

Ethical Considerations: Artificial Intelligence & The Law

2nd Annual Western Agricultural and Environmental Law Conference
National Agricultural Law Center Reno, NV June 2024

THE PROGRAM

Our inboxes are filled with news of artificial intelligence – its dangers, threats, promises, and possibilities. During this program, we will consider the practical uses for AI, generative and otherwise, in the delivery of legal services. Beginning with a demonstration and overview of categories of AI-powered legal tools, we'll discuss the possibilities for their use, as well as the ethical guardrails we must consider to best ensure client protection. We'll look at the states that have provided guidance on the use of generative AI and peek at what the future might hold.

The materials below are designed to provide an overview of these topics, including the current state ethics opinions and other guidance, relevant judicial orders on the use of AI, current articles related to AI in the delivery of legal services, and selected relevant law review articles.

PRESENTER BIOGRAPHY

[David A. McCarville](#) focuses his practice on succession planning, business and finance issues for partnerships and family-owned businesses. This often involves working closely with his clients and their professional advisors, as well as other Fennemore attorneys in various practice areas to provide expertise as needed. His clients appreciate his ability to identify and resolve legal issues in a timely and efficient manner.

Since 2018, he has been an Adjunct Professor at Arizona State University's Sandra Day O'Connor College of Law teaching a course titled Blockchain & Cryptocurrencies: Law and Policy.

David is co-chair of the Arizona Banking Association's Emerging Technology Committee, and is a member of the firm's Diversity and Inclusion Committee, as well as the Fennemore Foundation.

MATERIALS

State Ethics Opinions and Other Guidance Materials

The State Bar of California Standing Committee on Professional Responsibility and Conduct, *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law*, available at:

<https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-PracticalGuidance.pdf> (last visited April 13, 2024).

Florida Bar Ethics Opinion, *Opinion 24-1* (January 19, 2024), available at:

<https://www.floridabar.org/etopinions/opinion-24-1/>

Massachusetts Board of Bar Overseers (by Afton Pavletic), *The Wild West of Artificial Intelligence: Ethical Considerations for the Use of A.I. in the Practice of Law*, available at:

[https://bbopublic.massbbo.org/web/f/The Wild West of Artificial Intelligence.pdf](https://bbopublic.massbbo.org/web/f/The_Wild_West_of_Artificial_Intelligence.pdf) (last visited April 13, 2024).

State Bar of Michigan JI-155 (October 27, 2023), available at:

https://www.michbar.org/opinions/ethics/numbered_opinions/JI-155 (last visited April 13, 2024).

New Jersey Supreme Court Committee on AI and the Courts, *Preliminary Guidelines on New Jersey Lawyers' Use of Artificial Intelligence* (January 25, 2024), available at:

<https://www.njcourts.gov/sites/default/files/notices/2024/01/n240125a.pdf>

New York State Bar Association, *Report and Recommendations of the Task Force on Artificial Intelligence* (April 2024), available at:

<https://nysba.org/app/uploads/2022/03/2024-April-Report-andRecommendations-of-the-Task-Force-on-Artificial-Intelligence.pdf>

North Carolina State Bar, Proposed 2024 Formal Ethics Opinion 1 Use of Artificial Intelligence in a Law Practice, available at: <https://www.ncbar.gov/for-lawyers/ethics/proposed-opinions/> (last visited April 13, 2024).

Virginia State Bar, *Guidance on Generative Artificial Intelligence*, available at: <https://vsb.org/Site/Site/lawyers/ethics.aspx?hkey=bc8a99e2-7578-4e60-900f45991d5c432b> (last visited April 13, 2024).

Judicial Orders on AI

RAILS (Responsible AI in Legal Services) Compilation of Court Orders on AI, available at: <https://rails.legal/resource-ai-orders/> (last visited April 13, 2024).

Selected Current Articles on AI in the Delivery of Legal Services

What is Artificial Intelligence (AI)?, via IBM, ibm.com/topics/artificial-intelligence (last visited April 13, 2024).

AI Terms for Legal Professionals: Understanding What Powers Legal Tech, LexisNexis (March 20, 2023), available at: lexisnexis.com/community/insights/legal/b/thought-leadership/posts/ai-terms-for-legal-professionals-understanding-what-powers-legal-tech (last visited April 13, 2024).

John Villasenor, *How AI Will Revolutionize the Practice of Law*, Brookings Institution (March 20, 2023), available at: brookings.edu/articles/how-ai-willrevolutionize-the-practice-of-law/ (last visited April 13, 2024).

Selected Recent Law Review Articles

Murray, Michael D., *Artificial Intelligence and the Practice of Law Part 1: Lawyers Must be Professional and Responsible Supervisors of AI* (June 14, 2023), available at: <https://ssrn.com/abstract=4478588>

Perlman, Andrew, *The Legal Ethics of Generative AI* (February 22, 2024). Suffolk University Law Review, Forthcoming, available at: <https://ssrn.com/abstract=4735389>

Model Rules of PR Relevant to AI

Model Rule 1.1: Competence

A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.

Cmt [8] To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject.

Model Rule 1.6 – Confidentiality of Information

(c) A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.

Cmt [18] Paragraph (c) requires a lawyer to act competently to safeguard information relating to the representation of a client against unauthorized access by third parties and against inadvertent or unauthorized disclosure by the lawyer or other persons who are participating in the representation of the client or who are subject to the lawyer's supervision. See Rules 1.1, 5.1 and 5.3. The unauthorized access to, or the inadvertent or unauthorized disclosure of, information relating to the representation of a client does not constitute a violation of paragraph (c) if the lawyer has made reasonable efforts to prevent the access or disclosure. Factors to be considered in determining the reasonableness of the lawyer's efforts include, but are not limited to, the sensitivity of the

information, the likelihood of disclosure if additional safeguards are not employed, the cost of employing additional safeguards, the difficulty of implementing the safeguards, and the extent to which the safeguards adversely affect the lawyer's ability to represent clients (e.g., by making a device or important piece of software excessively difficult to use). A client may require the lawyer to implement special security measures not required by this Rule or may give informed consent to forgo security measures that would otherwise be required by this Rule. Whether a lawyer may be required to take additional steps to safeguard a client's information in order to comply with other law, such as state and federal laws that govern data privacy or that impose notification requirements upon the loss of, or unauthorized access to, electronic information, is beyond the scope of these Rules. ***

Cmt [19] When transmitting a communication that includes information relating to the representation of a client, the lawyer must take reasonable precautions to prevent the information from coming into the hands of unintended recipients. This duty, however, does not require that the lawyer use special security measures if the method of communication affords a reasonable expectation of privacy. Special circumstances, however, may warrant special precautions. Factors to be considered in determining the reasonableness of the lawyer's expectation of confidentiality include the sensitivity of the information and the extent to which the privacy of the communication is protected by law or by a confidentiality agreement. A client may require the lawyer to implement special security measures not required by this Rule or may give informed consent to the use of a means of communication that would otherwise be prohibited by this Rule. Whether a lawyer may be required to take additional steps in order to comply with other law, such as state and federal laws that govern data privacy, is beyond the scope of these Rules.

Model Rule 5.3: Responsibilities Regarding Nonlawyer Assistance

With respect to a nonlawyer employed or retained by or associated with a lawyer:

(a) a partner, and a lawyer who individually or together with other lawyers possesses comparable managerial authority in a law firm shall make reasonable efforts to ensure that the firm has in effect measures giving reasonable assurance that the person's conduct is compatible with the professional obligations of the lawyer;

(b) a lawyer having direct supervisory authority over the nonlawyer shall make reasonable efforts to ensure that the person's conduct is compatible with the professional obligations of the lawyer; and

(c) a lawyer shall be responsible for conduct of such a person that would be a violation of the Rules of Professional Conduct if engaged in by a lawyer if:

(1) the lawyer orders or, with the knowledge of the specific conduct, ratifies the conduct involved; or

(2) the lawyer is a partner or has comparable managerial authority in the law firm in which the person is employed, or has direct supervisory authority over the person, and knows of the conduct at a time when its consequences can be avoided or mitigated but fails to take reasonable remedial action.

The Legal Ethics of Generative AI

Andrew M. Perlman*

I. Introduction

The legal profession is notoriously conservative when it comes to change.¹ From email to outsourcing,² lawyers have been slow to embrace new methods and quick to point out potential problems, especially ethics-related concerns.

The legal profession's approach to generative artificial intelligence (generative AI) is following a similar pattern. Many lawyers have readily identified the legal ethics issues associated with generative AI,¹ often citing the New York lawyer who cut and pasted fictitious citations from ChatGPT into a federal court filing.² Some judges have gone so far as to issue standing orders requiring lawyers to reveal when they use generative AI or to ban the use of most kinds of artificial intelligence (AI) outright.³ Bar associations are chiming in on the subject as well, though they have

* Dean & Professor of Law, Suffolk University Law School. I am grateful to multiple colleagues, including Sarah Boonin and Jeffrey Lipshaw, for their helpful suggestions on a draft of this essay. I also benefited in numerous ways from the work of research assistant Robert Massaro Stockard and the rest of the *Suffolk University Law Review* editorial staff.

¹ See generally RICHARD SUSSKIND, *TOMORROW'S LAWYERS: AN INTRODUCTION TO YOUR FUTURE* 1-15 (2d ed. 2017) (discussing the legal profession's slow adoption of new technologies).

² See, e.g., *ABA Comm. on Ethics and Prof'l Responsibility, Formal Op.* 99-413, at 11 n.40 (Mar. 10, 1999) (noting earlier ethics opinions that cautioned lawyers against the use of unencrypted email); *ABA COMM'N. ON ETHICS 20/20, Report on Resolution 105(c)*, at 2 (2012)

¹ LexisNexis, *Generative AI and the Legal Profession Survey Report* 8 (2023)

<https://www.lexisnexis.co.uk/pdf/generative-ai-and-the-legal-profession-report.pdf> (finding that 87% of surveyed lawyers were significantly concerned about the ethical implication of generative AI); Matt Reynolds, *Majority of Lawyers Have no Immediate Plans to use Generative AI*, LexisNexis Survey Finds, *ABA J.* (Mar. 24, 2023) <https://www.abajournal.com/web/article/survey-finds-majority-of-lawyers-have-no-immediate-plans-to-use-generative-ai> [<https://perma.cc/PN7P-YM7Y>] (reporting that 60% of surveyed lawyers had no plans to use generative AI at that time).

² *Mata v. Avianca*, No. 22-cv-1461, 2023 U.S. Dist. LEXIS 108263, at *3 (S.D.N.Y. June 22, 2023) (sanctioning lawyers for filing "false and misleading statements to the Court").

³ See Sara Merken, *Another US Judge Says Lawyers Must Disclose AI Use*, *REUTERS* (Feb. 24, 2023), <https://www.reuters.com/legal/another-us-judge-says-lawyers-must-disclose-ai-use-2023-02-24/> [<https://perma.cc/7Q2X-TS75?type=standar>] (comparing standing orders issued by Judge Stephen Vaden and U.S. District Judge Brantley Starr); Cedra Mayfield, *Judicial Crackdown: 'This Is Why I Have a Standing Order on the Use of AI'*, *ALM LAW.COM* (July 27, 2023), <https://www.law.com/2023/07/27/judicial-crackdown-this-is-why-i-have-a-standing-order-on-the-use-of-ai/> [<https://perma.cc/325M-AJSA>] (discussing generative AI standing orders issued by federal judges in four states); *infra* note 66 (listing standing orders on generative AI).

https://www.americanbar.org/content/dam/aba/administrative/ethics_2020/2012_hod_annual_meeting_105c_filed_may_2012.pdf (last visited Feb. 19, 2024) (acknowledging that the Commission’s proposals regarding outsourcing were controversial).
(so far) taken an admirably open-minded approach to the subject.⁶

Part II of this essay explains why the Model Rules of Professional Conduct (Model Rules) do not pose a regulatory barrier to lawyers’ careful use of generative AI, just as the Model Rules did not ultimately prevent lawyers from adopting many now-ubiquitous technologies.⁴ Drawing on my experience as the Chief Reporter of the ABA Commission on Ethics 20/20 (Ethics 20/20 Commission), which updated the Model Rules to address changes in technology, I explain how lawyers can use generative AI while satisfying their ethical obligations.⁵ Although this essay does not cover every possible ethics issue that can arise or all of generative AI’s law-related use cases, the overarching point is that lawyers can use these tools in many contexts if they employ appropriate safeguards and procedures.⁶

Part III describes some recent judicial standing orders on the subject and explains why they are ill-advised.⁷

The essay closes in Part IV with a potentially provocative claim: the careful use of generative AI is not only consistent with lawyers’ ethical duties, but the duty of competence may eventually *require* lawyers’ use of generative AI.⁸ The technology is likely to become so important to the delivery of legal services that lawyers who fail to use it will be considered as incompetent as lawyers today who do not know how to use computers, email, or online legal research tools.

II. Model Rules Implicated by Lawyers’ Use of Generative AI

Generative AI refers to technologies “that can generate high-quality text, images, and other content based on the data they were trained on.”⁹ The tools have the potential to reshape law practice,¹⁰ but lawyers necessarily need to consider a number of ethics-related issues. Although the list below is not comprehensive, the

⁶ See, e.g. FL. Eth. Op. 24-1, 2024 WL 271230, at *1 (Fla. State Bar Ass’n. Jan 19., 2024) (identifying some of the ethical issues that lawyers need to address when using generative AI). Cal. State Bar Standing Comm. On Pro. Responsibility and Conduct, *Practical Guidance for the Use of Generative*

⁴ See *infra* note 77 and accompanying text (discussing adoption of email).

⁵ See N.J. COURTS, *supra* note 6, at 3-4 (making similar observation).

⁶ See *infra* Part II (describing implicated Model Rules).

⁷ See *infra* Part III (focusing on current standing orders).

⁸ See *infra* Part IV (making the case for vision of the future).

⁹ Kim Martineau, *What is Generative AI* (Apr. 20, 2023), <https://research.ibm.com/blog/what-isgenerative-AI> (last visited Feb. 22, 2024).

¹⁰ Andrew Perlman, *The Implications of ChatGPT for Legal Services and Society*, 30 MICH. TELECOMM. & TECH. L. REV. (forthcoming 2024).

Artificial Intelligence in the Practice of Law, STATE BAR OF CAL. 1, 1 (Nov. 16, 2023), <https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf> [<https://perma.cc/B3X4-FAEC>] (same); N.J. COURTS, NOTICE TO THE BAR LEGAL PRACTICE: PRELIMINARY GUIDELINES ON THE USE OF ARTIFICIAL INTELLIGENCE BY NEW JERSEY LAWYERS 1-2 (2024), <https://www.njcourts.gov/sites/default/files/notices/2024/01/n240125a.pdf> [<https://perma.cc/LK7V-KY2R>] (same).

primary takeaway is that the Model Rules offer a useful roadmap for the ethical use of generative AI.

A. The Duty of Confidentiality Under Model Rule 1.6

Lawyers have to address several confidentiality issues when inputting or uploading client-related information into a generative AI tool. These issues, however, are not especially novel.¹⁴ For many years, lawyers have faced conceptually similar situations when using third-party, cloud-based technology, such as online document storage systems (e.g., Microsoft OneDrive or Dropbox) and email services (e.g., Gmail).¹⁵ Lawyers have also had to navigate confidentiality issues when inputting information into third-party tools, such as when querying online legal research tools like Westlaw and Lexis. Just as lawyers can adopt appropriate safeguards when using these kinds of services, they can do so when using generative AI.

The Ethics 20/20 Commission proposed amendments to the Model Rules in order to help lawyers address these kinds of confidentiality concerns.¹⁶ Model Rule 1.6(c), which was added in 2012, explains that “[a] lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.”¹¹ Comment 18 then refers lawyers to Model Rule 5.3, Comments 3-4 for guidance on how to comply with the duty when sharing information with third-parties outside the lawyer’s firm.¹²

Rule 5.3, Comment 3 is especially instructive. It counsels a lawyer to make “reasonable efforts to ensure” that outside service providers act in ways that are compatible with the lawyer’s professional obligations.¹³ The scope of this obligation

¹⁴ See Fla. Bar Standing Comm. on Pro. Ethics, *supra* note 6, at *1 (reaching a similar conclusion).

¹¹ See MODEL RULES OF PRO. CONDUCT r 1.6(c) (AM. BAR ASS’N 2020).

¹² See MODEL RULES OF PRO. CONDUCT r 1.6(c) cmt. [18] (AM. BAR ASS’N 2020) (referring readers to Model Rule 5.3, Comments 3-4); MODEL RULES OF PRO. CONDUCT r 5.3 cmt. [3]-[4] (AM. BAR ASS’N 2020) (commenting on how lawyers should obtain client consent before using third party nonlawyers).

¹³ See MODEL RULES OF PRO. CONDUCT r 5.3 cmt. [3] (AM. BAR ASS’N 2020) (asserting standard). The Comment provides as follows:

When using ... services outside the firm, a lawyer must make reasonable efforts to ensure that the services are provided in a manner that is compatible with the lawyer’s professional obligations. The extent of this obligation will depend upon the circumstances, including the education, experience and reputation of the nonlawyer;

¹⁵ See *id.*; Andrew C. Budzinski, *Clinics, the Cloud, and Protecting Client Data in the Age of Remote Lawyering*, 29 CLINICAL L. REV. 201, 201-03 (2023) (weighing cloud storage and professional responsibility considerations). Because most client data is now electronic, “the ethical lawyer must protect that data under their duty of confidentiality, to safeguard client property, and to protect the attorney-client privilege and work-product doctrine.” See *id.* at 202-03.

¹⁶ See ABA COMM’N. ON ETHICS 20/20 (2012)

https://www.americanbar.org/groups/professional_responsibility/committees_commissions/abacommission-on-ethics-20-20/ (last visited Feb. 19, 2024) (offering background and updates on Commission activities).

varies depending on the nature of the services involved, the terms of any arrangements concerning client information, and the “legal and ethical environments of the jurisdictions where the services are performed.”¹⁴ Put simply, lawyers can satisfy their confidentiality obligations when using generative AI tools (i.e., a “service outside the firm”) as long as they “make reasonable efforts to ensure that the services are provided in a manner that is compatible with the lawyer’s professional obligations.”¹⁵

This prescription means that, in the absence of informed client consent, lawyers should not insert or upload confidential information into most publicly available versions of generative AI services (like ChatGPT) because the companies operating those services typically have the right to review the prompts that are used.¹⁶ The companies also can train their models on any information that a lawyer shares.¹⁷

In contrast, lawyers can satisfy their duty of confidentiality when using thirdparty generative AI tools by making reasonable efforts to ensure that the third parties do not access the prompts or train their models from those prompts. For example, OpenAI has a version of ChatGPT (ChatGPT Enterprise) that includes data protection procedures that likely satisfy a lawyer’s duty of confidentiality.¹⁸ In that case, the use of generative AI would be analogous to a lawyer’s use of Microsoft OneDrive or a query on Westlaw or Lexis.

