Transportation Lawyers Association

Understanding the Ethics of Generative AI

2025 Chicago Regional Conference

By: Megan K. Baker,[1] William M. Davis[2] and Sophia Rago[3]

How is AI Different from What I've Used in the Past

All attorneys are familiar with database systems. A database is an organized collection of data, stored electronically, that allows for efficient access and management of large amounts of information. Litigators will be familiar with the online offerings from Lexis and Westlaw which have replaced libraries of law books. Transactional lawyers are familiar with corporate databases that house thousands of contracts, government repositories of data such as property ownership records, *etc*.

Traditional database systems are static repositories of information. Data are manually entered and retrieved using "structured queries," such as Boolean searches in Westlaw or Lexis Nexis.[4] For average users who aren't database programmers, querying a traditional database is challenging because explicit instructions are required to fetch specific records. Crafting a good structured query requires technical expertise. Additionally, once records are retrieved, they are not necessarily presented in a format that is helpful to the user. For example, the results may simply be ordered in the sequence in which they appear in the database rather than in an order that is relevant to the user.

Traditional databases do not identify patterns or trends in data. Analysis of data is possible with a traditional database, but it is entirely dependent upon the user formulating a specific query to generate the analysis, and further relies upon the user to analyze the results.

Traditional database systems are limited to storing and retrieving data. The user is left to manually process and analyze the retrieved information.

Finally, traditional databases struggle with analyzing very large datasets in real time. The amount of computing power required for real time responses to queries on a huge dataset is considerable.

AI solves these problems with traditional databases by creating a system that not only retrieves information but also interprets, processes and learns from it. AI can generate

---

[1] Chartwell Law (Boston, MA).

[2] Bovis, Kyle, Burch & Medlin, LLC (Atlanta, GA).

[3] Canty Novy Bertkau Gordon, LLC (Chicago, IL).

[4] A Boolean query is a search that uses Boolean operators to combine keywords and create logical phrases. It is named after George Boole, a 19th-century mathematician who developed the system of logic. Here's an example of such a query: tax taxation % income /3 tax taxation. This search would return database entries that contain the words "tax" or "taxation", but it would exclude documents where those terms appear within 3 words of "income."

insights, summarize data, predict outcomes and adapt to new data over time. Additionally, AI can interpret a user's natural-language query in order to generate results, thus obviating the need for the "structured queries" that were necessary to easily find records in traditional databases.

AI can automate repetitive tasks and order the presentation of information in a way that is responsive to the user's query, rather than relying on the user to sift through multiple results to find the answer. Results are prioritized and tailored to the user's intent automatically.

AI uses machine learning algorithms to detect patterns, trends and correlations in vast amounts of data. AI can learn, for example, your firm's preferences for document formatting and can produce documents in that format.

What is AI

AI—Artificial Intelligence—refers to computer systems which perform tasks that would ordinarily require human intelligence, such as reasoning, decision-making, understanding language, learning and problem-solving. AI systems are built to analyze data, recognize patterns in the data, and make predictions based on those data.

Although AI systems can differ widely in terms of how they are built, most AI systems have the following core components:

1. Machine Learning: This process allows a system to learn from data to improve its performance without being explicitly programmed to do so. This aspect of AI is what allows a system to predict future outcomes automatically based on prior outcomes.

2. Computer Vision: This process focuses on interpreting visual data, such as analyzing scanned contracts for specific clauses, or analyzing charts which visually depict data, such as a pie chart.

3. Natural Language Processing: This process enables AI to understand and generate human language. This is what powers tools like legal document review software or chatbots.

4. Automation: AI typically employs additional standard computer systems to execute repetitive tasks faster and with fewer errors.

AI requires large datasets to "train" itself. AI "learns" by analyzing the datasets to identify patterns and trends. For example, in the legal context an AI system might review thousands of contracts to flag non-standard clauses or to calculate the risk of litigation. Predictive AI analyzes past outcomes (from litigated cases, for example) and assesses the likelihood of success in similar cases.

As has always been the case with data analysis, the AI's ability to function effectively depends largely on the quality of the data it ingests for analysis. Biases in the data can lead

to biased outcomes. Incomplete data can also impair the functioning of an AI. Human oversight is crucial for tasks requiring nuanced judgment.

What Are Some Concerns About the Use of AI in Legal Practice

AI's learning component works because it analyzes and stores users' queries. Because lawyers are entrusted with their clients' secrets, a query session with an AI may inevitably require that an attorney input a question containing confidential information or setting forth a legal strategy that would be considered attorney work product.

