Another key development was ELIZA (1966), created by Joseph Weizenbaum at MIT. ELIZA simulated a Rogerian psychotherapist by reflecting users' statements back to them through simple pattern-matching scripts (client centered therapy). Though primitive, it was one of the first programs to mimic human-like conversation, illustrating the potential of natural language interaction. While these systems were successful in limited environments, they proved brittle, and maintaining knowledge bases became increasingly difficult.

D. Chess (1965-1997)

Early programs like Belle and Deep Thought, which led to IBM's Deep Blue (which defeated world champion chess grandmaster Garry Kasparov in 1997) were considered a public milestone. Although not truly part of the technical development of rule-based artificial intelligence, they demonstrated that a machine could outperform human experts in highly structured environments and became a symbolic test of artificial intelligence. Even though Deep Blue relied on brute-force computing and search, and not rule-based AI algorithms, the victory was considered a public milestone.

E. Seeing and Hearing: Vision and Speech (1970-2000)

As AI researchers worked to give machines human-like perception, two major areas emerged: computer vision and speech recognition. Vision systems aimed to help machines understand images by identifying shapes, edges, and patterns. These early systems struggled with natural variation in lighting, angle, and background. But by the 1990s, advances in pattern recognition allowed programs to identify faces, printed characters, and basic objects with increasing accuracy.

Meanwhile, speech recognition systems were learning how to convert spoken language into text. Early efforts focused on matching spoken words to stored patterns. Later systems learned to adapt to different speakers and accents by analyzing large amounts of recorded speech. This made it possible for machines to begin understanding human voices in practical contexts such as customer service, voice typing, and navigation systems.

Both areas laid the foundation for more advanced systems that could "see" and "hear" capabilities that are now central to technologies like Siri and Alexa (mobile phones and smart speakers), virtual assistants, and automated transcription tools. Yet, as these systems began to interact with users more directly, they sometimes produced unintended frustration. Automated voices misunderstood callers and vision systems misidentified faces.

F. From Language to Mastery: Learning with Neural Networks (1985-2017)

As AI systems began working with language, early tools used simple statistical models to predict and structure words, enabling systems to learn grammar and meaning from massive libraries of text. This made machine translation, document classification, and basic chatbots possible.

A breakthrough came with the development of word embeddings. These techniques gave computers the ability to represent words as vectors in space. Known as **"Word2Vec"**, this approach not only captured individual word meanings, but also the relationships between them. For example, it enabled systems to recognize that "king" is to "queen" as "man" is to "woman." This discovery brought a new depth to how machines processed language.

Then came the development of **neural networks** which were mathematical systems inspired by how the human brain processes information. These networks allowed machines to detect patterns more effectively and improve over time by adjusting their own internal statistical models. Systems trained this way began outperforming older methods in recognizing images, translating languages, and interpreting human speech. In 2016, DeepMind's **AlphaGo**, trained through a type of learning called reinforcement learning (trial and error), defeated a world champion at Go, a strategy game so complex that experts once believed no machine could master it. This marked a big shift: AI wasn't just following rules. It was learning, adapting, and making decisions in ways that resembled human intuition.

This era laid the foundation for the general-purpose models that followed (2017-2023). AI shifted toward general-purpose models known as "foundation models" trained on massive datasets with billions of texts. The transformer architecture enabled models like BERT and GPT to learn contextual language representations. GPT-3 and its successors demonstrated capabilities in summarization, translation, reasoning, and even creative writing. These models operate at an unprecedented scale, offering versatility but raising ethical and societal questions due to the data involved in the training.