¹⁴ *Id.* (describing multiple factors).

¹⁵ *Id.*

¹⁶ See David Canellos, *What to Know About Sharing Company Data with Generative AI*, FORBES (Aug. 10, 2023), <https://www.forbes.com/sites/forbestechcouncil/2023/08/10/what-to-know-about-sharing-company-data-with-generative-ai/?sh=1ec0fff60229> [<https://perma.cc/DZV5-DA2>] (describing the dangers of using generative AI, including data leakage and exposing personally identifiable information); Michael Schade, *How Your Data is Used to Improve Model Performance*, OPENAI (2023), <https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-modelperformance> (last visited Feb. 19, 2024) (explaining how the company uses consumer data). With regard to Open AI’s Enterprise service, authorized employees are permitted to view stored inputs and outputs as are “specialized third-party contractors who are bound by confidentiality and security obligations.” See OpenAI, *API Platform FAQ*, <https://openai.com/enterprise-privacy> [<https://perma.cc/Y8VZ-KQWW>] (describing OpenAI’s policies regarding enterprise data).

¹⁷ See Schade, *supra* note 22 (describing OpenAI training policies).

¹⁸ See OpenAI, *supra* note 22 (highlighting ChatGPT Enterprise data protection procedures).

Other factors that lawyers need to consider include the reputation and location of the provider. For example, lawyers should be more wary of using a generative AI tool owned and operated in China versus one owned and operated in the United

the nature of the services involved; the terms of any arrangements concerning the protection of client information; and the legal and ethical environments of the jurisdictions in which the services will be performed, particularly with regard to confidentiality.

Id.
States.

In the absence of purchasing an instance of a third-party tool with appropriate privacy protections in place, lawyers have three other options for satisfying their confidentiality obligations. First, they could use the tools without uploading or sharing client confidences. Generative AI can be quite useful even without disclosing confidential information, just as legal research tools can be helpful without disclosing client confidences.

Second, lawyers could build their own generative AI tools. Although few law firms and legal departments currently have sufficient resources to do so on their own, the expense of deploying these tools internally may not be as expensive as many lawyers believe.¹⁹

A third option is for a lawyer to obtain a client's informed consent under Rule 1.6(a).²⁰ Rule 1.0(e) defines "informed consent" as "the agreement by a person to a proposed course of conduct after the lawyer has communicated adequate information and explanation about the material risks of and reasonably available alternatives to the proposed course of conduct."²¹ Rule 1.0 Comment 6 elaborates on the meaning of informed consent, but the essential idea is that the client must have sufficient information to make an informed decision, with lawyers having a greater obligation

¹⁹ See Robert J. Ambrogi, *Four Months After Launching Its 'Homegrown' GenAI Tool, Law Firm Gunderson Dettmer Reports on Results so far, New Features, and a Surprise on Cost*, LAW SITES (Dec. 20, 2023) <https://www.lawnext.com/2023/12/four-months-after-launching-its-homegrown-genai-toollaw-firm-gunderson-dettmer-reports-on-results-so-far-new-features-and-a-surprise-on-cost.html> [<https://perma.cc/6N35-GVD4>] (commenting on Gunderson Dettmer's recent launch of "ChatGD"). Gunderson's Chief Innovation Officer projects that the total annual cost for providing ChatGD to the entire firm "will be less than \$10,000." See *id.*

²⁰ See MODEL RULES OF PRO. CONDUCT r 1.6(a) (AM. BAR ASS'N 2020) (providing that "[a] lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorized in order to carry out the representation or the disclosure is permitted by paragraph (b)").

²¹ See MODEL RULES OF PRO. CONDUCT r 1.0(e) (AM. BAR ASS'N 2020) (defining informed consent).

²⁸ See MODEL RULES OF PRO. CONDUCT r 1.0 cmt. [6] (AM. BAR ASS'N 2020) (elaborating on the definition of informed consent).

to disclose information to unsophisticated clients than to those who are experienced regarding the conduct for which consent is sought.²⁸ For example, before sharing confidential information with a generative AI tool, a lawyer would have to explain the implications of doing so in more detail to the typical client than to the executive of an AI company. That said, given the current lack of technological sophistication of most lawyers and clients, it may not be possible in some instances to obtain informed consent to share sensitive information with many generative AI tools.

In sum, lawyers can comply with their duty of confidentiality when using generative AI tools either by not sharing confidential information (e.g., by prompting the tool with generic information) or by using tools owned and controlled by

companies that have appropriate terms and conditions on how the information can be used. An increasing number of well-established, reputable companies that have long served the legal industry are already launching generative AI tools in an attempt to satisfy these requirements.²⁹ Building a proprietary service is another option that is likely to become increasingly cost effective, and informed consent offers yet another possibility depending on the sophistication of the lawyer and the client.

B. Consulting with Clients Under Model Rule 1.4

Rule 1.4 imposes a number of duties on lawyers to keep clients informed about a pending matter.³⁰ As applied to generative AI, the most relevant portion may be Rule 1.4(a)(2). It explains that “a lawyer shall reasonably consult with the client about the means by which the client's objectives are to be accomplished.”³¹ Comment [3] elaborates on the duty this way:

In some situations—depending on both the importance of the action under consideration and the feasibility of consulting with the client—this duty will require consultation prior to taking action. In other circumstances, such as during a trial when an immediate decision must be made, the exigency of the situation may require the lawyer to act without prior consultation. In such cases the lawyer must nonetheless act reasonably to inform the client of actions

²⁹ See LexisNexis, *LexisNexis Launches Lexis+ AI, a Generative AI Solution with Linked Hallucination-Free Legal Citations*, LEXISNEXIS (Nov. 14, 2023) <https://www.lexisnexis.com/community/pressroom/b/news/posts/lexisnexis-launches-lexis-ai-generative-ai-solution-with-hallucination-free-linked-legal-citations> [https://perma.cc/T82P-R2QY] (explaining development and capabilities of Lexis+ AI); Thomson Reuters, *Thomson Reuters*

Launches Generative AI-Powered Solutions to Transform how Legal Professionals Work, THOMSON REUTERS (Nov. 15, 2023) <https://www.thomsonreuters.com/en/pressreleases/2023/november/thomson-reuters-launches-generative-ai-powered-solutions-to-transform-how-legal-professionals-work.html> [https://perma.cc/KS42-BY4Y] (debuting AI-Assisted Research on Westlaw Precision).

³⁰ See MODEL RULES OF PRO. CONDUCT r 1.4 (AM. BAR ASS'N 2020). Rule 1.4 provides as follows:

- (a) A lawyer shall:
 - (1) promptly inform the client of any decision or circumstance with respect to which the client's informed consent, as defined in Rule 1.0(e), is required by these Rules;
 - (2) reasonably consult with the client about the means by which the client's objectives are to be accomplished;
 - (3) keep the client reasonably informed about the status of the matter;
 - (4) promptly comply with reasonable requests for information; and (5) consult with the client about any relevant limitation on the lawyer's conduct when the lawyer knows that the client expects assistance not permitted by the Rules of Professional Conduct or other law.
- (b) A lawyer shall explain a matter to the extent reasonably necessary to permit the client to make informed decisions regarding the representation. *See id.*

³¹ See MODEL RULES OF PRO. CONDUCT r 1.4(a)(2) (AM. BAR ASS'N 2020) (explaining that lawyers must reasonably consult with their clients to accomplish clients' objectives).
the lawyer has taken on the client's behalf.²²

Because the use of generative AI can be viewed as one of the “means to be used to accomplish the client’s objectives,” Rule 1.4(a)(2) arguably imposes on a lawyer the duty to consult with a client before using such services.²³ Thus, even if a lawyer can overcome the confidentiality issues described earlier—such as by deploying a tool within the law firm that contains appropriate privacy protections—a lawyer may still have to inform the client about the tool’s use in the client’s matter. Indeed, some lawyers have begun to inform clients about these uses in their engagement letters.²⁴

Such a consultation is only arguable because it is not entirely clear that a lawyer’s use of generative AI is sufficiently important to warrant a consultation in all circumstances. For example, lawyers already take advantage of some basic forms of generative AI without even realizing it—such as when they use the autocomplete feature in Microsoft Word—and lawyers should not need to consult clients before using such tools.²⁵

²² MODEL RULES OF PRO. CONDUCT r. 1.4 cmt. [3] (AM. BAR ASS'N.).

²³ See MODEL RULES OF PRO. CONDUCT r. 1.4(a)(2) (AM. BAR ASS'N.) (requiring lawyers to “reasonably consult” with their client about the means used to accomplish a client’s objectives).

²⁴ See Isabel Gottlieb, *Law Firms Wrestle with How Much to Tell Clients About AI Use*, BLOOMBERG LAW (Nov. 29, 2023) <https://news.bloomberglaw.com/business-and-practice/law-firms-wrestle-with-how-much-to-tell-clients-about-ai-use> [<https://perma.cc/YBN6-MQUE>] (asking numerous firms about how they disclose the use of generative AI to their clients).

²⁵ See generally, Andrea Eoanou, *Introducing New AI Enhancements in Microsoft 365: New Features Coming to Microsoft Editor and More!*, MICROSOFT (Oct. 12, 2022) <https://techcommunity.microsoft.com/t5/microsoft-365-blog/introducing-new-ai-enhancements-inmicrosoft-365-new-features/ba-p/3643499> [<https://perma.cc/7R84-U5B2>] (describing new autocomplete features in Outlook and Word); Microsoft, *Welcome to Copilot in Word*, MICROSOFT, <https://support.microsoft.com/en-us/office/welcome-to-copilot-in-word-2135e85f-a467-463b-b2f0c51a46d625d1> [<https://perma.cc/4QMA-JQCV>] (announcing how Word customers can use Copilot AI to draft documents).

Even when lawyers use more sophisticated forms of generative AI (e.g., using it to draft a legal memo), it is not obvious that a lawyer should have to consult with the client before doing so.²⁶ Assuming the lawyer is appropriately protecting client confidences and carefully reviewing the outputs, one could conclude that lawyers should have no greater obligation to consult with clients before using generative AI than before using online legal research tools, querying Google, or storing client documents on a network drive.

That said, at least for now, lawyers are well-advised to consult with clients before using generative AI to assist with anything other than the *de minimis* case of autocompleting simple text. Consultation aligns with the principle of transparency that underlies Rule 1.4 and aids in managing client expectations about the nature and source of the legal services provided.³⁷ Given the novelty and evolving nature of

generative AI, clients may not be fully aware of its capabilities and limitations, so for the time being, lawyers should typically consult with clients before using generative AI in more substantive ways.

That said, this duty may evolve considerably in the future. Even if a duty of consultation currently exists under Rule 1.4, generative AI tools are likely to become so ubiquitous in the years to come that consultation is likely to become unnecessary. In the meantime, however, such a consultation is highly advisable for anything other than the most basic of drafting tasks.

C. Oversight of Nonlawyer Services Under Model Rule 5.3

In 2012, the Ethics 20/20 Commission proposed a two-letter change to the title of Rule 5.3 from “Responsibilities Regarding Nonlawyer Assistants” to “Responsibilities Regarding Nonlawyer Assistance.”²⁷ The change signaled that lawyers use an increasingly wide range of non-human forms of assistance when representing clients and should consider several factors when using those services.²⁸ The Ethics 20/20 Commission also proposed (and the ABA adopted) several new Comments that were designed to guide lawyers with regard to the use of such thirdparty services.²⁹

²⁶ See N.J. Guidance, *supra* note 6, at 4-5 (reaching a similar conclusion). ³⁷

See MODEL RULES OF PRO. CONDUCT r 1.4 (AM. BAR ASS’N 2020).

²⁷ See ABA COMM’N. ON ETHICS 20/20, RES. 105A REVISED, REPORT TO THE HOUSE OF DELEGATES 2 (2012) (describing change from “Assistants” to “Assistance”) [hereinafter RES. 105A REVISED]; MODEL RULES OF PRO. CONDUCT r 5.3 (AM. BAR ASS’N 2020) (stating modified title).

²⁸ See ABA COMM’N ON ETHICS 20/20, RES. 105C, REPORT TO THE HOUSE OF DELEGATES 2 (2012) [hereinafter RES. 105C] (introducing change to Rule 5.3).

²⁹ See ABA Commission on Ethics 20/20, AM. BAR ASS’N.

As discussed earlier in the context of the duty of confidentiality, Comment 3 is especially helpful in understanding how Rule 5.3 applies to a lawyer's use of generative AI.³⁰ The Comment has implications well beyond issues of confidentiality and suggests that lawyers who use third-party services must make reasonable efforts to ensure that those services are performed in a manner that is consistent with the lawyer's own obligations.³¹ The extent of the lawyer's obligation will necessarily turn on the "education, experience and reputation of the nonlawyer; the nature of the services involved; the terms of any arrangements concerning the protection of client information; and the legal and ethical environments of the jurisdictions in which the services will be performed, particularly with regard to confidentiality."³²

These factors suggest that lawyers will have varying duties of oversight

depending on the nature of the generative AI service that they use. For example, if a lawyer is simply using Microsoft's autocomplete feature, the lawyer would not have an obligation to take any particular action. The feature typically inserts only a few words at the end of a sentence, making it easy for a lawyer to determine the reasonableness of the suggested wording and to either accept, reject, or modify it. The "nature of the service involved" in this example is modest and should not require a lawyer to take any additional steps under Rule 5.3.³³

In contrast, if a lawyer uses more sophisticated forms of generative AI, there will be additional oversight obligations. Among other considerations, the lawyer would have to understand the "education, experience, and reputation" of the generative AI before using it.³⁴ For example, a lawyer might look into how the generative AI service was trained and what procedures are used to ensure the accuracy of outputs. The lawyer might also investigate the reputation of the tool by reviewing the increasing number of studies that document how reliable various generative AI services are (i.e., the extent to which the tool "hallucinates").³⁵ A lawyer can have more confidence when using a generative AI tool that has a reputation for accuracy in the context of legal services than when using a tool that does not have any indicators of reliability. Moreover, as the Comment suggests and as discussed

https://www.americanbar.org/groups/professional_responsibility/committees_commissions/abacommission-on-ethics-20-20/ (last visited Feb. 19, 2024) (describing all accepted and proposed changes to Model Rules).

³⁰ See MODEL RULES OF PRO. CONDUCT r. 5.3 cmt. [3] (AM. BAR ASS'N. 2020) (explaining how to use nonlawyer assistance outside firm).

³¹ See *id.* (noting how lawyers must make reasonable efforts to ensure nonlawyer compliance with Model Rule 5.3).

³² See *id.* (describing standard of Model Rule 5.3, Comment 3).

³³ See *id.* (tying lawyer's disclosure obligations to the nature of the services involved).

³⁴ See MODEL RULES OF PRO. CONDUCT r. 5.3 cmt. [3] (AM. BAR ASS'N. 2020).

³⁵ See IBM, *What are AI Hallucinations?*, IBM, <https://www.ibm.com/topics/ai-hallucinations> [<https://perma.cc/WMD4-GU6P>] (explaining what leads to generative AI hallucinations).

earlier, the lawyer will have to assess the confidentiality implication of using the generative AI service.

A lawyer might reasonably decide to use a generative AI tool after considering these factors, but the lawyer should still carefully review all AI-generated content for accuracy before relying on it. To be clear, the high likelihood of errors does not mean that Rule 5.3 prohibits lawyers from using the service. Rather, in much the same way that lawyers have to check the work of paralegals or inexperienced summer associates (who often make mistakes), lawyers will have to do the same when generating content through AI. A high probability of error does not mean a lawyer is prohibited from using a particular service; it just means that the lawyer must vet the content more carefully.

D. The Duty of Competence Under Rule 1.1

All of the preceding ethical obligations arguably fall under the more general obligation to act competently with regard to technology. Prior to the work of the Ethics 20/20 Commission, the word “technology” did not even appear in the Model Rules, so the Commission decided that the Model Rules should address the issue and that a comment related to the duty of competence was the appropriate place to do

so.⁴⁷

The new language (in italics) says that, “[t]o maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, *including the benefits and risks associated with relevant technology...*”⁴⁸ The idea here is that, to maintain competence, lawyers necessarily need to remain aware of both the benefits and the risks associated with existing and emerging technologies.

In the context of generative AI, this obligation means that lawyers should understand the potential advantages and risks from the tools.⁴⁹ Lawyers can quite reasonably conclude that, under some circumstances, generative AI does not present a sufficient benefit to outweigh the risks and vice versa. This assessment is a necessary part of a lawyer’s ongoing duty of competence.³⁶

In sum, lawyers have to navigate a number of ethical issues when using generative AI, including some not even referenced here. For example, lawyers may have to deal with issues involving the unauthorized practice of law, duties to prospective clients under Rule 1.18 (e.g., when generative AI is used to interact with potential clients) and duties related to fees under Rule 1.5 (e.g., how lawyers charge

³⁶ See generally *id.* (finding 66% of surveyed attorneys believe that the use of AI does not violate ABA Model Rules).

for their time when using generative AI and the prohibition against lawyers billing for time that they did not spend on a matter).³⁷ Moreover, the legal profession is likely to face other ethics-related issues going forward, such as whether to have mandatory training on generative AI for both law students and practicing lawyers, as the California Committee on Professional Responsibility and Conduct recently suggested.⁵² The overarching point, however, is that the ethics rules will not impede the steady advance of generative AI in the delivery of legal services.

⁴⁷ See ABA COMM'N ON ETHICS 20/20, *supra* note 40 (proposing changes to the Comments to Model Rule 1.1). See *infra* note 48.

⁴⁸ See MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. [8].