There are a number of ways in which confidential information could be inadvertently disclosed when using AI. Any AI system that would currently be used by an attorney is cloud-based.[5] Query data could be stored across multiple computer systems in different geographical areas, raising concerns about who has access to the data and whether it is subject to a breach by hackers.

There are ethical considerations with respect to the use of any software by attorneys. Attorneys must stay informed about the benefits and risks of technology to maintain competence in their practice. Unlike other forms of software that have entered the legal marketplace in the past, AI presents some particular challenges due to its cloud-based nature and ability to store and learn from queries. From a practical standpoint, it can be difficult for lawyers to appreciate the functionality of a particular AI software given that there are various types of AI systems, each of which function differently. Add to that the fact that "AI" is used ubiquitously as a marketing term in the current economy, even for products that don't possess true AI abilities.

The algorithms used by various AI platforms are often proprietary, and even when they are not, understanding how they function requires specialized training. The result is that AI output comes from a "black box," making it difficult for an attorney to understand or explain how decisions or recommendations were made. Additionally, while AI tools can be highly efficient, they can also produce errors and misinterpretations—particularly in areas requiring nuanced judgment such as legal analysis. These errors are commonly called "hallucinations" in the AI industry.

When AI responds to a prompt, it generates output based on patterns it has learned, but it doesn't actually "know" or "understand" the truth or correctness of the information. Hallucinations can occur due to gaps in training data, ambiguous prompts, overgeneralization by the AI and a lack of contextual understanding by the AI. In the legal world, AI systems have been known to fabricate case law in wholesale fashion and present it as genuine. AI might state that a legal principle applies universally, when in fact it varies from jurisdiction to jurisdiction. AI might also invent procedural steps which don't exist in reality (*i.e.* a "mandatory arbitration form").

[5] The term "cloud" as used in this sense means that the software is hosted on remote servers and accessed via the internet, rather than being installed and run locally on a user's computer. The term "cloud" refers to a network of servers that store, manage and process data in a centralized or distributed manner, making resources available on demand.

A study conducted in 2023 by Stanford University's RegLab and Institute for Human-Centered Artificial Intelligence found that large language models from the top tech companies hallucinated 69% to 88% of the time when answering legal questions, and their performance worsened as the difficulty of the questions increased.[6]

What Areas of the Law Will AI Most Likely Impact

This is the $100,000 question. No one knows for sure, but a review of current use cases suggests the following:

1. Transactional Work: Because AI can process massive amounts of data much faster than humans, it is useful in transactional legal work where there are repetitive tasks and repetitive language. AI can scan thousands of contracts and identify nonstandard language much more quickly than humans.

2. E-Discovery: In litigation involving hundreds of thousands of pages of data (for example, the email inboxes of multiple employees), AI can produce digests and summaries of vast quantities of data which are useful to litigators who are looking for a needle in a haystack.
3. Legal Research: Westlaw and Lexis have enhanced legal research tools which are AI-powered. This makes an attorney's job easier by allowing natural language queries and providing more relevant results that are tailored to the attorney's needs.

4. Analysis of Large Sets of Reliable Data: In cases where data sets are reliably maintained, AI's abilities to analyze the data really shine. Areas of the law which frequently use reliable data sets include tax law, contract review/drafting, legal research, regulatory compliance, and intellectual property.

As to the last item above, an example of a reliable data set would be an online law library such as that offered by Westlaw or Lexis. The data set is reliable because it has been professionally-created, has been reviewed for errors, and the data have been formatted and categorized to ensure that a query will retrieve accurate results.

An example of an unreliable database is a local government database of driver's licenses. That data set is likely unreliable for a number of factors, such as: a) data are entered by different employees, at different locations, with little or no quality control; b) the database relies upon drivers to update their information as it changes, which is rarely done (many people do not get a new driver's license when they move, even though it is required); c) much of the data are self-reported and are not confirmed, even though the data are subject to change (height/weight information, hair color, *etc*.).

One of the best reasons for relying on the advice of local counsel in a jurisdiction is local counsel's knowledge of aspects of the law that are not currently amenable to incorporation in a reliable database. Such areas include: 1) juror tendencies in a given jurisdiction; 2)

[6] Researchers Find Legal Errors 'Pervasive' In Top AI Models - Law360 Pulse

judge tendencies in a particular courtroom; 3) negotiation style of counterparties; 4) settlement value of cases; 5) likelihood of litigation; 6) appropriate child custody outcomes; 7) appropriate criminal sentencing outcomes; 8) novel legal issues; 9) application of law to fact-intensive scenarios; 10) evaluation of highly emotional or moral issues; 11) international law and cross-border disputes; 12) equitable remedies; and,13) small claims.