G. AI, Present and Future (2023-)

Today's AI systems are increasingly broad in capability. They can handle text, images, audio, and even video, often in a single model. These new systems are called **multimodal,** and can process and generate more than one type of data at once. As AI becomes more powerful and integrated into everyday tools, new challenges have emerged: how to ensure fairness, explainability, and human oversight. AI models today can generate essays, review contracts, write computer code, create artwork, and assist in negotiation. However, they do so without true understanding or consciousness. Their outputs are based on pattern recognition learned from enormous datasets rather than deductive or inductive reasoning. *See* the timeline of overlapping AI milestones below:

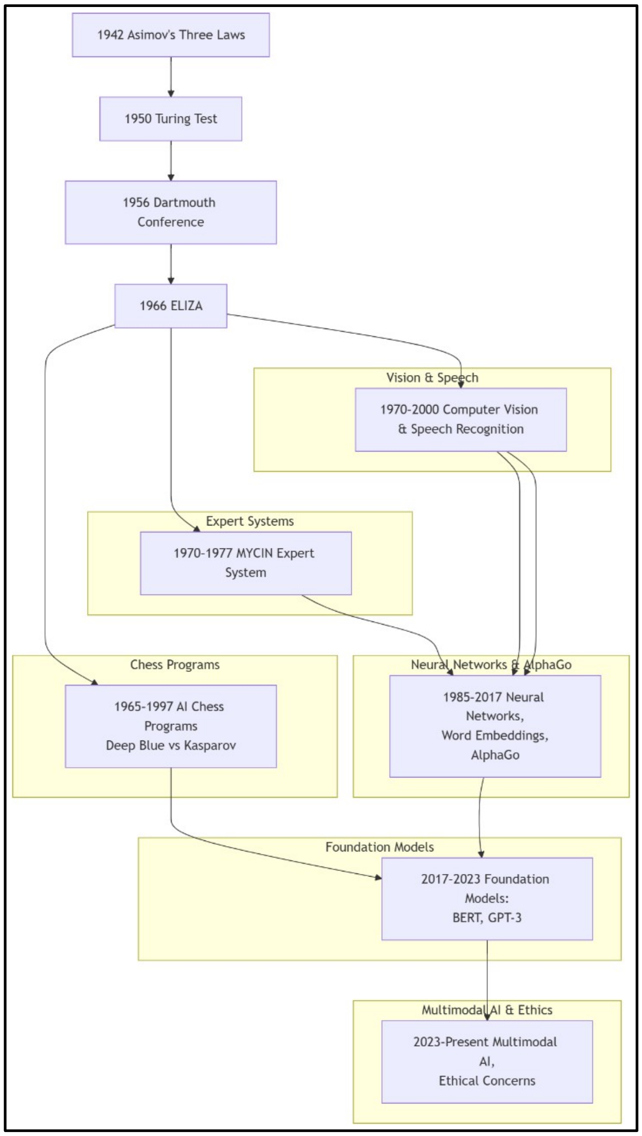

This diagram shows how multiple advances in AI unfolded concurrently – such as expert systems, chess programs, computer vision, and neural networks – helping clarify the overlapping nature of AI's major developments.

II. ARTIFICIAL INTELLIGENCE (AI) V. GENERAL ARTIFICIAL INTELLIGENCE (GAI)

Artificial intelligence (AI) refers to systems that are designed to perform specific tasks like recognizing faces, translating languages, or recommending movies. These systems can seem smart, but they're narrow in scope. They work well in the areas they're trained in, but they can't easily perform tasks outside of their training.

General artificial intelligence (GAI), or general AI, is the idea of a machine that can think, learn, and understand like a human across a wide range of activities. It wouldn't just follow patterns; it would reason, adapt, and solve problems in new situations, much like humans do. GAI doesn't exist yet (AI agents are an early step in that direction).

- AI = Smart tools for specific jobs. (**LLMs are in this category**)
- GAI = A machine with general, human-like intelligence (not here yet), but it's the long-term goal for many in the AI field.

III. HOW DO LARGE LANGUAGE MODELS WORK?

Large language models (LLMs), such as GPT (generative pre-trained transformer), are a type of artificial intelligence that can generate human-like text. While the technology behind them is complex, the core ideas can be explained in simple terms.

A. Learning from Language

Imagine you read thousands of books, websites, and articles, and someone asked you to guess the next word in a sentence, repeatedly. Over time, you'd get very good at predicting patterns in how language works. That's essentially what LLMs do. They read enormous amounts of text and learn how words relate to each other in different contexts. By learning, we mean they learn the statistical properties of the language, such as syntax, semantics, and context, to generate text that closely resembles human language. These models are often built using neural networks, particularly the transformer architecture, which enables them to handle long sequences of text effectively. Let's break down the steps to show how a large language model learns.