⁴⁹ See Jessica R. Blaemire, *Analysis: Lawyers Recognize Ethical Duty to Understand Gen AI*, BL ANALYSIS (Oct. 19, 2023) <https://news.bloomberglaw.com/bloomberg-law-analysis/analysis-lawyersrecognize-ethical-duty-to-understand-gen-ai> [<https://perma.cc/542A-T2LR>] (explaining results of study). In fact, many attorneys have already concluded that they can use generative AI in their practice without violating an ethical duty. See *id.* For example, Bloomberg Law asked 452 attorneys for their opinion on legal ethics and the use of generative AI and “almost 70% said that it’s possible to use generative AI in legal practice without violating an ethical duty, and almost as many (66%) said it can be used without violating the ABA Model Rules or state equivalents.” See *id.* These results suggest that, while the Model Rules may not currently have provisions that directly address generative AI, the legal profession recognizes that the rules of professional conduct are unlikely to impede the legal profession’s adoption of generative AI. See *generally id.*

III. Obligations Imposed by Court Order

Some courts have responded to the emergence of generative AI by issuing standing orders that impose near-outright bans on lawyers’ use of AI or require lawyers to disclose when they have used the technology for court filings.³⁸ Both types of orders are overly broad and unnecessary.

A. The Problems with Banning AI

One example of a ban comes from Judge Michael J. Newman of the United States District Court for the Southern District of Ohio.³⁹ Judge Newman has a standing order that not only prohibits the use of generative AI tools to prepare a court

³⁷ See MODEL RULES OF PRO. CONDUCT r. 1.18 (AM. BAR ASS’N. 2020) (describing duties to prospective clients); MODEL RULES OF PRO. CONDUCT r. 1.5 (AM. BAR ASS’N. 2020) (explaining lawyer fee schedules and arrangements); Fla. Bar Standing Comm. on Pro. Ethics, *supra* note 6 (describing billing-related issues arising from lawyers’ use of generative AI). ⁵² See *infra* note 78 (recommending such training).

³⁸ See, e.g., J. Michael J. Newman, *Artificial Intelligence (“AI”) Provision in Both Civil and Criminal Cases* (S.D. Ohio July 14, 2023); J. Roy Ferguson, *Standing Order Regarding Use of Artificial Intelligence* (394th Jud. Dist. Tex, June 9, 2023); J. Stephen Alexander Vaden, *Order on Artificial Intelligence*, (U.S. Ct. Int’l. Trade, June 6, 2023).

³⁹ Newman, *supra* note 53.

filing but extends that prohibition to the use of nearly all forms of artificial intelligence.⁴⁰ The standing order provides as follows:

No attorney for a party, or a *pro se* party, may use Artificial Intelligence (“AI”) in the preparation of any filing submitted to the Court. Parties and their counsel who violate this AI ban may face sanctions including, *inter alia*, striking the pleading from the record, the imposition of economic sanctions or contempt, and dismissal of the lawsuit. The Court does not intend this AI ban to apply to information gathered from legal search engines, such as Westlaw or LexisNexis, or Internet search engines, such as Google or Bing. All parties and their counsel have a duty to immediately inform the Court if they discover the use of AI in any document filed in their case.⁵⁶

This ban is problematic for two reasons. First, by prohibiting the use of nearly all forms of AI—and not just generative AI—the order is dramatically overbroad. The definition of “artificial intelligence” varies, but it commonly “refers to the ability of machines and computers to perform tasks that would normally require human intelligence.”⁴¹ Using this definition, the order would prohibit lawyers from using most types of professional productivity software, such as Microsoft Word, Outlook, and Gmail, given that most of these tools perform tasks (like spellchecking and

grammar checking) that used to require human-level intelligence.⁵⁸ The order also would seem to extend to e-discovery services, which almost always rely on some form of AI.⁵⁹ Since those e-discovery services do not fall within the safe harbor of “legal search engines,” lawyers would presumably be prohibited from using them to find relevant information when preparing a court filing.

Not only is the court order overbroad, but it is also unnecessary. Lawyers are already subject to sanctions or discipline for filing inaccurate or false documents using AI.⁶⁰ For example Rule 11(b) of the Federal Rules of Civil Procedure (FRCP) requires lawyers to thoroughly research their pleadings, filings, or motions to a court using “an inquiry reasonable under the circumstances.”⁴² In other words, lawyers must

⁴⁰ *Id.*

⁵⁶ *Id.*

⁴¹ Jennifer Monahan, *Artificial Intelligence, Explained*, <https://www.heinz.cmu.edu/media/2023/July/artificial-intelligence-explained> (Jul. 2023) (last visited Feb. 22, 2024). *See also* Clara Pilato, *Artificial Intelligence vs Machine Learning: What's the difference?*, <https://professionalprograms.mit.edu/blog/technology/machine-learning-vs-artificialintelligence/> (last visited Feb. 23, 2024) (describing artificial intelligence as the ability of “computers to imitate cognitive human functions” and noting that “artificial intelligence is everywhere”).

⁴² *See* Fed. R. Civ. P. 11(b) (imposing obligations on lawyers when filing documents with the court). The Rule provides as follows:

certify that their filings do not contain fictitious legal contentions, citations, or claims.⁴³ Model Rule 3.1, which has been adopted in nearly every U.S. jurisdiction, imposes almost identical obligations.⁴⁴

These provisions were more than adequate to discipline and sanction the infamous New York lawyer who cut and pasted bogus citations from ChatGPT into a court document.⁴⁵ In fact, the judge in that case (Judge P. Kevin Castel) acknowledged “there is nothing inherently improper about using a reliable artificial

⁵⁸ John Roach, *How AI is making people’s workday more productive*, <https://news.microsoft.com/source/features/ai/microsoft-365-intelligent-workday-productivity/> (May 6, 2019) (explaining how artificial intelligence was infused in Microsoft products in 2019 through spellchecking and grammar checking).

⁵⁹ *See AI for Lawyers: How Law Firms are Leveraging AI for Document Review*, CASEPOINT, <https://www.casepoint.com/resources/spotlight/leveraging-ai-document-review-law-firms/> [<https://perma.cc/V3PP-WRPF>] (offering ways to use AI throughout the e-discovery process); *Casetext Launches AllSearch, Powerful Document Search Technology for Litigators*, CASETEXT (June 6, 2022) <https://casetext.com/blog/allsearch-launch/> [<https://perma.cc/XG2N-RWWH>] (promoting AllSearch’s ability to streamline e-discovery workflows).

⁶⁰ *See Mata v. Avianca*, No. 22-cv-1461, 2023 U.S. Dist. LEXIS 108263, at *3 (S.D.N.Y. June 22, 2023) (sanctioning attorney under FRCP 11 for submitting document with fictitious citations generated by ChatGPT).

intelligence tool for assistance.”⁴⁶ Judge Castel correctly recognized that an across-the-board ban is unnecessary because both the Model Rules and the Federal Rules of Civil Procedure provide sufficient protections against a lawyer’s careless use of AI.

B. The Overbreadth of Orders Requiring Disclosure

Some courts have adopted a more targeted approach by simply requiring lawyers to disclose when they have used generative AI to prepare a court filing.⁴⁷ For

By presenting to the court a pleading, written motion, or other paper—whether by signing, filing, submitting, or later advocating it—an attorney or unrepresented party certifies that to the best of the person’s knowledge, information, and belief, formed after an inquiry reasonable under the circumstances ... (2) the claims, defenses, and other legal contentions are warranted by existing law or by a nonfrivolous argument for extending, modifying, or reversing existing law or for establishing new law; [and] (3) the factual contentions have evidentiary support or, if specifically so identified, will likely have evidentiary support after a reasonable opportunity for further investigation or discovery . . . *See id.*

⁴³ *Id.*

⁴⁴ *Compare* MODEL RULES OF PRO. CONDUCT r. 3.1 (AM. BAR ASS’N. 2020) (describing a lawyer’s obligations with regard to meritorious claims & contentions), *with* Fed. R. Civ. P. 11(b) (outlining similar standards).

⁴⁵ *See Mata*, 2023 U.S. Dist. LEXIS 108263, at *45-46 (sanctioning attorney for false citations).

⁴⁶ *Id.* at *1 (noting the effective and ethical applications of AI in legal work).

⁴⁷ *See* Magis. J. Gabriel A. Fuentes, *Standing Order for Civil Cases Before Magistrate Judge Fuentes*, (N.D. Ill. May 5, 2023) (requiring any party to disclose the use of generative AI in courtfiled

example, U.S. Magistrate Judge Gabriel Fuentes of the United States District Court for the Northern District of Illinois has a standing order with the following directive: “[a]ny party using any generative AI tool to conduct legal research or to draft documents for filing with the Court must disclose in the filing that AI was used, with the disclosure including the specific AI tool and the manner in which it was used.”⁴⁸

The United States Court of Appeals for the Fifth Circuit similarly specifies that:

Counsel and unrepresented filers must ... certify that no generative artificial intelligence program was used in drafting the document presented for filing, or to the extent such a program was used, all generated text, including all citations and legal analysis, has been reviewed for accuracy and approved by a human.⁴⁹

Other courts have adopted conceptually similar approaches.⁵⁰

These directives are an improvement over Judge Newman’s order, but they are still overly broad.⁵¹ One problem is that lawyers are now using generative AI without even realizing it. Take, for example, this very essay, which was drafted using Microsoft Word 365. At various times while drafting the piece, Microsoft suggested ways to autocomplete a sentence (including while writing this sentence). These autocomplete features are a form of “generative AI,” and they are now incorporated into a wide range of professional software. Does a lawyer have to disclose to a court

each time a filed document may have had some words generated by commonly used tools? If courts only intend to require lawyers to disclose when they use AI to generate more substantive content, how much more substantive does it need to be? The lines are difficult to draw already, but they will become increasingly so as generative AI is incorporated more deeply and widely into professional tools.

Another problem with these orders is that they would require lawyers to disclose when they have used generative AI just to brainstorm ideas. The tools are

documents to court); J. Brantley Starr, Mandatory Certification Regarding Generative Artificial Intelligence, (N.D. Tex. May 30, 2023) (requiring all attorneys or *pro se* litigants to certify that generative AI did not draft any portion of filing).

⁴⁸ Fuentes, *supra* note 66.

⁴⁹ 5th Cir. R. 32.3 (proposed Amendment, Dec. 1, 2023) [hereinafter Fifth Circuit Standing Order].

⁵⁰ See, e.g., Fuentes, *supra* note 66; Starr, *supra* note 66; Vaden, *supra* note 53; Ferguson, *supra* note 53; J. Michael M. Baylson, Standing Order RE: Artificial Intelligence (“AI”) in Cases Assigned to Judge Baylson (E.D. Penn., June 6, 2023).

⁵¹ Compare Newman, *supra* note 53 (creating generative AI standing order), with 5th Cir. R. 32.3 (proposed Amendment, Dec. 1 2023) (allowing for use of generative AI with human oversight for accuracy), and Fuentes, *supra* note 66 (requiring attorneys or *pro se* litigants to disclose the use of generative AI, but not banning it).

often quite useful in helping to think through possible arguments or to suggest weaknesses in wording. There is no clear public policy rationale for why a lawyer should have to disclose such uses, but most of the standing orders effectively impose such a disclosure requirement.⁵²

The standing orders are not only worded too broadly, but like Judge Newman's order, they are unnecessary. As noted earlier, the rules of professional conduct and rules of civil procedure impose sufficient duties on lawyers with regard to their filings. A notification requirement will not only cause increasing confusion as generative AI tools become ubiquitous, but courts have ample tools to ensure that lawyers fulfill their ethical and legal duties to the court.⁷²

Judges have expressed their concerns about generative AI in a variety of ways, with Judge Brantley Starr of the United States District Court for the Northern District of Texas offering among the most elaborate explanations:

These platforms are incredibly powerful and have many uses in the law: form divorces, discovery requests, suggested errors in documents, anticipated questions at oral argument. But legal briefing is not one of them. Here's why. These platforms in their current states are prone to hallucinations and bias. On hallucinations, they make stuff up—even quotes and citations. Another issue is reliability or bias. While attorneys swear an oath to set aside their personal prejudices, biases, and beliefs to faithfully uphold the law and represent their clients, generative artificial intelligence is the product of programming devised by humans who did not have to swear such an oath. As such, these systems hold no allegiance to any client, the rule of law, or the laws and Constitution of the United States (or, as addressed above, the truth). Unbound by any sense of duty, honor, or justice, such programs act according to computer code rather than conviction, based on programming rather than

principle. Any party believing a platform has the requisite accuracy and reliability for legal briefing may move for leave and explain why. Accordingly, the Court will strike any filing from a party who fails to file a certificate on the docket attesting that they have read the Court's judge-specific requirements and understand that they will be held responsible under Rule 11 for the

⁵² See Fuentes, *supra* note 66 (requiring any party that uses generative AI in research or drafting documents to disclose its use); Vaden, *supra* note 53 (mandating disclosure of use of generative AI in any submission to Judge Vaden); Baylson, *supra* note 69 (requiring any attorney or *pro se* litigant to disclose generative AI use in any submitted filing); see also Maura R. Grossman et al., *Is Disclosure and Certification of the Use of Generative AI Really Necessary?*, 107 JUDICATURE 69, 76 (2023) (arguing that current standing orders with disclosure requirements unnecessarily burden litigants).

⁷² See, e.g., *Mata v. Avianca*, No. 22-cv-1461, 2023 U.S. Dist. LEXIS 108263, at *45-46 (S.D.N.Y. June 22, 2023) (using existing provisions to impose sanctions).

contents of any filing that they sign and submit to the Court, regardless of whether generative artificial intelligence drafted any portion of that filing. A template Certificate Regarding Judge-Specific Requirements is provided here.⁵³

The problem with this reasoning is that it proves too much. Lawyers have long used a variety of methods to prepare court filings that trigger conceptually similar concerns, yet courts do not impose any new certification obligations. Consider, for example, lawyers who use summer associates to help prepare the first draft of a court filing, including a brief. The summer associate is much more likely to make mistakes than a lawyer (i.e., summer associates do not have “requisite accuracy and reliability for legal briefing”), but despite this risk of error, courts do not require lawyers to separately certify that have adequately supervised summer associates who worked on the filing. Lawyers understand their obligations to provide appropriate oversight and review before filing a document with a court. That obligation is sufficient in the context of summer associates, and it is sufficient with regard to generative AI.

Having said that, there is arguably no downside to courts reminding lawyers to comply with their existing ethical and legal obligations when using generative AI, especially given the nascent nature of the technology. Most of the existing orders, however, go beyond such a reminder. They institute notification requirements or outright bans, which cause increasing confusion and impose unnecessary new obligations as these tools become more widespread. For now, the best approach is for courts to rely on their existing ability to sanction lawyers or to simply remind lawyers that they should be careful when using generative AI.

IV. The Future of the Duty of Competence

The contention of this essay so far has been fairly modest and can be summarized by two basic points. First, lawyers can typically use generative AI in ethically compliant ways by adopting appropriate procedures and protocols. Second, judicial efforts to prohibit these tools or impose notification requirements are either problematic or unnecessary.

The final section of this essay makes an even more provocative claim: generative AI is advancing so rapidly that we may eventually move away from saying that lawyers are ethically permitted to use it, to saying that lawyers are ethically *required* to do so. The idea here is that, just as we would question the competence of a lawyer who pulls out a typewriter to prepare a client document, we will at some point question the competence of a lawyer who begins drafting legal documents by opening a word processing program to a blank screen and typing from scratch.

⁵³ 73 Starr, *supra* note 66.

Lawyers will be expected to use generative AI tools—or whatever they will be called in the future—as part of the modern, competent practice of law.

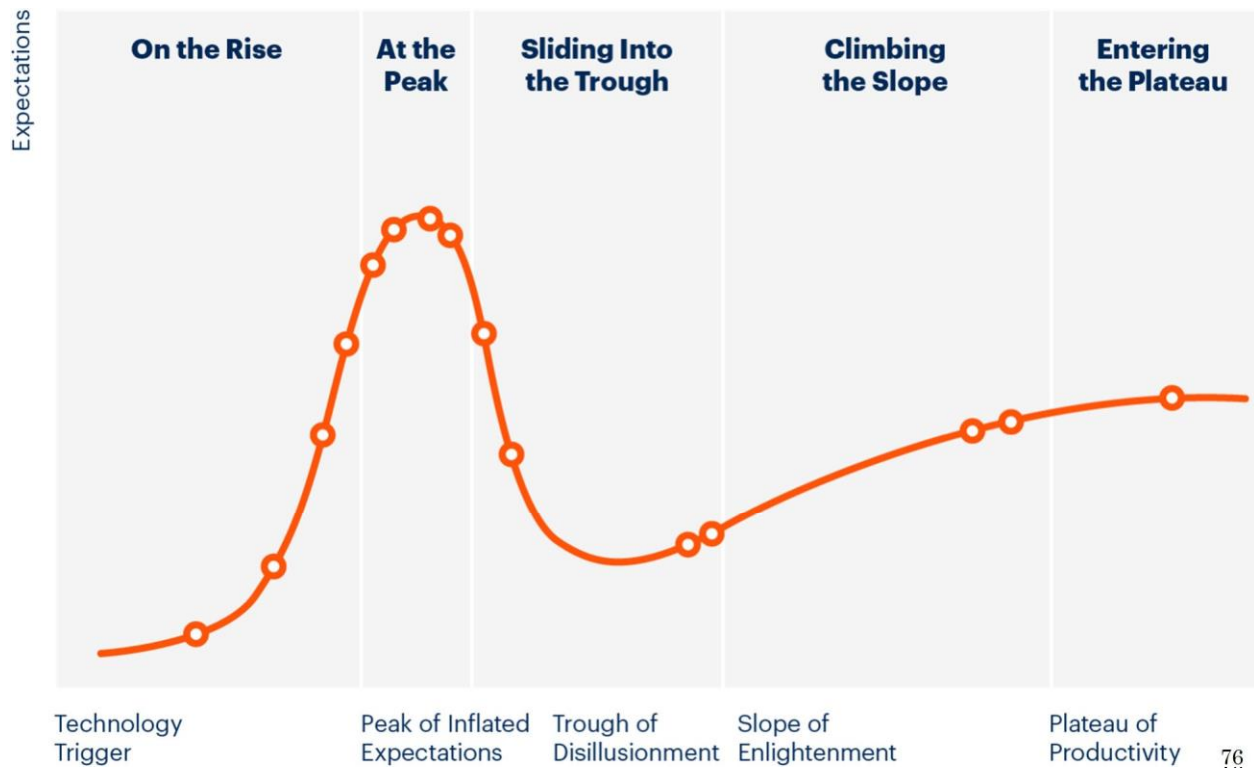
Lawyers already have begun to use these tools to improve the quality of their work or make it more efficient. For example, generative AI tools are helping lawyers draft clauses and phrases in transactional documents; summarize large collections of documents in litigation and transactional work; draft and respond to emails; brainstorm possible arguments to raise in litigation or identify weaknesses in existing arguments; draft interrogatories and document requests; draft simple transactional documents; prepare first drafts of simple motions and briefs; identify inconsistencies in deposition and trial testimony in real time; prepare first drafts of legal memos; and identify possible deposition topics and questions.⁵⁴ These use cases have emerged within only one year of ChatGPT’s release, when these tools are in their relative infancy. The level of sophistication is likely to grow significantly in the future, making these tools indispensable to modern law practice.