Take small claims as an example. Most small claims, such as landlord-tenant disputes, are handled in administrative "small claim" courts in various jurisdictions. Such courts are not "courts of record," so very little data are generated in such cases (if any at all). Even though one of these courts might process thousands of such cases in a given year, very little data are left behind in a database for AI to learn from. The same is true for verdict outcomes, juror analysis, negotiation styles, *etc*. Unless a large, reliable data set of information is available, AI has nothing to analyze.

What is the ABA's Position on AI?

On July 29, 2024, the ABA issued Formal Opinion 512 which covers Generative Artificial Intelligence Tools. The term "generative AI" refers to artificial intelligence models that are capable of generating new content, such as text, images, music, code, 3D models, *etc*. based on prompts by a user. The word "generative" highlights the ability of these models to create rather than just analyze or classify data. Generative AI is thus different from most e-discovery AI platforms that analyze and tag data.

The fifteen-page opinion emphasizes that the use of Generative AI implicates several of the existing Model Rules of Professional Conduct.

Model Rule 1.1

Model Rule 1.1 obligates lawyers to provide competent representation to clients. The rule requires, among other things, that lawyers understand "the benefits and risks associated" with any software/technology used to deliver legal services to clients.

The opinion notes that there is no requirement that a lawyer using Generative AI become an expert in the technology.

> Rather, lawyers must have a reasonable understanding of the capabilities and limitations of the specific [Generative AI] technology that the lawyer might use. This means that the lawyers should either acquire a reasonable understanding of the benefits and risks of the [Generative AI] tools that they employ in their practices or draw on the expertise of others who can provide guidance about the relevant [Generative AI] tools' capabilities and limitation.[7]

[7] Formal Opinion 512, pp. 2–3 (*citing* Pa. Bar Ass'n, Comm. on Legal Ethics & Prof'l Resp. Op. 2020-300, 2020 WL 2544268, at *2–3 (2020); Cal. State Bar, Standing Comm. on Prof'l Resp. & Conduct Op. 2023-208, 2023 WL 4035467, at 2 (2023)).

Importantly, a lawyer's reliance on Generative AI output—without and appropriate independent verification of that output—could run afoul of the duties imposed by Rule 1.1. Wholesale reliance on Generative AI can easily amount to an abdication of the duty to exercise professional judgment.

Model Rule 1.6

Model Rule 1.6 requires that attorneys keep all information relating to the representation of a client confidential unless the client gives informed consent, or if the disclosure is impliedly authorized in order to carry out the representation, or if the disclosure is permitted by an exception to the rule. There are similar proscriptions in Model Rules 1.9(c) and 1.18(b) affecting confidentiality for former and prospective clients, respectively.

The drafters of Opinion 512 noted that the risk of disclosure of confidential information exists not only when information is entered into a Generative AI prompt (as discussed earlier in this paper), but also when it is disclosed or accessed by others "*inside* the firm who will not adequately protect the information from improper disclosure or use because, for example, the are unaware of the source of the information and that it originated with a client of the firm."[8]

The Opinion takes the position that, because Generative AI is self-learning, a disclosure of confidential information could easily occur when the AI responds to prompts by users outside the firm, or by other users in the same firm who would otherwise be prohibited from using the information. For that reason, the Opinion advises that "a client's informed consent is required prior to inputting information relating to the representation into [a Generative AI tool]."[9] The Opinion specifically states that "boiler-plate" provisions in engagement letters will be insufficient to create informed consent. Rather, the lawyer must reveal to the client, in his best judgment: 1) why the use of Generative AI is appropriate; 2) provide specific information about the risk of disclosure; 3) including details of the types of information that could be disclosed; 4) indicate the ways in which other parties might use disclosed information against the client's interests; and, 5) provide a clear explanation as to the benefit of using Generative AI in the representation.

In order to be competent to provide informed consent, the Opinion states that as a *baseline*, "all lawyers should read and understand the Terms of Use, privacy policy, and related contractual terms and policies of any [Generative AI] tool they use. . . or consult with a colleague or external expert who has read and analyzed those terms and policies."[10]

As of December 11, 2024, the Terms of Service for ChatGPT—one of the more popular free Generative AI offerings in the marketplace—are eight pages long. One short provision

---

[8] *Id.* at p. 6.