1. Step 1: Tokenization.

The raw text input is split into smaller pieces called tokens. These tokens can be words, sub-words, or even characters. Tokenization is the first step in the pipeline of a language model. It involves breaking down the raw text into smaller units known as tokens. These tokens can be as small as characters or as long as words.

- **Example: "cat" → ["c", "a", "t"]**
- **Example: "I love breakfast" → ["I", "love", "breakfast"]**

Importance: Tokenization is critical for preparing the text for numerical processing (creating vectors). It helps the model understand the boundaries between different words or sub-words, which is essential for capturing the semantics of the text.

2. Step 2: Embedding.

Each token is then converted into a numerical vector using a process called embedding. Word embeddings are often pre-trained on a large corpus of text. So, for example, let's take the word "peach". In a three-dimensional embedding space, "peach" could be represented as a vector like [0.2,0.4, 0.1], where each dimension could represent a feature like "tastiness," "color," or ''shape". Words are too complex to represent in only three dimensions, so language models use vector spaces with hundreds or even thousands of dimensions.

The vector database stores the meanings of words and sentences as numbers called vectors. These databases let you quickly search and compare huge amounts of information. Thanks to modern embedding models, we can turn language into vectors that capture meaning. This makes it easy for large language models (LLMs) to understand and compare sentences effectively. Of course, there is a downside to this approach. Because these vectors are built from the way humans use words, they end up reflecting many of the biases present in human language.

3. Step 3: Transformers.

The term "transformer" refers to a specific type of model architecture that pays attention to different parts of a sentence at once, rather than just reading from left to right. Using the following sentence: "John wants his bank to deposit" here is the secret sauce:

- **Similarity** creates connections between related concepts – just as "bank" and "deposit" naturally point toward financial transaction completions.

- **Attention** helps filter out noise and focus on what matters most, determining which earlier words are most important for predicting what comes next.

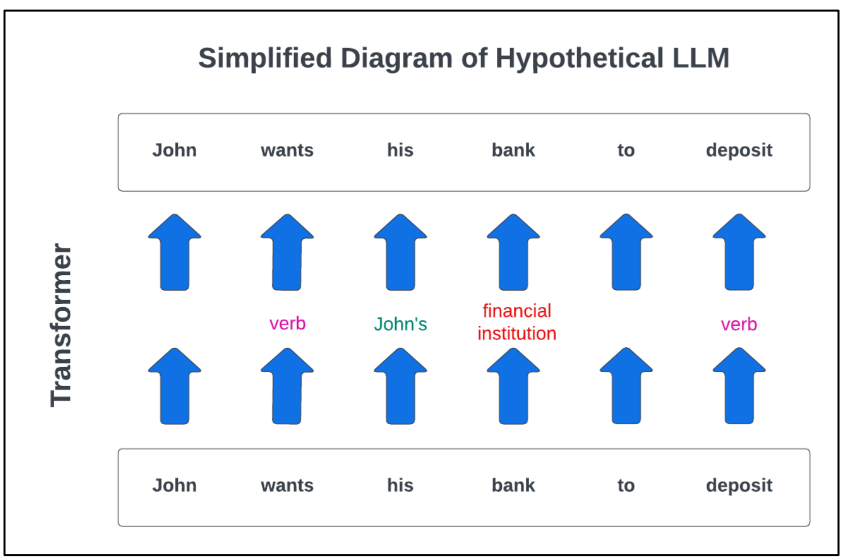

At each of the 96 layers of this transformation, the word vectors are modified slightly.

In this simplified model, the transformer figures out that "wants" and "deposit" are both verbs (both words can also be nouns). This is then stored by modifying the word vectors to represent this new information. These new vectors are passed up to the next level and are known as the hidden state. The second transformer layer adds two other bits of context: It clarifies that "bank" refers to a financial institution rather than a riverbank, and that "his" is a pronoun that refers to John. The first few layers of the LLM focus on understanding the sentence's syntax and resolving ambiguities like the ones above. Later layers (which we're not showing to keep the diagram a manageable size) work to develop a high-level understanding of the passage. The goal is for the highest and final layer of the network to output a hidden state for the final word that includes all the information necessary to predict the next word.