Is this transition likely to happen soon? The answer is almost certainly, “no.” As Bill Gates once said, “People often overestimate what will happen in the next two years and underestimate what will happen in ten.”⁵⁵ Generative AI’s potential to transform the legal profession is enormous, but it will not lead to seismic changes in the immediate future. The tools are evolving; their reliability is still improving; and the use cases are still emerging. Law firms, legal departments, and legal services providers are understandably cautious about deploying these tools, and they are waiting to see how the market evolves in the coming years.

Put another way, generative AI is going through some version of the so-called Gartner hype cycle, where we expect a new technology to be more transformative than we can reasonably expect it to be in the short term. We may soon enter the “trough of disillusionment” if we are not there already.

⁵⁴ See, e.g., Patrick Smith, *Sullivan & Cromwell’s Investments in AI Lead to Discovery, Deposition Assistants*, ALM LAW.COM (Aug. 21, 2023), <https://www.law.com/americanlawyer/2023/08/21/sullivan-cromwell-investments-in-ai-lead-to-discovery-deposition-assistants/> [<https://perma.cc/TUX4-UK2L>] (describing current and future uses of generative AI at Sullivan & Cromwell); *How To . . . Use AI to Ace Your Next Deposition*, CASETEXT (Aug. 31, 2023), <https://casetext.com/blog/4-steps-to-acing-your-next-deposition-using-ai/> [<https://perma.cc/TY3D-38X3>] (explaining how AI helps litigators efficiently and effectively prepare for depositions).

⁵⁵ BILL GATES ET AL., *THE ROAD AHEAD* 316 (2d. ed. 2023).



That said, generative AI will very likely become ubiquitous in much the same way as email and online legal research. Competent lawyers are now expected to know how to use those tools, and the same will eventually be true for generative AI (i.e., the technology will reach the right side of the curve, but perhaps with a steeper upward slope).

The email analogy may be especially apt. When the technology first became available, ethics opinions urged considerable caution and even suggested that lawyers might violate their duty of confidentiality by using it.⁷⁷ We have now reached the point where lawyers *must* have an email address in order to remain licensed to practice law.⁵⁶ We are likely to see a similar transition for generative AI, as we move

⁷⁶ See *Decide Which Technologies Are Crucial to Future Proof Your Business*, GARTNER, <https://www.gartner.com/en/marketing/research/hype-cycle> [<https://perma.cc/EQQ5-G9PF>] (explaining and illustrating Gartner hype cycle).

⁷⁷ See Laurel S. Terry, *30th Anniversary Commemorative Issue: Commemorative Contributions: The Impact of Global Developments on U.S. Legal Ethics During the Past Thirty Years*, 30 GEO. J. LEGAL ETHICS 365, 372 (2017) (explaining the history behind the legal profession's treatment of email); ABA Comm. on Ethics & Pro. Resp., Formal Op. 99-413 (1999) (concluding that lawyers can use email and

⁵⁶ See *Attorneys Must Provide E-mail Address to the Bar by Feb. 1*, STATE BAR OF CAL., <https://www.calbarjournal.com/January2010/TopHeadlines/TH3.aspx> [<https://perma.cc/TUA6-2NPQ>] (announcing change to Rule 9.7 and requiring attorneys to provide e-mail addresses); *Service: It's the*

fulfill their ethical obligations under Rule 1.6); ABA Comm. on Ethics & Pro. Resp., Formal Op. 477 (2017) (concluding that lawyers may transmit information about their client over the internet without violating the Model Rules).
from urging caution to expecting usage.

V. Conclusion

The Model Rules offer an adaptable framework for guiding lawyers on their use of generative AI. This adaptability is by design. When the Ethics 20/20 Commission proposed amendments to the Model Rules more than a decade ago, it understood that the amendments needed to offer sufficiently flexibility to accommodate future technological developments.⁷⁹

This flexible approach implies that we can expect the assessment of generative AI to evolve in the future as the tools become more reliable and useful. At some point, generative AI is likely to become so critical to the effective and efficient delivery of legal services that lawyers will have an ethical obligation to use it. We may even come to see generative AI as an important way to serve the public's unmet legal needs and as a powerful tool for addressing the access-to-justice crisis.⁸⁰

The first sentence of the preamble to the Model Rules says that “[a] lawyer, as a member of the legal profession, is a representative of clients, an officer of the legal system and a public citizen having special responsibility for the quality of justice.”⁵⁷ If we take this obligation seriously, we necessarily need to consider how new technologies can help us to better serve our clients and the public. Generative AI is such a technology and may have more potential in this regard than any technology ever invented.

Law, ILL. STATE BAR ASS'N (Sept. 17, 2017), <https://www.isba.org/barnews/2017/09/27/email-serviceit-s-law> [<https://perma.cc/7WG7-2Y5R>] (explaining recent update to Illinois Supreme Court Rule 11); *Annual Regulatory Compliance*, VA. STATE BAR, <https://vsb.org/Site/Site/lawyers/compliance.aspx> [<https://perma.cc/6S8F-AXKZ>] (mandating all attorneys to keep an “email of record” to maintain their license).

⁷⁹ See Letter from ABA Comm'n. on Ethics 20/20 Working Group, to ABA Entities, Courts, Bar Associations (state, local, specialty and international), Law Schools, Individuals, and Entities (Sept. 20, 2010) (on file with author) (discussing the Commission's goal of offering recommendations and proposals for ethically integrating technology into practice).

⁵⁷ MODEL RULES OF PRO. CONDUCT Preamble (AM. BAR ASS'N 2020).

⁸⁰ See *WJP Rule of Law Index, United States*, WORLD JUST. PROJECT, <https://worldjusticeproject.org/rule-of-law-index/country/2022/United%20States/Civil%20Justice> [<https://perma.cc/B4QS-BQ75>] (ranking United States 115 out of 140 countries in access to civil justice); Ashwin Telang, Article, *The Promise and Peril of AI Legal Services to Equalize Justice*, 2023 HARV. J.L. & TECH. 1, 3 (Mar. 14, 2023) <https://jolt.law.harvard.edu/digest/the-promise-and-peril-of-ai-legal-services-to-equalize-justice> [<https://perma.cc/8XUB-4S5Z>] (describing AI's ability to answer legal questions and offer low-cost legal assistance).

FLORIDA BAR ETHICS OPINION
OPINION 24-1
January 19, 2024

Advisory ethics opinions are not binding.

Lawyers may use generative artificial intelligence (“AI”) in the practice of law but must protect the confidentiality of client information, provide accurate and competent services, avoid improper billing practices, and comply with applicable restrictions on lawyer advertising. Lawyers must ensure that the confidentiality of client information is protected when using generative AI by researching the program’s policies on data retention, data sharing, and selflearning. Lawyers remain responsible for their work product and professional judgment and must develop policies and practices to verify that the use of generative AI is consistent with the lawyer’s ethical obligations. Use of generative AI does not permit a lawyer to engage in improper billing practices such as double-billing. Generative AI chatbots that communicate with clients or third parties must comply with restrictions on lawyer advertising and must include a disclaimer indicating that the chatbot is an AI program and not a lawyer or employee of the law firm. Lawyers should be mindful of the duty to maintain technological competence and educate themselves regarding the risks and benefits of new technology.

RPC: 4-1.1; 4-1.1 Comment; 4-1.5(a); 4-1.5(e); 4-1.5(f)(2); 4-1.5(h); 4-1.6; 4-1.6 Comment; 4-1.6(c)(1); 4-1.6(e); 4-1.18 Comment; 4-3.1; 4-3.3; 4-4.1; 4-4.4(b); Subchapter 4-7; 4-7.13; 4-7.13(b)(3); 4-7.13(b)(5); 4-5.3(a)

OPINIONS: 76-33 & 76-38, Consolidated; 88-6; 06-2; 07-2; 10-2; 12-3; ABA Comm. on Ethics and Prof’l Responsibility, Formal Op. 498 (2021); ABA Comm. on Ethics and Prof’l Responsibility, Formal Op. 93-379 (1993); Iowa Ethics Opinion 11-01; New York State Bar Ethics Opinion 842

CASES: *Mata v. Avianca*, 22-cv-1461, 2023 WL 4114965, at 17 (S.D.N.Y. June 22, 2023); *Bartholomew v. Bartholomew*, 611 So. 2d 85, 86 (Fla. 2d DCA 1992); *The Florida Bar v. Carlon*, 820 So. 2d 891, 899 (Fla. 2002); *Att’y Grievance Comm’n of Maryland v. Manger*, 913 A.2d 1 (Md. 2006)

The Florida Bar Board of Governors has directed the Board Review Committee on Professional Ethics to issue an opinion regarding lawyers’ use of generative artificial intelligence (“AI”). The release of ChatGPT-3 in November 2022 prompted wide-ranging debates regarding lawyers’ use of generative AI in the practice of law. While it is impossible to determine the impact generative AI will have on the legal profession, this opinion is intended to provide guidance to Florida Bar members regarding some of the ethical implications of these new programs.

Generative AI are “deep-learning models” that compile data “to generate statistically probable outputs when prompted.” IBM, [What is generative AI?](https://research.ibm.com/blog/what-is-generative-AI), (April 20, 2023), <https://research.ibm.com/blog/what-is-generative-AI> (last visited 11/09/2023). Generative AI can create original images, analyze documents, and draft briefs based on written prompts. Often, these programs rely on large language models. The datasets utilized by generative AI large

language models can include billions of parameters making it virtually impossible to determine how a program came to a specific result. Tsedel Neeley, 8 Questions About Using AI Responsibly, Answered, Harv. Bus. Rev. (May 9, 2023).

While generative AI may have the potential to dramatically improve the efficiency of a lawyer's practice, it can also pose a variety of ethical concerns. Among other pitfalls, lawyers are quickly learning that generative AI can "hallucinate" or create "inaccurate answers that sound convincing." Matt Reynolds, vLex releases new generative AI legal assistant, A.B.A. J. (Oct. 17, 2023), <https://www.abajournal.com/web/article/vlex-releases-new-generative-ai-legal-assistant> (last visited 11/09/2023). In one particular incident, a federal judge sanctioned two unwary lawyers and their law firm following their use of false citations created by generative AI. *Mata v. Avianca*, 22-cv-1461, 2023 WL 4114965, at 17 (S.D.N.Y. June 22, 2023).

Even so, the judge's opinion explicitly acknowledges that "[t]echnological advances are commonplace and there is nothing inherently improper about using a reliable artificial intelligence tool for assistance." *Id.* at 1.

Due to these concerns, lawyers using generative AI must take reasonable precautions to protect the confidentiality of client information, develop policies for the reasonable oversight of generative AI use, ensure fees and costs are reasonable, and comply with applicable ethics and advertising regulations.

Confidentiality

When using generative AI, a lawyer must protect the confidentiality of the client's information as required by Rule 4-1.6 of the Rules Regulating The Florida Bar. The ethical duty of confidentiality is broad in its scope and applies to all information learned during a client's representation, regardless of its source. Rule 4-1.6, Comment. Absent the client's informed consent or an exception permitting disclosure, a lawyer may not reveal the information. In practice, the most common exception is found in subdivision (c)(1), which permits disclosure to the extent reasonably necessary to "serve the client's interest unless it is information the client specifically requires not to be disclosed[.]" Rule 4-1.6(c)(1). Nonetheless, it is recommended that a lawyer obtain the affected client's informed consent prior to utilizing a third-party generative AI program if the utilization would involve the disclosure of any confidential information.

Rule 4-1.6(e) also requires a lawyer to "make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the client's representation." Further, a lawyer's duty of competence requires "an understanding of the benefits and risks associated with the use of technology[.]" Rule 4-1.1, Comment.

When using a third-party generative AI program, lawyers must sufficiently understand the technology to satisfy their ethical obligations. For generative AI, this specifically includes knowledge of whether the program is "self-learning." A generative AI that is "self-learning" continues to develop its responses as it receives additional inputs and adds those inputs to its existing parameters. Neeley, *supra* n. 2. Use of a "self-learning" generative AI raises the

possibility that a client's information may be stored within the program and revealed in response to future inquiries by third parties.

Existing ethics opinions relating to cloud computing, electronic storage disposal, remote paralegal services, and metadata have addressed the duties of confidentiality and competence to prior technological innovations and are particularly instructive. In its discussion of cloud computing resources, Florida Ethics Opinion 12-3 cites to New York State Bar Ethics Opinion 842 and Iowa Ethics Opinion 11-01 to conclude that a lawyer should:

- Ensure that the provider has an obligation to preserve the confidentiality and security of information, that the obligation is enforceable, and that the provider will notify the lawyer in the event of a breach or service of process requiring the production of client information;
- Investigate the provider's reputation, security measures, and policies, including any limitations on the provider's liability; and
- Determine whether the provider retains information submitted by the lawyer before and after the discontinuation of services or asserts proprietary rights to the information.

While the opinions were developed to address cloud computing, these recommendations are equally applicable to a lawyer's use of third-party generative AI when dealing with confidential information.

Florida Ethics Opinion 10-2 discusses the maintenance and disposition of electronic devices that contain storage media and provides that a lawyer's duties extend from the lawyer's initial receipt of the device through the device's disposition, "including after it leaves the control of the lawyer." Opinion 10-2 goes on to reference a lawyer's duty of supervision and to express that this duty "extends not only to the lawyer's own employees but over entities outside the lawyer's firm with whom the lawyer contracts[.]" Id.

Florida Ethics Opinion 07-2 notes that a lawyer should only allow an overseas paralegal provider access to "information necessary to complete the work for the particular client" and "should provide no access to information about other clients of the firm." Additionally, while "[t]he requirement for informed consent from a client should be generally commensurate with the degree of risk involved[.]" including "whether a client would reasonably expect the lawyer or law firm to personally handle the matter and whether the non-lawyers will have more than a limited role in the provision of the services." Id. Again, this guidance seems equally applicable to a lawyer's use of generative AI.

Finally, Florida Ethics Opinion 06-2 provides that a lawyer should take reasonable steps to safeguard the confidentiality of electronic communications, including the metadata attached to those communications, and that the recipient should not attempt to obtain metadata information that they know or reasonably should know is not intended for the recipient. In the event that the recipient inadvertently receives metadata information, the recipient must "promptly notify the sender," as is required by Rule 4-4.4(b). Similarly, a lawyer using generative AI should take reasonable precautions to avoid the inadvertent disclosure of confidential information and should not attempt to access information previously provided to the generative AI by other lawyers.

It should be noted that confidentiality concerns may be mitigated by use of an inhouse generative AI rather than an outside generative AI where the data is hosted and stored by a thirdparty. If the use of a generative AI program does not involve the disclosure of confidential information to a third-party, a lawyer is not required to obtain a client's informed consent pursuant to Rule 4-1.6.

Oversight of Generative AI

While Rule 4-5.3(a) defines a nonlawyer assistant as a "a person," many of the standards applicable to nonlawyer assistants provide useful guidance for a lawyer's use of generative AI.

First, just as a lawyer must make reasonable efforts to ensure that a law firm has policies to reasonably assure that the conduct of a nonlawyer assistant is compatible with the lawyer's own professional obligations, a lawyer must do the same for generative AI. Lawyers who rely on generative AI for research, drafting, communication, and client intake risk many of the same perils as those who have relied on inexperienced or overconfident nonlawyer assistants.

Second, a lawyer must review the work product of a generative AI in situations similar to those requiring review of the work of nonlawyer assistants such as paralegals. Lawyers are ultimately responsible for the work product that they create regardless of whether that work product was originally drafted or researched by a nonlawyer or generative AI.

Functionally, this means a lawyer must verify the accuracy and sufficiency of all research performed by generative AI. The failure to do so can lead to violations of the lawyer's duties of competence (Rule 4-1.1), avoidance of frivolous claims and contentions (Rule 4-3.1), candor to the tribunal (Rule 4-3.3), and truthfulness to others (Rule 4-4.1), in addition to sanctions that may be imposed by a tribunal against the lawyer and the lawyer's client.

Third, these duties apply to nonlawyers "both within and outside of the law firm." ABA Comm. on Ethics and Prof'l Responsibility, Formal Op. 498 (2021); see Fla. Ethics Op. 07-2. The fact that a generative AI is managed and operated by a third-party does not obviate the need to ensure that its actions are consistent with the lawyer's own professional and ethical obligations.

Further, a lawyer should carefully consider what functions may ethically be delegated to generative AI. Existing ethics opinions have identified tasks that a lawyer may or may not delegate to nonlawyer assistants and are instructive. First and foremost, a lawyer may not delegate to generative AI any act that could constitute the practice of law such as the negotiation of claims or any other function that requires a lawyer's personal judgment and participation.

Florida Ethics Opinion 88-6 notes that, while nonlawyers may conduct the initial interview with a prospective client, they must:

- Clearly identify their nonlawyer status to the prospective client;
- Limit questions to the purpose of obtaining factual information from the prospective client; and

- Not offer any legal advice concerning the prospective client’s matter or the representation agreement and refer any legal questions back to the lawyer.

This guidance is especially useful as law firms increasingly utilize website chatbots for client intake. While generative AI may make these interactions seem more personable, it presents additional risks, including that a prospective client relationship or even a lawyer-client relationship has been created without the lawyer’s knowledge.

The Comment to Rule 4-1.18 (Duties to Prospective Client) explains what constitutes a consultation:

A person becomes a prospective client by consulting with a lawyer about the possibility of forming a client-lawyer relationship with respect to a matter. Whether communications, including written, oral, or electronic communications, constitute a consultation depends on the circumstances. For example, a consultation is likely to have occurred if a lawyer, either in person or through the lawyer’s advertising in any medium, specifically requests or invites the submission of information about a potential representation without clear and reasonably understandable warnings and cautionary statements that limit the lawyer’s obligations, and a person provides information in response. In contrast, a consultation does not occur if a person provides information to a lawyer in response to advertising that merely describes the lawyer’s education, experience, areas of practice, and contact information, or provides legal information of general interest. A person who communicates information unilaterally to a lawyer, without any reasonable expectation that the lawyer is willing to discuss the possibility of forming a client-lawyer relationship, is not a “prospective client” within the meaning of subdivision (a).