[9] *Id.*at p. 7.

[10] *Id.* (*citing* Stephanie Pacheco, *Three Considerations for Attorneys Using Generative AI*, BLOOMBERG LAW ANALYSIS (June 16, 2023, 4:00 pm), https://news.bloomberglaw.com/bloomberg-law-analysis/analysis-three-considerations-forattorneys-using-generative-ai?context=search&index=7).

addresses "opting out" from allowing ChatGPT to use prompts to train itself. That provision directs the user to another set of polices for ChatGPT which is an interesting read and can be found here: How your data is used to improve model performance | OpenAI Help Center. For business accounts, Open AI (the developer of ChatGPT) has the following policy as of the date of this paper:

> By default, we do not train on any inputs or outputs from our products for business users, including ChatGPT Team, ChatGPT Enterprise, and the API. We offer API customers a way to opt-in to share data with us, such as by providing feedback in the Playground, which we then use to improve our models. Unless they explicitly opt-in, organizations are opted out of data-sharing by default.[11]

However, it is likely that attorneys at law firms may use their own ChatGPT account, which is free. For individual accounts, the default setting is for ChatGPT to use prompt information to train itself. A user must expressly opt-out of this feature. This would be necessary in order to protect client information from being subsequently exposed.

One area which the Opinion leaves open to interpretation is: What information is "related to the representation" such that entering it as a Generative Ai prompt could potentially expose client confidentiality? It seems clear, for example, that entering information directly obtained from a confidential client document into a prompt would potentially expose that data to others in a self-learning AI model. But what about prompts that are related to the representation of the client, but don't actually use client-specific information? Or what about prompts that ask hypothetical questions about a planned future action by a high-profile client, but which don't specifically identify the client?

These are daunting questions considering that the comments to Rule 1.6 define "information relating to the representation" to include not just matters communicated by the client in confidence, but "all information relating to the representation, whatever its source."[12]

While Generative AI is a new and emerging technology, concerns over confidential client information being exposed to third parties is hardly a new issue. Twenty-nine years ago, the ABA issued an ethics opinion concerning the use of what has become known as "cloud databases." Formal Opinion 95-398 is entitled *Access of Nonlawyers to a Lawyer's Database*, and it addresses the use of database services that are now commonly known as "Dropbox," "OneDrive," "SharePoint," *etc*.

Prior to that, there are various ethics opinions regarding the use of copy services, document disposal services, *etc*. who handled physical client documents.

---

[11] https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance

[12] Model Rule 1.6, cmt. 3.

There is a common thread to all of these opinions, whether they are addressing the handling of old-fashioned paper documents, or sending prompts to a Generative AI: a lawyer must make sure that the third parties with whom the information is being shared will keep it confidential or the lawyer must otherwise take steps to ensure confidentiality. With paper documents, law firms shred or redact documents internally so that they cannot fall into the wrong hands once they are in the trash bin. Similarly with paper documents, law firms who engage outside shredding contractors require contractual assurances that the information will not be read and will not be disclosed. At the end of the day, this is no different than "opting out" of Generative AI using prompt data or reading the terms of service for a Generative AI to ensure that there are reasonable assurances that will protect confidential information.

Model Rule 1.4

Model Rule 1.4 addresses a lawyer's duty to communicate with her client. The Opinion notes that the duties prescribed by Rule 1.4 may require lawyers to disclose the use of Generative AI even when the Generative AI is not being prompted with information "relating to the representation" such that Rule 1.6 is triggered.

While the Opinion notes that the "facts of each case will determine whether Model Rule 1.4 requires lawyers to disclose their [Generative AI] practices to clients…"[13] the Opinion goes on to list some circumstances where disclosure would be required:

1. If the client inquires as to how legal work product was generated, the firm must disclose how Generative AI was used;
2. If the client requires in its terms of engagement or outside counsel guidelines that the firm disclose its use of Generative AI, the firm must disclose;
3. If the use of Generative AI is relevant to the basis or reasonableness of a lawyer's fee, then its use must be disclosed to the client;
4. If the output of Generative Ai will be used to influence a significant decision in the presentation, the use of Generative AI must be disclosed to the client;

As a best practice, the Opinion notes:

> Even when Rule 1.6 does not require informed consent and Rule 1.4 does not require a disclosure regarding the use of [Generative AI], lawyers may tell clients how they employ [Generative AI] tools to assist in the delivery of legal services. Explaining this may serve the interest of effective client communication. The engagement agreement is a logical place to make such disclosures and to identify any client

[13] Formal Opinion 512, p. 8.