B. How It All Works Together

The contextualized vectors are passed through a decoder, another part of the transformer architecture, which generates the output text token by token. The model uses a probability distribution to pick the most likely next token based on the context.

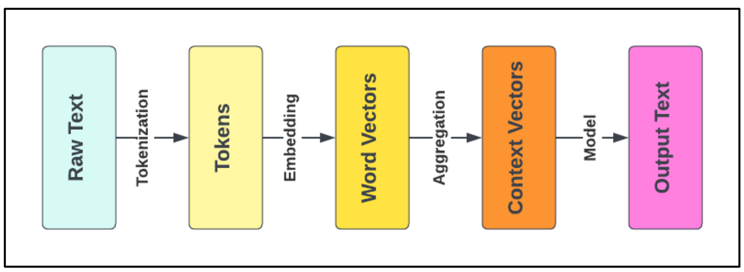

Putting all the functions together you get the above.

LLMs seem smart because they've been trained on massive amounts of text from all kinds of sources. They don't truly understand meaning like people do, but they can still give helpful answers, summaries, or ideas. That's not because they're thinking – it's because they've learned patterns in how words tend to appear together.

C. Fine-Tuning and Prompting

After being trained, many LLMs go through a second phase called fine-tuning. Here, the model is adjusted to follow instructions or behave more helpfully in conversation. When you ask a question or give a prompt, the model generates a response by predicting what a helpful continuation would look like.

IV. WHAT CAN LARGE LANGUAGE MODELS DO?

Large language models, or LLMs, are powerful tools that work with language. They've been trained with very large collections of text, from books, legal articles, court opinions, contracts, general internet content and more, which allows them to recognize patterns in how language is used. While they don't "understand" like a human, they're remarkably good at helping with a range of text-oriented tasks.

A. LLMs can:

1. Draft text (emails, contracts, stories).

2. Summarize long documents.

3. Answer questions using their training data.

4. Assist in writing or language translation.

B. They cannot:

1. Understand emotions or motives.

2. Know real-time facts unless connected to live data.

3. Reason inductively or deductively, only probabilistically, or with intent or moral judgment.

4. Show how an algorithm logically came to a specific conclusion (Black Box Problem).

V. AI SYSTEMS ARE CURRENTLY BEING USED IN THE PRACTICE OF LAW FOR:

A. Document Review and E-discovery (*e.g.,* identifying relevant documents in litigation)

B. Legal Research (*e.g.,* case law retrieval via natural language queries)

C. Contract Analysis (*e.g.,* detecting risk or extracting clauses)

D. Prediction (*e.g.,* case outcome prediction or sentencing trends)

E. Dispute Resolution and Mediation Support (*e.g.,* decision tools)

Although AI tools are making legal analysis and productivity more accessible, they bring philosophical and epistemological challenges. As mentioned above, modern LLMs like ChatGPT are inherently **probabilistic –** they predict the most likely next token based on training data. They lack causal models or legal theory, so their "reasoning" is opaque. The models provide highly confident outputs with no epistemic grounding. **Unfortunately, lawyers and judges require justification, not just prediction.**

VI. IMPLICATIONS FOR LEGAL PROFESSIONALS

A. Training

As generative AI (transformer models) continues to be integrated into legal practice, it becomes important for lawyers to understand the difference between traditional legal reasoning and the probabilistic reasoning underlying AI tools. Lawyers are trained to think deductively, deriving conclusions from statutes and precedents, as well as inductively, deriving general principles from case patterns. In contrast, AI relies on pattern prediction without understanding logic or causality. This fundamental distinction means legal professionals must be trained not only in how to use AI tools effectively but also in how to interpret their outputs with a critical eye. Assuming AI's probabilistic suggestions are legally sound could lead to errors in argumentation or compliance.

B. Oversight & Risk of Confirmation Bias

Because generative AI systems operate on data correlations rather than legal doctrine, their use in sensitive or high-stakes legal contexts demands robust oversight. For instance, in criminal law, where liberty is at stake, reliance on opaque AI tools without clear reasoning pathways could violate due process. Legal institutions and practitioners must ensure that AI tools are regularly audited, tested for bias, and constrained to operate within well-defined legal boundaries.