Similarly, the existence of a lawyer-client relationship traditionally depends on the subjective reasonable belief of the client regardless of the lawyer’s intent. *Bartholomew v. Bartholomew*, 611 So. 2d 85, 86 (Fla. 2d DCA 1992).

For these reasons, a lawyer should be wary of utilizing an overly welcoming generative AI chatbot that may provide legal advice, fail to immediately identify itself as a chatbot, or fail to include clear and reasonably understandable disclaimers limiting the lawyer’s obligations.

Just as with nonlawyer staff, a lawyer should not instruct or encourage a client to rely solely on the “work product” of generative AI, such as due diligence reports, without the lawyer’s own personal review of that work product.

Legal Fees and Costs

Rule 4-1.5(a) prohibits lawyers from charging, collecting, or agreeing to fees or costs that are illegal or clearly excessive while subdivision (b) provides a list of factors to consider when determining whether a fee or cost is reasonable. A lawyer must communicate the basis for fees and costs to a client and it is preferable that the lawyer do so in writing. Rule 4-1.5(e).

Contingent fees and fees that are nonrefundable in any part must be explained in writing. Rule 41.5(e); Rule 4-1.5(f)(2).

Regarding costs, a lawyer may only ethically charge a client for the actual costs incurred on the individual client's behalf and must not duplicate charges that are already accounted for in the lawyer's overhead. *See, The Florida Bar v. Carlon*, 820 So. 2d 891, 899 (Fla. 2002) (lawyer sanctioned for violations including a \$500.00 flat administrative charge to each client's file); ABA Comm. on Ethics and Prof'l Responsibility, Formal Op. 93-379 (1993) (lawyer should only charge clients for costs that reasonably reflect the lawyer's actual costs); Rule 4-1.5(h) (lawyers accepting payment via a credit plan may only charge the actual cost imposed on the transaction by the credit plan).

Regarding fees, a lawyer may not ethically engage in any billing practices that duplicate charges or that falsely inflate the lawyer's billable hours. Though generative AI programs may make a lawyer's work more efficient, this increase in efficiency must not result in falsely inflated claims of time. In the alternative, lawyers may want to consider adopting contingent fee arrangements or flat billing rates for specific services so that the benefits of increased efficiency accrue to the lawyer and client alike.

While a lawyer may separately itemize activities like paralegal research performed by nonlawyer personnel, the lawyer should not do so if those charges are already accounted for in the lawyer's overhead. Fla. Ethics Op. 76-33 & 76-38, Consolidated. In the alternative, the lawyer may need to consider crediting the nonlawyer time against the lawyer's own fees. *Id.* Florida Ethics Opinion 07-2 discusses the outsourcing of paralegal services in contingent fee matters and explains:

The law firm may charge a client the actual cost of the overseas provider [of paralegal services], unless the charge would normally be covered as overhead. However, in a contingent fee case, it would be improper to charge separately for work that is usually otherwise accomplished by a client's own attorney and incorporated into the standard fee paid to the attorney, even if that cost is paid to a third-party provider.

Additionally, a lawyer should have sufficient general knowledge to be capable of providing competent representation. *See, e.g., Att'y Grievance Comm'n of Maryland v. Manger*, 913 A.2d 1 (Md. 2006). "While it may be appropriate to charge a client for case-specific research or familiarization with a unique issue involved in a case, general education or background research should not be charged to the client." *Id.* at 5.

In the context of generative AI, these standards require a lawyer to inform a client, preferably in writing, of the lawyer's intent to charge a client the actual cost of using generative AI. In all instances, the lawyer must ensure that the charges are reasonable and are not duplicative. If a lawyer is unable to determine the actual cost associated with a particular client's matter, the lawyer may not ethically prorate the periodic charges of the generative AI and instead should account for those charges as overhead. Finally, while a lawyer may charge a client for the reasonable time spent for case-specific research and drafting when using generative AI, the

lawyer should be careful not to charge for the time spent developing minimal competence in the use of generative AI.

Lawyer Advertising

The advertising rules in Subchapter 4-7 of the Rules Regulating The Florida Bar include prohibitions on misleading content and unduly manipulative or intrusive advertisements.

Rule 4-7.13 prohibits a lawyer from engaging in advertising that is deceptive or inherently misleading. More specifically, subdivision (b) includes prohibitions on:

(3) comparisons of lawyers or statements, words, or phrases that characterize a lawyer's or law firm's skills, experience, reputation, or record, unless the characterization is objectively verifiable; [and]

* * *

(5) [use of] a voice or image that creates the erroneous impression that the person speaking or shown is the advertising lawyer or a lawyer or employee of the advertising firm unless the advertisement contains a clear and conspicuous disclaimer that the person is not an employee or member of the law firm[.]

As noted above, a lawyer should be careful when using generative AI chatbot for advertising and intake purposes as the lawyer will be ultimately responsible in the event the chatbot provides misleading information to prospective clients or communicates in a manner that is inappropriately intrusive or coercive. To avoid confusion or deception, a lawyer must inform prospective clients that they are communicating with an AI program and not with a lawyer or law firm employee. Additionally, while many visitors to a lawyer's website voluntarily seek information regarding the lawyer's services, a lawyer should consider including screening questions that limit the chatbot's communications if a person is already represented by another lawyer.

Lawyers may advertise their use of generative AI but cannot claim their generative AI is superior to those used by other lawyers or law firms unless the lawyer's claims are objectively verifiable. Whether a particular claim is capable of objective verification is a factual question that must be made on a case-by-case basis.

Conclusion

In sum, a lawyer may ethically utilize generative AI technologies but only to the extent that the lawyer can reasonably guarantee compliance with the lawyer's ethical obligations. These obligations include the duties of confidentiality, avoidance of frivolous claims and contentions, candor to the tribunal, truthfulness in statements to others, avoidance of clearly excessive fees and costs, and compliance with restrictions on advertising for legal services. Lawyers should be cognizant that generative AI is still in its infancy and that these ethical concerns should not be treated as an exhaustive list. Rather, lawyers should continue to develop competency in their use of new technologies and the risks and benefits inherent in those technologies.

RULES, PROCEDURE, COMMENTS

All opinions of the Ethics Committee are predicated upon the North Carolina Rules of Professional Conduct. Any interested person or group may submit a written comment – including comments in support of or against the proposed opinion – or request to be heard concerning a proposed opinion. The Ethics Committee welcomes and encourages the submission of comments, and all comments are considered by the committee at the next quarterly meeting. Any comment or request should be directed to the Ethics Committee at ethicscomments@ncbar.gov no later than March 30, 2024.

Council Actions

At its meeting on January 19, 2024, the State Bar Council adopted the ethics opinion summarized below:

2023 Formal Ethics Opinion 4

Use of a Lawyer’s Trade Name for Keyword Advertisements in an Internet Search Engine

Proposed opinion rules that the intentional selection of another lawyer’s unique firm trade name in a keyword advertisement campaign is prohibited, but that prohibition does not apply when the trade name is also a common search term.

Ethics Committee Actions

At its meeting on January 18, 2024, the Ethics Committee considered a total of six inquiries, including the opinion noted above. Four inquiries were sent or returned to subcommittee for further study, including an inquiry addressing a lawyer’s ability to obligate a client’s estate to pay the lawyer for any time spent defending the lawyer’s work in drafting and executing the client’s will and an inquiry exploring a lawyer’s duty of confidentiality when inheriting confidential client information. Additionally, in October 2023 the Ethics Committee published Proposed 2023 Formal Ethics Opinion 3, Installation of Third Party’s Self-Service Kiosk in Lawyer’s Office and Inclusion of Lawyer in Third Party’s Advertising Efforts; based on comments received during publication, the committee voted to return the inquiry to subcommittee for further study. The committee also approved the publication of one new proposed formal ethics opinion on a lawyer’s use of artificial intelligence in a law practice, which appears below.

Proposed 2024 Formal Ethics Opinion 1 Use of Artificial Intelligence in a Law Practice

January 18, 2024

Proposed opinion discusses a lawyer’s professional responsibility when using artificial intelligence in a law practice.

Editor’s Note: There is an increasingly vast number of helpful resources on understanding Artificial Intelligence and the technology’s interaction with the legal profession. The resources referenced in this opinion are not exhaustive but are intended to serve as a starting point for a lawyer’s understanding of the topic. Over time, this editor’s note may be updated as additional resources are published that staff concludes would be beneficial to lawyers.

Background

“Artificial intelligence” (hereinafter, “AI”) is a broad and evolving term encompassing myriad programs and processes with myriad capabilities. While a single definition of AI is not yet settled (and likely impossible), for the purposes of this opinion, the term “AI” refers to “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments.” Nat’l Artificial Intelligence Initiative Act of 2020, Div. E, sec. 5002(3) (2021). Said in another, over-simplified way, AI is the use of computer science and extensive data sets to enable problem solving or decision-making, often through the implementation of sophisticated algorithms. AI encompasses, but is not limited to, both extractive and generative AI,¹ natural language processing, large language models, and any number of machine learning processes.² Examples of law-related AI programs range from online electronic legal research and case management software to ediscovery tools and programs that draft legal documents (e.g., a trial brief, will, etc.) based upon the lawyer’s input of information that may or may not be client-specific.

Most lawyers have likely used some form of AI when practicing law, even if they didn’t realize it (e.g., widely used online legal research subscription services utilize a type of extractive AI, or a program that “extracts” information relevant to the user’s inquiry from a large set of existing data upon which the program has been trained). Within the year preceding the date of this opinion, generative AI programs that create products in response to a user’s request based upon a large set of existing data upon which the program has been trained (e.g., Chat-GPT) have grown in capability and popularity, generating both positive and negative reactions regarding the integration of these technological breakthroughs in the legal profession.³ It is unquestioned that AI can be used in the practice of law to increase efficiency and consistency in the provision of legal services. However, AI and its work product can be inaccurate or unreliable despite its appearance of reliability when used during the provision of legal services.⁴ **Inquiry #1:**

Considering the advantages and disadvantages of using AI in the provision of legal services, is a lawyer permitted to use AI in a law practice?

Opinion #1:

Yes, provided the lawyer uses any AI program, tool, or resource competently, securely to protect client confidentiality, and with proper supervision when relying upon or implementing the AI’s work product in the provision of legal services.

On the spectrum of law practice resources, AI falls somewhere between programs, tools, and processes readily used in law practice today (e.g. case management systems, trust account management programs, electronic legal research, etc.) and nonlawyer support staff (e.g. paralegals, summer associates, IT professionals, etc.). Nothing in the Rules of Professional Conduct specifically addresses, let alone prohibits, a lawyer’s use of AI in a law practice. However, should a lawyer choose to employ AI in a practice, the lawyer must do so competently, the lawyer must do so securely, and the lawyer must exercise independent judgment in supervising the use of such processes.

Rule 1.1 prohibits lawyers from “handl[ing] a legal matter that the lawyer knows or should know he or she is not competent to handle[.]” and goes on to note that “[c]ompetent representation

requires the legal knowledge, skill, thoroughness, and preparation reasonably necessary for the representation.” Comment 8 to Rule 1.1 recognizes the reality of advancements in technology impacting a lawyer’s practice, and states that part of a lawyer’s duty of competency is to “keep abreast of changes in the law and its practice, including the benefits and risks associated with the technology relevant to the lawyer’s practice[.]” Rule 1.6(c) requires a lawyer to “make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.” Rule 5.3 requires a lawyer to “make reasonable efforts to ensure that the firm or organization has in effect measures giving reasonable assurance that the nonlawyer’s conduct is compatible with the professional obligations of the lawyer[.]” and further requires that “a lawyer having direct supervisory authority over the nonlawyer shall make reasonable efforts to ensure that the nonlawyer’s conduct is compatible with the professional obligations of the lawyer[.]” Rules 5.3(a) and (b). The requirements articulated in Rule 5.3 apply to nonlawyer assistants within a law firm as well as those outside of a law firm that are engaged to provide assistance in the lawyer’s provision of legal services to clients, such as third-party software companies. *See* 2011 FEO 6 (“Although a lawyer may use nonlawyers outside of the firm to assist in rendering legal services to clients, Rule 5.3(a) requires the lawyer to make reasonable efforts to ensure that the services are provided in a manner that is compatible with the professional obligations of the lawyer.”).

A lawyer may use AI in a variety of manners in connection with a law practice, and it is a lawyer’s responsibility to exercise independent professional judgment in determining how (or if) to use the product of an AI tool in furtherance of the representation of a client. From discovery and document review to legal research, drafting contracts, and aggregating/analyzing data trends, the possibilities for employing AI in a law practice are increasingly present and constantly evolving. A lawyer’s decision to use and rely upon AI to assist in the lawyer’s representation of a client is generally hers alone and one to be determined depending upon a number of factors, including the impact of such services, the cost of such services, and the reliability of the processes.⁵ This opinion does not attempt to dictate when and how AI is appropriate for a law practice.

Should a lawyer decide to employ AI in the representation of a client, however, the lawyer is fully responsible for the use and impact of AI in the client’s case. The lawyer must use the AI tool in a way that meets the competency standard set out in Rule 1.1. Like other software, the lawyer employing an AI tool must educate herself on the benefits and risks associated with the tool, as well as the impact of using the tool on the client’s case. Educational efforts include, but are not limited to, reviewing current and relevant resources on AI broadly and on the specific program intended for use during the provision of legal services. A lawyer that inputs confidential client information into an AI tool must take steps to ensure the information remains secure and protected from unauthorized access or inadvertent disclosure per Rule 1.6(c). Additionally, a lawyer utilizing an outside third-party company’s AI program or service must make reasonable efforts to ensure that the program or service used is compatible with the lawyer’s responsibilities under the Rules of Professional Conduct pursuant to Rule 5.3. Whether the lawyer is reviewing the results of a legal research program, a keyword search of emails for production during discovery, proposed reconciliations of the lawyer’s trust account prepared by a long-time assistant, or a risk analysis of potential borrowers for a lender-client produced by an AI process,

the lawyer is individually responsible for reviewing, evaluating, and ultimately relying upon the work produced by someone—or something—other than the lawyer.

Inquiry #2:

May a lawyer provide or input a client’s documents, data, or other information to a third-party company’s AI program for assistance in the provision of legal services?

Opinion #2:

Yes, provided the lawyer has satisfied herself that the third-party company’s AI program is sufficiently secure and complies with the lawyer’s obligations to ensure any client information will not be inadvertently disclosed or accessed by unauthorized individuals pursuant to Rule 1.6(c).

At the outset, the Ethics Committee does not opine on whether the information shared with an AI tool violates the attorney-client privilege, as the issue is a legal question and outside the scope of the Rules of Professional Conduct. A lawyer should research and resolve any question on privilege prior to engaging with a third-party company’s AI program for use in the provision of legal services to a client, particularly if client-specific information will be provided to the AI program.

This inquiry is akin to any lawyer providing confidential information to a third-party software program (practice management, cloud storage, etc.), on which the Ethics Committee has previously opined. As noted above, a lawyer has an obligation to “make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating the representation of the client.” Rule 1.6(c). What constitutes “reasonable efforts” will vary depending on the circumstances related to the practice and representation, as well as a variety of factors including the sensitivity of the information and the cost or benefit of employing additional security measures to protect the information. Rule 1.6, cmt. [19]. Ultimately, “[a] lawyer must take steps to minimize the risk that confidential client information will be disclosed to other clients or to third parties” when using technology to handle, communicate, analyze, or otherwise interact with confidential client information. 2008 FEO 5; *see also* 2005 FEO 10; 2011 FEO 6.

The Ethics Committee in 2011 FEO 6 recognized that employing a third-party company’s services/technology with regards to confidential client information requires a lawyer to exercise reasonable care when selecting a vendor. The opinion states:

[W]hile the duty of confidentiality applies to lawyers who choose to use technology to communicate, this obligation does not require that a lawyer use only infallibly secure methods of communication. Rather, the lawyer must use reasonable care to select a mode of communication that, in light of the circumstances, will best protect confidential client information and the lawyer must advise effected parties if there is reason to believe that the chosen communications technology presents an unreasonable risk to confidentiality....A lawyer must fulfill the duties to protect confidential client information and to safeguard client files by applying the same diligence and competency to manage the risks of [technology] that the lawyer is required to apply when representing clients.

2011 FEO 6 (internal citations omitted). In exercising reasonable care, the opinion discusses a sample of considerations for evaluating whether a particular third-party company's services are compatible with the lawyer's professional responsibility, including:

- The experience, reputation, and stability of the company;
- Whether the terms of service include an agreement on how the company will handle confidential client information, including security measures employed by the company to safeguard information provided by the lawyer; and
- Whether the terms of service clarify how information provided to the company will be retrieved by the lawyer or otherwise safely destroyed if not retrieved should the company go out of business, change ownership, or if services are terminated.

2011 FEO 6; *see* Rule 5.3. A proposed ethics opinion from the Florida Bar on a lawyer's use of AI adds that lawyers should "[d]etermine whether the provider retains information submitted by the lawyer before and after the discontinuation of services or asserts proprietary rights to the information" when determining whether a third-party company's technological services are compatible with the lawyer's duty of confidentiality. *See* Florida Bar Proposed Advisory Opinion 24-1 (published Nov. 13, 2023).

Furthermore, this duty of reasonable care continues beyond initial selection of a service, program, or tool and extends throughout the lawyer's use of the service. A lawyer should continuously educate herself on the selected technology and developments thereto—both individually and by "consult[ing] periodically with professionals competent in the area of online security"—and make necessary adjustments (including abandonment, if necessary) when discoveries are made that call into question services previously thought to be secure. 2011 FEO 6.

The aforementioned considerations—including the consideration regarding ownership of information articulated by the Florida Bar opinion—are equally applicable to a lawyer's selection and use of a third-party company's AI service/program. Just as with any third-party service, a lawyer has a duty under Rule 5.3 to make reasonable efforts to ensure the third-party AI program or service is compatible with the lawyer's professional responsibility, particularly with regards to the lawyer's duty of confidentiality pursuant to Rule 1.6. Importantly, some current AI programs are publicly available to all consumers/users, and the nature of these AI programs are to retain and train itself based on the information provided by any user of its program. Lawyers should educate themselves on the nature of any publicly available AI program intended to be used in the provision of legal services, with particular focus on whether the AI program will retain and subsequently use the information provided by the user. Generally, and as of the date of this opinion, lawyers should avoid inputting client-specific information into publicly available AI resources.

Inquiry #3:

If a firm were to have an AI software tool initially developed by a third-party but then used the AI tool in-house using law firm owned servers and related infrastructure, does that change the data security requirement analysis in Opinion #2?