> instructions on the use of [Generative AI] in the representation.[14]

Model Rules 3.1, 3.3 and 8.4(c)

These rules address a lawyer's ethical responsibilities to the court. Rule 3.1 requires, among other things, that "[a] lawyer shall not bring or defend a proceeding, or assert or controvert and issue therein, unless there is a basis in law or fact for doing so that is not frivolous." Rule 3.3 states that a lawyer cannot knowingly make any false statement of law or fact to a tribunal or fail to correct a material false statement of law or fact previously made to a tribunal. Finally, Rule 8.4(c) prohibits a lawyer from engaging in "conduct involving dishonesty, fraud, deceit or misrepresentation."

The previous discussions of AI hallucinations in this article are most relevant here. A lawyer who relies entirely on Generative AI content without verifying the accuracy of the output runs the risk of violating the above rules. There have already been several well-publicized examples of lawyers being sanctioned for filing briefs with fake case citations that were made up by Generative AI.[15]

Model Rules 5.1 and 5.3

These rules deal with the ethical duties of lawyers charged with supervisory responsibilities. The Opinion states, "Managerial lawyers must establish clear policies regarding the law firm's permissible use of [Generative AI], and supervisory lawyers must make reasonable efforts to ensure that the firm's lawyers and nonlawyers comply with their professional obligations when using [Generative AI] tools."[16]

Rule 5.3 extends some of these duties to persons outside the law firm when lawyers rely on such outside persons to provide services in connection with legal representation. Lawyers with direct supervisory authority over nonlawyers must make "reasonable efforts to ensure that" the nonlawyer's conduct conforms with the lawyer's own professional obligation.

The earlier discussion regarding ChatGPT's terms of service provides a good example of where Rule 5.1 might apply in a law firm context. It would seem to be a good practice for law firms to publish guidance either restricting the individual use of Generative AI programs which are not expressly authorized or requiring that users "opt out" of having prompt information used to train the particular Generative AI in use.

Model Rule 1.5

The Opinion concludes with a fairly lengthy analysis of the interaction between the use of Generative AI and Model Rule 1.5, which governs lawyers' fees and expenses.

---

[14] *Id.* at p. 9.

[15] *See* DC Bar Op. 388 (2024).

[16] Formal Opinion 512, p. 10.

For hourly billing, the Opinion observes: "If a lawyer uses a [Generative AI] tool to draft a pleading and expends 15 minutes to input the relevant information into the [Generative AI] program, the lawyer may charge for the 15 minutes as well as for the time the lawyer expends to review the resulting draft for accuracy and completeness."[17]

For contingent or flat fees, the Opinion was less instructive, noting only: "[I]f using a [Generative AI] tool enables a lawyer to compete tasks much more quickly than without the tool, it may be unreasonable under Rule 1.5 for the lawyer to charge the same flat fee when using the [Generative AI] tool as when not using it."[18]

Conclusion

Once upon a time, law firms struggled with the idea of how to incorporate computers into the workplace. Should there be just one computer in the library for everyone to use? Should everyone get a computer? Should we allow attorneys to take floppy disks home with them?

Then, one day, someone decided to connect those computers to a network, and beyond that, to the internet. Now there was a new host of questions. How should we prevent unauthorized access? How should we store our electronic data? How do we prevent exposing confidential information to the public domain through the internet?

Fast forward to the present day, and we find ourselves asking new questions as we grapple with Generative AI. In some ways, the questions can be answered in the same way that they were for old technologies. But there is one aspect of Generative AI that is different from the previous waves of technology that swept law firms: ease of access.

When computers were introduced, they were a substantial investment and were too big for a lawyer to tote around. These factors made it easy for firms to control and keep track of computers.

When networking and the internet were introduced, firms could employ firewalls and filters to control employee access.

Generative AI is a free service in many forms. ChatGPT, for example, is free and is accessible from desktop browsers, mobile phones, laptops—any electronic device. The ubiquity of Generative AI makes it difficult for law firms to keep track of when it is being used by attorneys and employees. Additionally, the rapid growth evolution of Generative AI models—with concomitant changes in terms of service and service providers—further complicates a law firm's efforts to track and manage the use of Generative AI within the firm.

Each firm will address these issues differently depending on practice areas, size, and a host of other factors. Regardless of the different needs of different firms, Formal Opinion 512 makes it clear that *every* lawyer must take some action to address the use of Generative AI.

---

[17] *Id.* at p. 12.

[18] *Id.*