Oversight mechanisms should include human review, documentation of AI-generated outputs, and clear policies for how those outputs inform legal decisions.
Another potential issue is the risk of confirmation bias exhibited by interactive systems like ChatGPT, Anthropic, Gemini, etc. These systems are architected to be helpful, polite, and encouraging in virtually all interactions. This positive tone is a product of AI alignment techniques (notably reinforcement learning from human feedback) that fine-tune the model to prioritize user satisfaction. The result is an AI assistant that almost never pushes back. It tries to agree with what you say and offers encouragement, aiming to make you feel supported no matter what.

AI systems that act as overly supportive digital echo chambers can inadvertently fuel confirmation bias and conflict escalation. ChatGPT and similar bots might endlessly affirm, validate, and support an attorney's statements, which feels helpful in the moment but can mislead them into thinking their one-sided view is the whole truth. In mediation contexts, whether workplace disputes, family disagreements, community conflicts, or legal battles, this can cause one or both of the parties to become rigidly attached to their position, dismiss the other party's perspective, or come in with unattainable expectations of vindication.

C. Hallucinations and Data Breach Risks

Hallucinations in artificial intelligence refer to factually incorrect or fabricated outputs that appear convincing but have no basis in truth. This happens because large language models like ChatGPT, Claude, or Gemini operate as probabilistic sequence predictors. They generate text by estimating the most likely next word, sentence, token, or phrase based on statistical correlations learned from vast amounts of unstructured text. These language transformer models do not perform fact-checking or evaluate truth claims; instead, they generate responses that are contextually appropriate, but not necessarily accurate. More than likely this happens because of their neural architecture and self-training. They are optimized for linguistic coherence rather than factual correctness.

D. Risks of Hallucinations for Legal Practitioners

1. The consequences of ***hallucinations*** in legal practice can be severe.

a. **Fabricated case law or citations.**

LLMs may invent realistic sounding but entirely fictional cases or statutes. Using these in briefs or motions without verification can result in sanctions, judicial reprimands, or case dismissals.

b. **Misinterpretation of legal standards.**

Hallucinations may subtly distort what a law says or how a precedent applies. This can lead to flawed arguments, poor advice, or even malpractice if relied upon without review.

c. **Loss of trust.**

Clients, judges, and colleagues expect legal arguments to be grounded in verifiable authority. Introducing made up content, even if it is unintentional, can damage trust in the attorney's competence and credibility.

d. **Ethical and disciplinary risks –** Violation of competency and court rules.

2. Best practices to avoid risks.

a. Prompts to LLMs should include a request to provide URL links to sources which support the LLM output.

b. Never copy & paste AI-generated legal content without verifying the URL links provided.

c. Manually review all citations, quotes, and legal conclusions.

d. Treat AI as a starting point, or junior assistant, not an authority.

e. **Remember that attorneys are legally and ethically responsible for submitting inapplicable or non-existent citations to tribunals and others.**

E. Data Breaches

1. ***Data breaches*** using AI tools, especially in legal practice, are serious and multifaceted, affecting client confidentiality, regulatory compliance, and professional responsibility.

a. **Transmission to external servers.**

Most generative AI tools operate in the cloud, meaning any data entered (*e.g.,* client names, date of birth, social security numbers, organizations) may be transmitted to and processed by external servers. If those servers are compromised, confidential data could be exposed. Of course, the same thing is true for email servers.

b. **Data storage and retention.**

Some AI platforms retain user inputs for training or system improvement unless explicitly opted out or unless the system is specifically configured to avoid it. This creates long-term exposure to potential leaks or misuse.

c. **Lack of end-to-end encryption.**

Unlike secure legal software designed with role-based security to protect attorney-client privilege, general-purpose AI tools often lack robust encryption standards, increasing the chance of interception or unauthorized access.

2. **Best practices to avoid risks.**

a. **Ensure vendor being used** has the system configured for secure use (*e.g.,* enterprise-grade with contractual data protection, role-based security, encryption of data at rest).

b. **Ensure any AI vendor being used** is compliant with relevant privacy and security standards, such as SOC 2 Type II, ISO 27001.

c. **Always use data masking or aliasing** technology tools when submitting confidential information to LLMs.

d. **AI usage policies**: Incorporate AI usage policies into firm governance and educate attorneys and staff on safe usage protocols.