Opinion #3:

No. Lawyer remains responsible for keeping the information secure pursuant to Rule 1.6(c) regardless of the program's location. While an in-house program may seem more secure because the program is maintained and run using local servers, those servers may be more vulnerable to attack because a lawyer acting independently may not be able to match the security features typically employed by larger companies whose reputations are built in part on security and customer service. A lawyer who plans to independently store client information should consult an information technology/cybersecurity expert about steps needed to adequately protect the information stored on local servers.

Relatedly, AI programs developed for use in-house or by a particular law practice may also be derivatives of a single, publicly available AI program; as such, some of these customized programs may continue to send information inputted into the firm-specific program back to the central program for additional use or training. Again, prior to using such a program, a lawyer must educate herself on the nuances and operation of the program to ensure client information will remain protected in accordance with the lawyer's professional responsibility. The list of considerations found in Opinion #2 offers a starting point for questions to explore when identifying, evaluating, and selecting a vendor.

Inquiry #4:

If a lawyer signs a pleading based on information generated from AI, is there variation from traditional or existing ethical obligations and expectations placed on lawyers signing pleadings absent AI involvement?

Opinion #4:

No. A lawyer may not abrogate her responsibilities under the Rules of Professional Conduct by relying upon AI. Per Rule 3.1, a lawyer is prohibited from bringing or defending "a proceeding, or assert[ing] or controvert[ing] an issue therein, unless there is a basis in law and fact for doing so that is not frivolous[.]" A lawyer's signature on a pleading also certifies the lawyer's good faith belief as to the factual and legal assertions therein. *See* N.C. R. Civ. Pro. 11 ("The signature of an attorney...constitutes a certificate by him that he has read the pleading, motion, or other paper; that to the best of his knowledge, information, and belief formed after reasonable inquiry it is well grounded in fact and is warranted by existing law or a good faith argument for the extension, modification, or reversal of existing law, and that it is not interposed for any improper purpose, such as to harass or to cause unnecessary delay or needless increase in the cost of litigation."). If the lawyer employs AI in her practice and adopts the tool's product as her own, the lawyer is professionally responsible for the use of the tool's product. *See* Opinion #1.

Inquiry #5:

If a lawyer uses AI to assist in the representation of a client, is the lawyer under any obligation to inform the client that the lawyer has used AI in furtherance of the representation or legal services provided? **Opinion #5:**

The answer to this question depends on the type of technology used, the intended product from the technology, and the level of reliance placed upon the technology/technology's product. Ultimately, the attorney/firm will need to evaluate each case and each client individually. Rule

1.4(b) requires an attorney to explain a matter to her client “to the extent reasonably necessary to permit the client to make informed decisions regarding the representation.” Generally, a lawyer need not inform her client that she is using an AI tool to complete ordinary tasks, such as conducting legal research or generic case/practice management. However, if a lawyer delegates substantive tasks in furtherance of the representation to an AI tool, the lawyer’s use of the tool is akin to outsourcing legal work to a nonlawyer, for which the client’s advanced informed consent is required. See 2007 FEO 12. Additionally, if the decision to use or not use an AI tool in the case requires the client’s input with regard to fees, the lawyer must inform and seek input from the client.

Inquiry #6:

Lawyer has an estate planning practice and bills at the rate of \$300 per hour. Lawyer has integrated an AI program into the provision of legal services, resulting in increased efficiency and work output. For example, Lawyer previously spent approximately three hours drafting standard estate planning documents for a client; with the use of AI, Lawyer now spends only one hour preparing those same documents for a client. May Lawyer bill the client for the three hours of work that the prepared estate documents represent?

Opinion #6:

No, Lawyer may not bill a client for three hours of work when only one hour of work was actually experienced. A lawyer’s billing practices must be accurate, honest, and not clearly excessive. Rules 7.1, 8.4(c), and 1.5(a); *see also* 2022 FEO 4. If the use of AI in Lawyer’s practice results in greater efficiencies in providing legal services, Lawyer may enjoy the benefit of those new efficiencies by completing more work for more clients; Lawyer may not inaccurately bill a client based upon the “time-value represented” by the end product should Lawyer not have used AI when providing legal services.

Rather than billing on an hourly basis, Lawyer may consider billing clients a flat fee for the drafting of documents—even when using AI to assist in drafting—provided the flat fee charged is not clearly excessive and the client consents to the billing structure. *See* 2022 FEO 4.

Relatedly, Lawyer may also bill a client for actual expenses incurred when employing AI in the furtherance of a client’s legal services, provided the expenses charged are accurate, not clearly excessive, and the client consents to the charge, preferably in writing. *See* Rule 1.5(b). Lawyer may not bill a general “administrative fee” for the use of AI during the representation of a client; rather, any cost charged to a client based on Lawyer’s use of AI must be specifically identified and directly related to the legal services provided to the client during the representation. For example, if Lawyer has generally incorporated AI into her law practice for the purpose of case management or drafting assistance upon which Lawyer may or may not rely when providing legal services to all clients, Lawyer may not bill clients a generic administrative fee to offset the costs Lawyer experiences related to her use of AI. However, if Lawyer employs AI on a limited basis for a single client to assist in the provision of legal services, Lawyer may charge those expenses to the client provided the expenses are accurate, not clearly excessive, and the client consents to the expense and charge, preferably in writing.

Endnotes

1. For a better understanding of the differences between extractive and generative AI, see Jake Nelson, *Combining Extractive and Generative AI for New Possibilities*, LexisNexis (June 6, 2023), [lexisnexis.com/community/insights/legal/b/thought-leadership/posts/combiningextractive-and-generative-ai-for-new-possibilities](https://www.lexisnexis.com/community/insights/legal/b/thought-leadership/posts/combiningextractive-and-generative-ai-for-new-possibilities) (last visited January 10, 2024).
2. For an overview of the state of AI as of the date of this opinion, see *What is Artificial Intelligence (AI)?*, IBM, [ibm.com/topics/artificial-intelligence](https://www.ibm.com/topics/artificial-intelligence) (last visited January 10, 2024). For information on how AI relates to the legal profession, see *AI Terms for Legal Professionals: Understanding What Powers Legal Tech*, LexisNexis (March 20, 2023), [lexisnexis.com/community/insights/legal/b/thought-leadership/posts/ai-terms-for-legal-professionals-understanding-what-powers-legal-tech](https://www.lexisnexis.com/community/insights/legal/b/thought-leadership/posts/ai-terms-for-legal-professionals-understanding-what-powers-legal-tech) (last visited January 10, 2024).
3. John Villasenor, *How AI Will Revolutionize the Practice of Law*, Brookings Institution (March 20, 2023), [brookings.edu/articles/how-ai-will-revolutionize-the-practice-of-law/](https://www.brookings.edu/articles/how-ai-will-revolutionize-the-practice-of-law/) (last visited January 10, 2024); Steve Lohr, *AI is Coming for Lawyers Again*, New York Times (April 10, 2023), [nytimes.com/2023/04/10/technology/ai-is-coming-for-lawyers-again.html](https://www.nytimes.com/2023/04/10/technology/ai-is-coming-for-lawyers-again.html) (last visited January 10, 2024).
4. Larry Neumeister, *Lawyers Blame ChatGPT for Tricking Them Into Citing Bogus Case Law*, AP News (June 8, 2023), [apnews.com/article/artificial-intelligence-chatgpt-courtse15023d7e6fdf4f099aa122437dbb59b](https://www.apnews.com/article/artificial-intelligence-chatgpt-courtse15023d7e6fdf4f099aa122437dbb59b) (last visited January 10, 2024).
5. In certain circumstances a lawyer may need to consult a client about employing AI in the provision of legal services to that client, see Opinion #5, below.

The Ethics Committee welcomes feedback on the proposed opinion; feedback should be sent to ethicscomments@ncbar.gov.

Tech Competence From the Courtroom to Cyberspace

“Tech Competence” is a lawyer's ability to understand and use technology in the practice of law. In today's digital age, tech competence is vital for lawyers to effectively represent their clients and comply with legal ethics rules.

For example, American Bar Association Model Rule 1.1 states that “a lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.” Comment 8 to Model Rule 1.1 further states that “a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology.” Thus, tech competence generally requires: staying informed about relevant technology, regularly assessing the benefits and risks of technology, and taking reasonable measures to protect client confidences and information in connection with new technology. See ABA Formal Ethics Op. 477R (2017); Lawyer-Client Relationship – Competence.

Considering the ABA's guidance and emerging technologies, this article discusses tech competence in four areas of growing importance: cybersecurity and cloud storage; e-discovery and electronically stored information (ESI); artificial intelligence (AI) and automation; and social media and software applications.

Cybersecurity & Cloud Storage

As lawyers increase their reliance on online software and cloud storage, their ethical obligations with regard to tech competence become more entangled with cybersecurity. Lawyers generally understand that they have a duty to take “reasonable steps” to safeguard client information and protect against unauthorized access. See ABA Formal Ethics Op. 477R (2017). In addition, as exemplified by Model Rule 1.6(c) lawyers should safeguard the “Client-Lawyer Relationship” and make reasonable efforts to prevent the inadvertent or unauthorized disclosure of client information.

Some lawyers may not appreciate, however, that their duty of competence may extend to protecting against data breaches and cybersecurity threats ABA Formal Ethics Op. 477R (2017) states that: “Each device and each storage location offer an opportunity for the inadvertent or unauthorized disclosure of information relating to the representation and thus implicates a lawyer's ethical duties.” Cybersecurity is a key issue for lawyers and law firms, which are generally targeted for two main reasons: “(1) they obtain, store and use highly sensitive information about their clients while at times utilizing safeguards to shield that information that may be inferior to those deployed by the client, and (2) the information in their possession is more likely to be of interest to a hacker and likely less voluminous than that held by the client.” ABA Formal Ethics Op. 477R (2017).

ABA Formal Ethics Op. 483 (2017) explains lawyers' obligations in the event of a data breach, including the duty to notify affected clients and mitigate harm. Lawyers should take reasonable precautions to safeguard client information against unauthorized access or disclosure, whether on a mobile device or in the cloud, and whether maintained by the law firm or a third-party vendor. If a breach occurs, the lawyer has "a duty to notify clients of the data breach ... in sufficient detail to keep clients 'reasonably informed' and with an explanation 'to the extent necessary to permit the client to make informed decisions regarding the representation.'" ABA Formal Ethics Op. 483 (2017).

Notably though, the legal profession recognizes that cloud storage is often necessary in today's digital age to store the exponentially increasing data of law firms and their clients. Thus, states do not attempt to per se prohibit the use of the cloud, and, rather, they focus on its reasonable and competent use with proper safeguards.

As the Professional Ethics Committee for the State Bar of Texas stated in September 2018, "Cloud-based electronic storage and software systems are in wide use among the general public and lawyers. While wide usage of an information storage method or software document creation system is not, in itself, justification for its use by lawyers, alternative methods of information storage and document preparation also have an inherent risk of disclosure or misuse." Thus, "a lawyer must take reasonable precautions in the adoption and use of cloud-based technology for client document and data storage or the creation of client-specific documents that require client confidential information." Texas Ethics Op. 680 (2018).

When considering what tech competence means as it relates to cloud storage, the Texas opinion provides that reasonable precautions include:

- (1) acquiring a general understanding of how the cloud technology works;
 - (2) reviewing the 'terms of service' to which the lawyer submits when using a specific cloud-based provider just as the lawyer should do when choosing and supervising other types of service providers;
 - (3) learning what protections already exist within the technology for data security;
 - (4) determining whether additional steps, including but not limited to the encryption of client confidential information, should be taken before submitting that client information to a cloud-based system;
 - (5) remaining alert as to whether a particular cloud-based provider is known to be deficient in its data security measures or is or has been unusually vulnerable to 'hacking' of stored information;
- and

(6) training for lawyers and staff regarding appropriate protections and considerations.”

Ethics opinions from many states illustrate the balance of necessary and reasonable use:

- **Washington Ethics Op. 2215 (2012).** “A lawyer may use online data storage systems to store and back up client confidential information as long as the lawyer takes reasonable care to ensure that the information will remain confidential and that the information is secure against risk of loss.”
- **Pennsylvania Ethics Op. 2011-200 (2011).** Describing the steps that a lawyer should take when dealing with “cloud” computing, including detailed lists of required steps and descriptions of what other states have held on this issue; ... Firms may be eager to capitalize on cloud services in an effort to promote mobility, flexibility, organization and efficiency, reduce costs, and enable lawyers to focus more on legal, rather than technical and administrative issues. However, lawyers must be conscientious about maintaining traditional confidentiality, competence, and supervisory standards.”
- **Alabama Ethics Op. 2010-02 (2010).** Analyzing various issues relating to client files; allowing lawyers to retain the client files in the “cloud” as long as they take reasonable steps to maintain the confidentiality of the data.
- **California Ethics Op. 2010-179 (2010).** “Whether an attorney violates his or her duties of confidentiality and competence when using technology to transmit or store confidential client information will depend on the particular technology being used and the circumstances surrounding such use.”
- **New York Ethics Op. 842 (2010).** “[T]he lawyer should stay abreast of technological advances to ensure that the storage system remains sufficiently advanced to protect the client's information, and should monitor the changing law of privilege to ensure that storing the information online will not cause loss or waiver of any privilege.”

E-Discovery & ESI

E-discovery has become vital to litigation and investigations in the digital age. As the amount of ESI continues to grow, to be considered competent, lawyers should have a basic understanding of the technology involved in e-discovery and be able to effectively manage e-discovery projects.

Lawyers have a duty to take reasonable steps to preserve potentially relevant ESI as soon as litigation or an investigation is reasonably anticipated, including through a litigation hold. This duty includes identifying and collecting relevant ESI, ensuring it is not destroyed or altered, and acting to prevent the loss of ESI through technological failures or other factors. Failure to properly preserve

relevant ESI can result in spoliation sanctions, adverse inferences, or even malpractice claims.

Another key issue in e-discovery is the review and production of ESI. Tech competent lawyers should be able to effectively review and produce large volumes of ESI in a cost-effective and defensible manner. This standard requires a knowledge of the technology and best practices involved in e-discovery, including the potential use of technology-assisted review, protective orders, and clawback agreements. Lawyers should also understand the legal and ethical issues that arise in the review and production of ESI. One important consideration is the duty to protect privileged or confidential information, whether stored internally or with third-party vendors. Lawyers should be able to effectively identify, protect, and sanitize privileged or confidential information during e-discovery, including information revealed in metadata.

Relatedly, under Model Rule 1.4 and ABA Formal Ethics Op. 481, lawyers should generally advise their current clients of material errors related to the representation, which may be implicated by their material mishandling of e-discovery projects and ESI. According to ABA Formal Ethics Op. 481 and Model Rule 1.6(b)(4), before potentially informing the client of the error, the lawyer may generally “consult with his or her law firm's general counsel, another lawyer, or the lawyer's professional liability insurer” regarding prompt compliance with their professional obligations.

AI & Automation Technology

AI and automation technologies are rapidly changing the legal industry and legal writing. Lawyers may want to use AI-powered tools to automate routine tasks, analyze large amounts of data, outline arguments, and even provide legal advice. Indeed, one AI program recently passed the Uniform Bar Exam with a score nearing the 90th percentile. While AI and automation can bring significant benefits to law firms and clients, they also raise important ethical and legal risks.

Lawyers must understand the technology they are using and ensure it is being used in a competent manner, including how to monitor, check, and assess common mistakes with the use of AI. For example, while an AI tool may be useful for outlining high level arguments and issue spotting, it may have significant shortcomings—and make express misstatements—regarding legal citations and references. In this way, the use of AI may be akin to the use of a nonlawyer assistant, which requires supervision. See Model Rule 5.3 cmt. [3].

Competence in the realm of AI includes discerning the many professional responsibility issues that can arise from its use. For example, lawyers should ensure they are meeting their duty of confidentiality when using AI and automation tools. See Model Rule 1.6. Confidential client information may be inadvertently disclosed through AI, especially a generative, open-source tool that ingests and re-uses content received across users. Even when the AI tool allegedly maintains confidentiality, the tool and its third-party provider must be vetted with reasonable scrutiny. Lawyers

should also consider Model Rule 5.5 and its prohibition on the unauthorized practice of law, to the extent that a sophisticated AI tool may qualify as the “practice of law.”

Further, excessive reliance on AI tools could lead to malpractice disputes. Model Rule 5.3 would conflict with a lawyer's ability to blame AI for incorrect or incomplete legal advice. Admittedly, however, the issue becomes more difficult if the client wanted the lawyer to use the AI tool to benefit from cost savings and was expressly advised of the tool's risks and shortcomings. See Model Rule 1.4(b) stating “A lawyer shall explain a matter to the extent reasonably necessary to permit the client to make informed decisions regarding the representation.”

As AI tools become more prevalent in the legal industry, tech competence will necessitate staying informed about the relevant ethical and legal issues implicated by AI. Lawyers should work to ensure that they are using AI and automation tools in a competent and ethical manner that protects client confidentiality and does not run afoul of professional responsibility rules.

Social Media & Software Applications

Social media platforms and software applications have become an important tool for lawyers to connect with clients, build professional networks, and promote their expertise. But the use of social media by lawyers also raises ethical and legal issues.

When it comes to the relationship between social media and tech competence, lawyers should understand the technology they are using and ensure they are using it appropriately. This includes understanding the privacy settings of social media platforms and taking appropriate measures to discover and investigate available evidence from these platforms. The outcome of a case could change completely based on the use—or non-use—of social media evidence related to the lawyers' client, the opposing party, and third parties.

For example, a competent attorney should generally understand how and where to gather vital evidence. With the prevalent use of mobile devices and software, attorneys may need to issue subpoenas to a wide variety of employers, phone carriers, data centers, and makers of social media applications. An attorney should be prepared to articulate the relevance, necessity, and scope of a subpoena for social media and app-collected data—wherever it may be stored—while also assuring the court that privacy concerns are adequately protected.

Another important issue in the competent use of social media is the duty of confidentiality. In line with Model Rule 1.6, lawyers should avoid disclosing confidential client information through their social media accounts. This may include information posted publicly, as well as private chat messages with third parties. It can also include excess commentary about an active case, appeal, or transaction, especially without client consent. The lawyers should also avoid and mitigate against

inadvertent disclosures through social media platforms that may store data on third-party servers. For instance, as stated in ABA Formal Ethics Op. 477R (2017), there are inherent risks in using electronic communications through “certain mobile applications or on message boards or via unsecured networks.”