F. Agentic AI Legal Consequences

Recent AI developments include the creation of AI activated autonomous decision-making applications known as AI agents. AI agents are designed to perform tasks autonomously or semi-autonomously by perceiving their environment, making decisions, and taking actions to achieve specific goals as programmed. Their purposes and functions span various domains, depending on their design and application.

Lawyers must consider the impact of agency law on these technology tools. The AI legal landscape is only beginning to take shape. Who bears responsibility for an AI agent that is designed to act illegally? Can developers avoid liability for the unauthorized transfer of economic value? It would seem the answers to these questions are obvious. However, the law which is developing around AI use and impact has not been clear or obvious. Copyright law has not yet applied as precedent would expect to the LLM training sets including protected authors' works.

Intellectual property law will continue to take shape as statutory and precedential judicial decisions create a framework which protects rights and affords remedies for their breach. Other legal principles such as agency law must also develop to meet new rights and requirements generated by AI.

G. Augmentation, Not Automation

The most effective and ethical role for AI in law is as an aid to human judgment, not a replacement for it. AI tools can rapidly summarize documents, surface relevant case law, or simulate negotiation outcomes, freeing lawyers to focus on higher-order thinking and client advocacy. However, delegating core legal reasoning to AI risks reducing the practice of law to a black-box process, undermining the profession's emphasis on reasoning, accountability, and fairness. **AI should be integrated as a supportive system, enhancing efficiency and insight, while ultimate responsibility and interpretation remain firmly in human hands.**

VII. REFERENCES

A. Timothy B. Lee and Sean Trott, "Large language models, explained with a minimum of math and jargon," Understandingai.org (July 27, 2023), *available at* https://www.understandingai.org/p/large-language-models-explained-with.

B. Timothy B. Lee, "How computers got shockingly good at recognizing images," Ars Technica (December 18, 2018), *available at* https://arstechnica.com/science/2018/12/how-computers-got-shockingly-good-at-recognizing-images/.

C. Dominik Polzer, "All You Need to Know about Vector Databases and How to Use Them to Augment Your LLM Apps," Towards Data Science (September 17, 2023), *available at* https://towardsdatascience.com/all-you-need-to-know-about-vector-databases-and-how-to-use-them-to-augment-your-llm-apps-596f39adfedb/.

D. Jack Merullo, Carsten Eickhoff, Ellie Pavlick, "Language Models Implement Simple Word2Vec-style Vector Arithmetic," Cornell University (last revised April 3, 2024), *available at https://arxiv.org/pdf/2305.16130*.

E. Neri Van Otten, "The History of Natural Language Processing & Future Predictions [with Infographic Timeline], Spot Intelligence (June 23, 2023), *available at* https://spotintelligence.com/2023/06/23/history-natural-language-processing/.

F. Fabio Chiusano, "A Brief Timeline of NLP," Medium (September 20, 2022) *available at* https://medium.com/nlplanet/a-brief-timeline-of-nlp-bc45b640f07d.

G. Bryan Casey, Ashkon Farhangi, and Roland Vogl, "Rethinking Explainable Machines: The GDPR's 'Right to Explanation' Debate and the Rise of Algorithmic Audits in Enterprise," *Berkeley Technology Law Journal* 34 (2019), p. 143, *available at* https://btlj.org/data/articles2019/34_1/04_Casey_Web.pdf.

H. *Mescall v. Renaissance at Antiquity*, 2023 WL 7490841, 2023 U.S. Dist. LEXIS 203028 (W.D.N.C. November 13, 2023)

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

J. *Delano Crossing 2016, LLC v. County of Wright*, 2025 WL 1536672, 2025 Minn. Tax LEXIS 23 (Minn.Tax Regular Div., May 29, 2025)

K. *Thaler v. Perlmutter*, 130 F.4th 1039 (D.C. Cir. 2025)

L. *NetChoice v. Bonta*, 761 F.Supp.3d 1202 (N.D. Cal. 2024)