When using social media, lawyers should also stay apprised of rules related to advertising and solicitation, which may vary across states:

- **New York Ethics Op. 972 (2013).** “A law firm may not list its services under the heading ‘Specialities’ on a social media site. A lawyer may not list services under that heading unless the lawyer is certified in conformity with the provisions of Rule 7.4(c).”
- **Ohio Ethics Op. 2013-2 (2013).** Treating text message marketing like direct mail rather than “real time” electronic communications; But “[b]ecause most text messages are received on cellular phones, which are often carried on one's person, lawyers should be sensitive to the fact that a text message may be perceived as more invasive than an email.”

Lawyers should also know about the risks of using social media to gather information about potential jurors, opposing parties, or witnesses. While social media can be a valuable tool for investigating and preparing for a case, lawyers should be careful not to engage in unethical or illegal behavior. For example, lawyers should not use fake social media profiles to gain access to private information, and should not attempt to contact potential witnesses or jurors through social media without following proper legal avenues:

- **San Diego Ethics Op. 2011-2 (2011).** Holding that a lawyer may not make a “friend request” to either an upper level executive of a corporate adversary—because the request is a “communication” about the subject matter of the representation—or even to an unrepresented person.
- **New York County Ethics Op. 745 (2013).** “Recent ethics opinions have concluded that accessing a social media page open to all members of a public network is ethically permissible. ... But an attorney's ability to access social media information is not unlimited. Attorneys may not make misrepresentations to obtain information that would otherwise not be obtainable.”

There are also practical considerations to keep in mind regarding social media. Lawyers should ensure that their profiles are professional and reflect positively on their practice. They should also be mindful of the time and resources required to maintain an active social media presence and should be careful not to engage in behavior that may reflect poorly on themselves or, especially, on their clients and ongoing representations.

Conclusion

Tech competence is an evolving concept, and lawyers should understand how new technology can impact their duties and professional obligations, including those of competence, confidentiality, and disclosure to the client. Tech competence involves awareness of potential benefits and risks of using—or not using—specific software or technology. By staying informed about relevant cases and ethical opinions, attorneys can make informed decisions and disclosures about how to best handle the wide range of emerging technology in today's digital age.

Bar Associations Begin to Tackle AI & the Practice of Law

Contributed by Amy Jane Longo, Shannon Capone Kirk, and Isaac Sommers, Ropes & Gray

December 2023

Since May 2023, a growing number of judges have issued standing orders regarding the use, generally, of artificial intelligence (AI) in litigation. We note that some of these guidances/advisory opinions are better than others at delineating between an overgeneralized use of the phrase “AI” and the item that should be of more concern to courts and regulators around “generative AI.”

Regulatory bodies and bar associations across the US are now taking initial but significant strides to address the ethical implications for attorneys, and to propose guidelines surrounding the use of AI in the legal profession. Whether those bodies will be precise in focusing on “generative AI,” will likely play out over the next year.

These initial moves, although narrower in scope, may set the trend for the legal industry's inevitable adoption of generative AI tools in everyday practice. States like California and Florida have begun their foray into providing guidance on the application of generative AI tools by focusing largely on how lawyers can apply existing ethical rules to such novel tools.

And such guidance may also be helpful for law firms and clients alike in crafting internal company policies regarding the application of generative AI tools, to be better positioned for proposed regulations and legislation down the road.

California Bar to Vote on AI Guidelines

Earlier this month, the California Bar's Committee on Professional Responsibility and Conduct (COPRAC) formally recommended that the Board of Trustees adopt recommendations regarding generative AI usage in legal practice. Specifically, COPRAC recommended that the Board:

- Develop a one-hour minimum continuing legal education course that addresses generative AI;
- Work with the state legislature and supreme court to evaluate the need for further regulations pertaining to the unauthorized practice of law and the use of legal generative AI products;
- Consider adding requirements regarding generative AI for law schools and the bar exam; and
- Publish a document called the “Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law.”

COPRAC's Practical Guidance is a four-page resource which explains how the existing Rules of

Professional Conduct “can be applied to generative AI use at this time,” and details how lawyers may use generative AI tools in a manner consistent with those rules. The Guidance includes suggestions that lawyers consider disclosing generative AI use to clients, review all generative AI outputs for accuracy, be aware of possible biases, among others.

Finally, COPRAC also intends to further study and potentially propose additional recommendations to the Board pertaining to using generative AI for the public good, how to supervise autonomous generative AI decision-making, whether the existing Rules should be updated to address generative AI, and what lawyers should be required to disclose to clients about the use of generative AI.

Although the recommendations are non-binding, the state bar association is also seeking to work with California's legislature and supreme court toward the eventual regulation of legal generative AI products. Depending on the breadth of future regulatory policies, such laws and rules could impact the legal industry in substantial ways, particularly as legal resource providers rush to integrate AI into their already existing suite of legal products.

Florida Bar Issues Proposed Advisory Opinion

On Nov. 13, 2023, the Florida Bar Association's Board of Governors' Review Committee on Professional Ethics issued Proposed Advisory Opinion (PAO) 24-1, and will consider any comments from bar members on Jan. 19, 2024. PAO 24-1 provides a variety of recommendations, including that attorneys obtain clients' consent before using third-party generative AI tools which may involve disclosure of confidential information, review and verify the accuracy of generative AI output, exercise caution when using generative AI-powered chatbots to handle client intake or advertising, and avoid overcharging for time spent using generative AI.

The Florida Bar also formed a Special Committee on AI Tools & Resources to further evaluate legal AI tools and their best uses in ways compliant with lawyers' ethical duties. However, unlike the recent PAO, which is specific to *generative* AI, the Special Committee does not have an explicitly generative AI-focused purpose. Rather, its mission statement speaks to AI in broad, general terms. Whether future proposals from the Special Committee will track the PAO's specificity to generative AI remains to be seen.

New York Bar Association Creates Task Force

The New York State Bar Association's Task Force on Artificial Intelligence, formed in July 2023, was created to delve into the intricacies of AI's role in the legal domain. Although the Task Force has yet to issue proposed guidelines, its mission is to proactively explore AI's intersection with the legal field. The Task Force will evaluate both positive and negative implications of the legal

community's use of AI, including generative AI tools specifically, but also encompassing other kinds of “AI-based software” and “machine learning tools.” Like California's, New York's Task Force will eventually propose and support legislation that regulates AI—although the nature of such legislation remains to be seen.

Other State Bar Associations & the ABA

In late summer of 2023, the American Bar Association (ABA) established the ABA Task Force on Law and Artificial Intelligence, dedicated to understanding the impact of AI on the legal profession and mitigating risks where necessary. This initiative recognizes the transformative nature of AI technologies and aims to navigate the challenges and opportunities they present. Like New York's Task Force, while the ABA Task Force's mission statement speaks of “AI” in broad terms, it has also acknowledged that unique legal challenges exist with regard to generative AI specifically.

By delving into the ethical dimensions of AI adoption, the ABA seeks to develop comprehensive guidelines that align with the evolving landscape of legal practice. In particular, the Task Force plans to evaluate how AI can impact lawyers' ethical duties, including how lawyers should handle possible confidentiality posed by AI tools, whether AI may create inadvertent waiver issues, how AI will impact attorneys' practice management, and even whether using AI tools “risks . . . the unauthorized practice of law.”

Other state bar associations are making similar moves.

New Jersey's State Bar Association created a task force to study how AI tools will affect the legal industry, with the Association's president predicting that the “end result” will be “guidelines and guidance” for attorneys regarding the use of AI. Other state bar associations in Illinois, Kentucky, Minnesota, and Texas have also formed working groups to examine legal and ethical issues relating to AI in the legal industry. A list of all such task forces identified by the ABA can be found [here](#). Whether these bar associations—and others—will seek to regulate AI broadly, or more prudently narrowly tailor proposed rules and policies to different kinds of AI—such as generative AI—will be determinative of how the legal industry develops with respect to artificial intelligence.

What's Next for Bar Associations?

These initiatives collectively underscore the legal community's recognition of the transformative potential of AI and the simultaneous commitment to ensuring its ethical and responsible integration into legal practice. As the legal landscape continues to evolve, these guidelines and frameworks—and others like them—will play an increasingly central role in navigating the intersection of law and AI.

It is foreseeable that bar associations and regulatory and lawmaking bodies will eventually work to

implement new rules and policies specifically tailored to generative AI. What is less clear is whether such organizations will consistently differentiate between AI in a broad sense and generative AI in particular.

The novelty and complexity of AI—and in particular generative AI and large language models—may prompt legislative responses from lawyers and legislators eager to establish regulatory frameworks. But if not guided by experts and lawyers who understand the technological nature of AI, and the differences between certain types of AI, such action could mirror recent moves by the SEC, which recently proposed broad rules that some fear could over-regulate or even stifle useful tools instead of carefully guiding innovation.

So far, however, at least California and Florida's intentional but cautious efforts seem to strike the right balance. California's and Florida's bar associations demonstrate a prudent, measured approach: prioritizing the application of existing rules specifically to new generative AI tools—as opposed to falling prey to overgeneralized concerns over “AI” as exhibited in the SEC's proposed rule from July 26, 2023.

Critically, many state bar associations are creating dedicated task forces to studying AI before making significant changes to ethics rules. This deliberate pace suggests a commitment to thoughtful and reasoned advancements in the integration of AI into legal practice, steering clear of impulsive regulatory measures that lack a comprehensive understanding of AI's nuances, which may well set the trend for the industry as a whole.

How Should Lawyers & Clients Respond?

Because generative AI legal tools are, at the end of the day, simply new tools that lawyers can employ for the benefit of their clients, it follows that many existing ethical rules and practices will apply to these new tools—even if that application may require some nuance. Accordingly, clients and attorneys should take particular note of bar associations' guidance on how to apply existing rules to new AI tools.

When considering how legal AI tools can be used to enhance business capabilities and outcomes, it will be important to simultaneously use practical guidance like the California or Florida bar associations' work to guide and craft the use of those AI tools. For example, a company that routinely uses legal AI tools to draft client-facing work should consider what kinds of disclosure policies it should adopt as part of that firm's regular practice.

Considering and implementing new internal policies may also better position the legal industry to provide substantive feedback and input when state bar associations eventually begin to propose actual legislative and regulatory reforms across the country. Firms that are proactive about

adopting clear internal policies and best practices regarding the use of legal AI tools will be among the best situated to provide experiential feedback and insight on eventual rulemaking and lawmaking around the US. The guidance from California and Florida provide an excellent starting point for organizations to begin that work.

ADOPTED

**AMERICAN BAR ASSOCIATION
CYBERSECURITY LEGAL TASK FORCE
ANTITRUST LAW SECTION
TORT, TRIAL & INSURANCE PRACTICE SECTION
SCIENCE & TECHNOLOGY LAW SECTION
STANDING COMMITTEE ON LAW AND NATIONAL SECURITY**

REPORT TO THE HOUSE OF DELEGATES

RESOLUTION

1 RESOLVED, That the American Bar Association urges organizations that design,
2 develop, deploy, and use artificial intelligence (“AI”) systems and capabilities to follow
3 these guidelines:
4

- 5 1) Developers, integrators, suppliers, and operators (“Developers”) of AI systems
6 and capabilities should ensure that their products, services, systems, and
7 capabilities are subject to human authority, oversight, and control;
8
9 2) Responsible individuals and organizations should be accountable for the
10 consequences caused by their use of AI products, services, systems, and
11 capabilities, including any legally cognizable injury or harm caused by their
12 actions or use of AI systems or capabilities, unless they have taken reasonable
13 measures to mitigate against that harm or injury; and
14
15 3) Developers should ensure the transparency and traceability of their AI products,
16 services, systems, and capabilities, while protecting associated intellectual
17 property, by documenting key decisions made with regard to the design and risk
18 of the data sets, procedures, and outcomes underlying their AI products,
19 services, systems and capabilities.
20

21 FURTHER RESOLVED, That the American Bar Association urges Congress, federal
22 executive agencies, and State legislatures and regulators, to follow these guidelines in
23 legislation and standards pertaining to AI.

REPORT

I. LEGAL ISSUES WITH AI

Artificial Intelligence (“AI”) systems and capabilities create significant new opportunities for technological innovation and efficiencies to benefit our society, but they also raise new legal and ethical questions. AI enables computers and other automated systems to perform tasks that have historically required human cognition, such as drawing conclusions and making predictions.¹ AI systems operate at much faster speeds than humans.²

With AI and machine learning (ML)³ already changing the way in which society addresses economic and national security challenges and opportunities, these technologies must be developed and used in a trustworthy and responsible manner. As private sector organizations and governments move rapidly to design, develop, deploy, and use AI systems and capabilities,⁴ now is a critical time for the American Bar Association (ABA) to articulate principles that are essential to ensuring that AI is developed and deployed in accordance with the law and well-accepted legal standards.⁵

¹ AI is not a single piece of hardware or software, but rather a constellation of technologies that give a computer system the ability to solve problems and to perform tasks that would otherwise require human intelligence. National Security Commission on Artificial Intelligence (NSCAI), *Final Report*, Artificial Intelligence in Context, pages 31-40, <https://www.nscai.gov/> [hereinafter “NSCAI Final Report”]. *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update* (Nov. 12, 2020), <https://catalog.data.gov/dataset/the-national-artificial-intelligence-research-and-development-strategic-plan-2019-update>.

According to the National Institute of Standards and Technology (NIST), AI is:

- (1) A branch of computer science devoted to developing data processing systems that performs functions normally associated with human intelligence, such as reasoning, learning, and self-improvement.
- (2) The capability of a device to perform functions that are normally associated with human intelligence such as reasoning, learning, and self-improvement.

NIST *U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools* (Aug. 2019), https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf.

² U.S. Government Accountability Office (GAO), *Artificial Intelligence: Status of Developing and Acquiring Capabilities for Weapons Systems*, GAO-22-104765 (Feb. 2022), <https://www.gao.gov/assets/gao-22-104765.pdf>. [hereinafter “GAO AI Report.”]

³ *Championing ethical and responsible machine learning through open-source best practices*, THE FOUNDATION FOR BEST PRACTICES IN MACHINE LEARNING, v. 1.0.0 (May 21, 2021), <https://www.nist.gov/system/files/documents/2021/08/18/ai-rmf-rfi-0010-attachment3.pdf>.

⁴ NSCAI Final Report at 28, *supra* note 1. (“We now know the uses of AI in all aspects of life will grow and the pace of innovation will accelerate.”)

⁵ This Resolution does not purport to alter lawyers’ obligations under applicable rules of professional conduct. Lawyers may wish to consider the issues raised in Daniel W. Linna Jr. and Wendy J. Muchma, *Ethical Obligations to Protect Client Data when Building Artificial Intelligence Tools: Wigmore Meets AI* (Oct. 2, 2020), https://www.americanbar.org/groups/professional_responsibility/publications/professional_lawyer/27/1/ethical-obligations-protect-client-data-when-building-artificial-intelligence-tools-wigmore-meets-ai/.

604

Fundamental concepts such as accountability, transparency, and traceability play an important role in ensuring the trustworthiness of AI systems. These concepts also play key roles in our legal system.⁶ This Resolution presents guidance on how the legal system and its participants, including attorneys, regulators, and stakeholders, such as developers, integrators, suppliers, and operators (“developers”) of AI systems and capabilities, should assess these fundamental issues with AI. It states that in the context of AI, individual and enterprise accountability and human authority, oversight, and control are required and it is not appropriate to shift legal responsibility to a computer or an “algorithm” rather than to responsible people and other legal entities.

This Resolution will ensure that courts and participants in the legal process have the capacity to evaluate and resolve legal questions and disputes by specifying the essential information that must be included in the design, development, deployment, and use of AI to ensure transparency and traceability.

By focusing on these principles related to AI, this Resolution will help to ensure that accountability, transparency, and traceability are built into AI products, services, systems, and capabilities “by design” from the beginning of the development process. Following the proposed guidelines will enhance AI by maximizing the benefits from the use of AI in a trustworthy and responsible manner and help to minimize the risks.

Further, the Resolution urges Congress, federal executive agencies, and State legislatures and regulators to follow the guidelines in legislation and standards pertaining to AI.

II. OVERVIEW OF AI

AI holds great potential to bring innovation and efficiency across a number of industry sectors. New AI-enabled systems are benefitting many parts of society and the economy, from commerce and healthcare to transportation and cybersecurity. Consider just a few examples of recent AI innovations:

- Artificial intelligence is being deployed as a dialog agent for customer service. Several of these efforts have passed the Turing test – the eponymous idea developed by early computer pioneer Alan Turing which posited that the true test of computer intelligence will be met when individuals cannot tell the difference between a computer and a human interaction;
- Self-driving cars are under wide development by virtually every major manufacturer in the world (as well as most of the larger tech companies). While they are still in the testing stage, there is every reason to anticipate that geofenced cars will be on the market within 5-10 years;

⁶ Other important legal issues with AI have been identified, such as intellectual property infringement, algorithmic bias, access to justice, fairness in decision-making, discrimination, unfairness, and privacy and data protection/ cybersecurity. These issues may be appropriate for future ABA resolutions.

- The AI product named Watson defeated the human champion in a game of Jeopardy and one named Alpha Go defeated the world Go champion;
- A system known as Deep Patient is now being deployed, successfully, as a diagnostic assistant to clinicians in a hospital setting, helping them make improved diagnoses in difficult cases. It is capable of predicting the onset of certain psychological diseases like schizophrenia in situations where the symptoms are not apparent to human clinicians;
- An artwork created by AI recently sold for over \$400,000 at auction;
- More than two years ago a TV station in China began using an AI-powered announcer as the news anchor;
- Recent tests of autonomous self-directed weapons systems have successfully demonstrated that military systems can identify and target adversaries without human intervention; and
- New AI programs that go by the generic name of Deep Fakes can create fake video that can be virtually indistinguishable from reality.

Recently, governments and other organizations have been working on proposed AI governance frameworks and principles with the goal of mitigating the risks that can result through implementation of AI systems and capabilities. For example, NIST has developed an AI Risk Management Framework to provide guidance regarding the trustworthiness of AI systems.⁷ Specifically, the framework is intended to help incorporate trustworthiness considerations into the design, development, use, and evaluation of AI systems, and it highlights accountability and transparency as two key guiding principles.⁸

The White House Office of Science and Technology Policy (OSTP) has acknowledged the “extraordinary promise of AI” as well as its pitfalls, and the need to “advance development, adoption, and oversight of AI in a manner that aligns with our democratic values.”⁹ In recognition of the importance of ensuring that the American public has appropriate protections in the age of AI, OSTP released its Blueprint for an AI Bill of Rights “for building and deploying automated systems that are aligned with democratic values and protect civil rights, civil liberties, and privacy.”¹⁰ OSTP explained:

⁷ NIST *AI Risk Management Framework: Second Draft* (August 2022), https://www.nist.gov/system/files/documents/2022/08/18/AI_RMF_2nd_draft.pdf. [hereinafter “NIST AI Risk Management Framework”].

⁸ *Id.* at 13.

⁹ L. Parker and R. Richardson, *OSTP’s Continuing Work on AI Technology and Uses That Can Benefit Us All*, OSTP Blog (Feb. 3, 2022), <https://www.whitehouse.gov/ostp/news-updates/2022/02/03/ostps-continuing-work-on-ai-technology-and-uses-that-can-benefit-us-all/>.

¹⁰ White House Office of Science and Technology Policy (OSTP), *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (October 2022), <https://www.whitehouse.gov/wp->

Our country should clarify the rights and freedoms we expect data-driven technologies to respect. What exactly those are will require discussion, but here are some possibilities: your right to know when and how AI is influencing a decision that affects your civil rights and civil liberties; your freedom from being subjected to AI that hasn't been carefully audited to ensure that it's accurate, unbiased, and has been trained on sufficiently representative data sets; your freedom from pervasive or discriminatory surveillance and monitoring in your home, community, and workplace; and your right to meaningful recourse if the use of an algorithm harms you.¹¹

III. **ACCOUNTABILITY AND HUMAN OVERSIGHT, AUTHORITY, AND CONTROL**

The ABA urges organizations that design, develop, deploy, and use AI systems and capabilities to follow these guidelines:

- Developers, integrators, suppliers, and operators (“developers”) of AI systems and capabilities should ensure that their products, services, systems, and capabilities are subject to human authority, oversight, and control.
- Responsible individuals and enterprises should be accountable for the consequences caused by their use of AI products, services, systems, and capabilities, including any legally cognizable injury or harm caused by their use, unless they have taken reasonable measures to mitigate against that harm or injury.

Accountability and human authority, oversight and control are closely interrelated legal concepts. In the context of AI, they present key concerns, given that AI is increasingly being used in a variety of contexts to make decisions that can significantly impact people's lives, including evaluating applicants for jobs, determining who receives access to loans, assessing criminal defendants' likelihood of being a repeat offender in connection with bail proceedings, screening rental applicants, and determining how self-driving cars should navigate through complex traffic and driving situations.

The Defense Advanced Research Projects Agency (DARPA) recently announced that it is starting a program to evaluate the use of AI to make complex decisions in modern military operations. DARPA explained that this In the Moment (ITM) program “aims to

<content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>. The Blueprint focuses on five principles for automated decision-making systems: (1) Safe and effective systems; (2) Algorithmic discrimination protections; (3) Data privacy; (4) Notice and explanation; and (5) Human alternatives, consideration and fallback.

¹¹ E. Lander & A. Nelson, *ICYMI: WIRED (Opinion): Americans Need a Bill of Rights For An AI-Powered World*, OTSP Blog (Oct. 22, 2022), <https://www.whitehouse.gov/ostp/news-updates/2021/10/22/icymi-wired-opinion-americans-need-a-bill-of-rights-for-an-ai-powered-world/>. See, Ben Winters, *AI Bill of Rights Provides Actionable Instructions for Companies, Agencies, and Legislators*, EPIC (Oct. 11, 2022), <https://epic.org/ai-bill-of-rights-leaves-actionable-instructions-for-companies-agencies-and-legislators/>.

evaluate and build trusted algorithmic decision-makers for mission-critical Department of Defense (DoD) operations.”¹²

Various organizations have recognized the importance of accountability with AI systems. In its draft AI Risk Management Framework, NIST stated that:

Determinations of accountability in the AI context are related to expectations for the responsible party in the event that a risky outcome is realized. Individual human operators and their organizations should be answerable and held accountable for the outcomes of AI systems, particularly adverse impacts stemming from risks.¹³

The Organization for Economic Cooperation and Development (OECD) Principles for AI includes Principle 1.5 on Accountability, which provides:

Organizations and individuals developing, deploying or operating AI systems should be held accountable for their proper functioning in line with the OECD’s values-based principles for AI.¹⁴

Australia has issued a voluntary framework of eight AI Ethics Principles which includes accountability, stating:

People responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled.¹⁵

In addition, large technology companies have also recognized the importance of accountability with regard to their AI products. For example, one of Microsoft’s Six Principles for Responsible AI is accountability: “people should be accountable for AI systems.”¹⁶ Similarly, Google includes accountability in its Objectives for AI Applications, and states that AI should “be accountable to people. We will design AI systems that provide appropriate opportunities for feedback, relevant explanations, and appeal. Our AI technologies will be subject to appropriate human direction and control.”¹⁷

Human accountability is of particular importance given that with ML, a subset of AI, computers are able to learn from data sets without being given explicit instructions from

¹² *Developing Algorithms That Make Decisions Aligned With Human Expert*, DARPA Notice (March 3, 2022), <https://www.darpa.mil/news-events/2022-03-03>.

¹³ NIST AI Risk Management Framework at 13, *supra* note 5.

¹⁴ OECD AI Principles, <https://oecd.ai/en/dashboards/ai-principles/P7>. [hereinafter “OECD AI Principles.”]

¹⁵ *Australia’s AI Ethics Principles, Principles at a Glance*, <https://www.industry.gov.au/data-and-publications/australias-artificial-intelligence-ethics-framework/australias-ai-ethics-principles>.

¹⁶ Microsoft *Responsible AI principles in practice*, <https://www.microsoft.com/en-us/ai/responsible-ai?activetab=pivot1%3aprimar6>, [hereinafter “Microsoft *Responsible AI Principles*”].

¹⁷ Google *AI Principles*, <https://ai.google/principles/>.

604

humans. Instead, the computer model learns from experience and trains itself to find patterns and make predictions.¹⁸ There has been widespread recognition of the critical role that humans should play in overseeing and implementing AI systems that are making such important decisions. For example, the term “human-centered artificial intelligence” has been used to describe the view that AI systems “must be designed with awareness that they are part of a larger system consisting of human stake-holders, such as users, operators, clients, and other people in close proximity.”¹⁹

Accountability is important given the increasing concern about understanding AI decision-making and ensuring fairness in AI models, including with regard to the potential discriminatory impact of certain AI systems. For example, Amazon started a program to automate hiring by using an algorithm to review resumes. However, the program had to be discontinued after it was discovered that it discriminated against women in certain technical positions, such as software engineer, because the software analyzed the credentials of its existing employee base, which was predominantly male.²⁰ In addition, researchers found a gender and skin-type bias with commercial facial analysis programs, with an error rate of 0.8 percent for light-skinned men, versus 34.7 for dark-skinned women.²¹

There have been recent efforts to prohibit AI systems from violating anti-discrimination and privacy laws. For example, the Equal Employment Opportunity Commission (EEOC) launched an initiative to ensure that AI used in hiring and other employment decisions does not violate anti-discrimination laws.²² New York City passed a new law to take effect in 2023 that prohibits the use of AI machine learning products in hiring and promotion decisions unless the tools have first been audited for bias.²³ In 2018, California passed the California Consumer Privacy Act (CCPA), a consumer protection law intended to protect the privacy of California residents. In 2020, it passed the California Privacy Rights Act (CPRA), amending the CCPA to add measures including the right to limit use and disclosure of sensitive personal information and the right to obtain information about how companies use automated decision-making technology.²⁴

¹⁸ S. Brown, *Machine Learning Explained*, MIT Management: Ideas Made to Matter (April 21, 2021), <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>.

¹⁹ M. Riedl, *Human-Centered Artificial Intelligence and Machine Learning*, arXiv:1901.11184[cs.AI].

²⁰ J. Dastin, *Amazon Scraps Secret AI Recruiting Tool That Shows Bias Against Women*, Reuters (Oct. 10, 2018), <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

²¹ L. Hardesty, *Study Finds Gender and Skin-Type Bias in Commercial Artificial Intelligence Systems*, MIT NEWS (Feb. 11, 2018), <https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-021>.

²² *EEOC Launches Initiative on Artificial Intelligence and Algorithmic Fairness*, EEOC press release (Oct. 28, 2021), <https://www.eeoc.gov/newsroom/eeoc-launches-initiative-artificial-intelligence-and-algorithmic-fairness>.

²³ N. Lee and S. Lai, *Why New York City Is Cracking Down on AI in Hiring*, BROOKINGS TECHTANK (Dec. 20, 2021), <https://www.brookings.edu/blog/techtank/2021/12/20/why-new-york-city-is-cracking-down-on-ai-in-hiring/>.

²⁴ B. Justice, *CPRA Countdown: It's Time to Brush Up on California's Latest Data Privacy Law*, NATIONAL LAW REVIEW (Dec. 18, 2021), <https://www.natlawreview.com/article/cpra-countdown-it-s-time-to-brush-california-s-latest-data-privacy-law>.

In addition, questions have also been raised about the protection of privacy because of the processing of personal data in AI systems.²⁵

Existing laws and regulations can be used to prevent potential violations of anti-discrimination and privacy laws by AI systems. For example, Federal Trade Commission (FTC) Commissioner Rebecca Kelly Slaughter explained her view that the FTC's existing tools, including section 5 of the FTC Act, the Equal Credit Opportunity Act, the Fair Credit Reporting Act, and the Children's Online Privacy Protection Act, can and should be used to protect consumers against algorithmic harms.²⁶

In light of the need to ensure compliance with laws and regulations being used to prevent harms from AI systems, it is essential that the humans and enterprises with responsibility for these AI systems be held accountable for the consequences of the uses of these systems.

Under our legal system, in order to be held accountable, an entity must have a specific legal status that allows it to be sued, such as being an individual human or a corporation. On the other hand, property, such as robots or algorithms, does not have a comparable legal status.²⁷ Thus, it is important that legally recognizable entities such as humans and corporations be accountable for the consequences of AI systems, including any legally cognizable injury or harm that their actions or those of the AI systems or capabilities cause to others, unless they have taken reasonable measures to mitigate against that harm or injury.²⁸

IV. TRANSPARENCY AND TRACEABILITY

The ABA urges organizations that design, develop, deploy, and use artificial intelligence ("AI") products, services, systems and capabilities to follow this guideline:

- Developers should ensure the transparency and traceability of their AI products, services, systems, and capabilities, while protecting associated intellectual property, by documenting key decisions made with regard to the design and risk of the data sets, procedures, and outcomes underlying their AI products, services, systems, and capabilities.

²⁵ C. Tucker, *Privacy, Algorithms and Artificial Intelligence*, in *The Economics of Artificial Intelligence: An Agenda*, NATIONAL BUREAU OF ECONOMIC RESEARCH (2019), <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda/privacy-algorithms-and-artificial-intelligence>.

²⁶ R. Slaughter, *Algorithms and Economic Justice*, ISP DIGITAL FUTURE WHITEPAPER & YALE JOURNAL OF LAW & TECHNOLOGY SPECIAL PUBLICATION (Aug. 2021)

²⁷ Michalski, Roger (2018), *How to Sue a Robot*, UTAH LAW REVIEW: Vol. 2018: No. 5, Article 3, <https://dc.law.utah.edu/ulr/vol2018/iss5/3>.

²⁸ In developing rules of liability, the supplier/component part doctrine would apply. Under that doctrine, the manufacturer of a non-defective component is not liable for harm caused by a defect in a larger system sold by a manufacturer into which the component was integrated.

604

A. Transparency

In the context of AI, transparency is about responsible disclosure to ensure that people understand when they are engaging with an AI system, product, or service and enable those impacted to understand the outcome and be able to challenge it if appropriate.²⁹ NIST stated that “explainable AI” is one of several properties that characterize trust in AI systems.³⁰

Lack of transparency with AI can negatively affect individuals who are denied jobs, refused loans, refused entry or are deported, imprisoned, put on no-fly lists or denied benefits. They are often not informed of the reasons other than the decision was processed using computer software. Human rights principles that may be impacted are rights to a fair trial and due process, effective remedies, social rights and access to public services, and rights to free elections.³¹

OECD has explained that the term transparency carries multiple meanings:

In the context of this Principle [1.3], the focus is first on disclosing when AI is being used (in a prediction, recommendation or decision, or that the user is interacting directly with an AI-powered agent, such as a chatbot). Disclosure should be made with proportion to the importance of the interaction. The growing ubiquity of AI applications may influence the desirability, effectiveness or feasibility of disclosure in some cases.

²⁹ OECD adopted Transparency and Explainability Principle 1.3 that states:

AI Actors should commit to transparency and responsible disclosure regarding AI systems. To this end, they should provide meaningful information, appropriate to the context, and consistent with the state of art:

- to foster a general understanding of AI systems,
- to make stakeholders aware of their interactions with AI systems, including in the workplace,
- to enable those affected by an AI system to understand the outcome, and,
- to enable those adversely affected by an AI system to challenge its outcome based on plain and easy-to-understand information on the factors, and the logic that served as the basis for the prediction, recommendation or decision.

OECD AI Principles, *supra* note 12.

³⁰ NIST *Artificial Intelligence*, <https://www.nist.gov/artificial-intelligence>; NIST *Four Principles of Explainable Artificial Intelligence*, NIST Interagency/Internal Report (NISTIR) - 8312, <https://doi.org/10.6028/NIST.IR.8312>.

Four principles of explainable AI – for judging how well AI decisions can be explained:

- *Explanation* – AI systems should deliver accompanying evidence or reasons for all their outputs.
- *Meaningful* – Systems should provide explanations that are meaningful or understandable to individual users.
- *Explanation Accuracy* – The explanation correctly reflects the system’s process for generating the output.
- *Knowledge Limits* – The system only operates under conditions for which it was designed or when the system reaches a sufficient confidence in its output. (The idea is that if a system has insufficient confidence in its decision, it should not supply a decision to the user.)

See, <https://www.nist.gov/artificial-intelligence/ai-fundamental-research-explainability>.

³¹ Rowena Rodrigues, *Legal and human rights issues of AI: Gaps, challenges and vulnerabilities*, JOURNAL OF RESPONSIBLE TECHNOLOGY, Vol. 4, Dec. 2020, 100005, <https://doi.org/10.1016/j.jrt.2020.100005>.

Transparency further means enabling people to understand how an AI system is developed, trained, operates, and deployed in the relevant application domain, so that consumers, for example, can make more informed choices. Transparency also refers to the ability to provide meaningful information and clarity about what information is provided and why. Thus transparency does not in general extend to the disclosure of the source or other proprietary code or sharing of proprietary datasets, all of which may be too technically complex to be feasible or useful to understanding an outcome. Source code and datasets may also be subject to intellectual property, including trade secrets.

An additional aspect of transparency concerns facilitating public, multi-stakeholder discourse and the establishment of dedicated entities, as necessary, to foster general awareness and understanding of AI systems and increase acceptance and trust.

Numerous organizations around the world have developed AI principles. A researcher who reviewed them reported that “[f]eatured in 73/84 sources, transparency is the most prevalent principle in the current literature.”³² Varied terminology is used to express this concept of transparency, comprising efforts to increase explainability, interpretability, intelligibility or other acts of communication and disclosure.

Intelligibility can uncover potential sources of unfairness, help users decide how much trust to place in a system, and generally lead to more usable products. It also can improve the robustness of AI systems by making it easier for data scientists and developers to identify and fix bugs.³³

The FTC published guidance regarding the commercial use of AI technology, acknowledging that while AI has significant positive potential, it also presents negative risks, such as unfair or discriminatory outcomes or the entrenchment of existing disparities.³⁴ The FTC urged companies to:

- Be transparent with consumers;
- Explain how algorithms make decisions;

³² Anna Jobin, *et. al.*, *Artificial Intelligence: the global landscape of ethics guidelines*, HEALTH ETHICS & POLICY LAB, ETH Zurich, 8092 Zurich, Switzerland (2019), https://www.researchgate.net/profile/Anna-Jobin/publication/334082218_Artificial_Intelligence_the_global_landscape_of_ethics_guidelines/links/5d19ec7d299bf1547c8d2be8/Artificial-Intelligence-the-global-landscape-of-ethics-guidelines.pdf?origin=publication_detail.

European Union member state reports on AI can be found at <https://futurium.ec.europa.eu/en/european-ai-alliance/pages/official-documents-and-reports>.

³³ Microsoft *Responsible AI principles*, *supra* note 14. Microsoft Research Collection: *Research Supporting Responsible AI* (April 13, 2020), <https://www.microsoft.com/en-us/research/blog/research-collection-research-supporting-responsible-ai/>.

³⁴ FTC *Using Artificial Intelligence and Algorithms* (April 8, 2020), <https://www.ftc.gov/business-guidance/blog/2020/04/using-artificial-intelligence-algorithms>; FTC, *Aiming for truth, fairness, and equity in your company's use of AI* (April 19, 2021), <https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>.

604

- Ensure that decisions are fair, robust, and empirically sound; and
- Hold themselves accountable for compliance, ethics, fairness and non-discrimination.

B. Traceability

It is important to ensure that the complex processes in data science — from data processing through modeling to deployment in production — can be documented in a way that is understood easily.³⁵ Traceability is considered a key requirement for trustworthy AI. It would allow companies to better understand the entire reasoning process, and builds trust with AI implementations.³⁶

According to NIST, “[t]rustworthy AI refers to AI capabilities that exhibit characteristics such as resilience, security, and privacy so that relevant people can adopt them without fear.”³⁷ An AI capability must be traceable, meaning that it is developed and deployed such that relevant personnel possess an appropriate understanding of the technology, development processes, and operational methods applicable to AI capabilities, including with transparent and auditable methodologies, data sources and design procedures and documentation.³⁸

³⁵ Andreas Gödde, *Traceability for Trustworthy AI: A Review of Models and Tools*, SAS, <https://www.mdpi.com/2504-2289/5/2/20/htm>.

<https://blogs.sas.com/content/hiddeninsights/2018/03/12/interpretability-traceability-clarity-ai-mandate/>.

See, Association for Computing Machinery, *Outlining Traceability: A Principle for Operationalizing Accountability in Computing Systems*, FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (March 2021), pages 758–771,

<https://dl.acm.org/doi/10.1145/3442188.3445937>.

³⁶ Sanjay Srivastava, *The path to explainable AI*, CIO (May 21, 2018),

<https://www.cio.com/article/221668/the-path-to-explainable-ai.html>.

³⁷ NIST, Draft – Taxonomy of AI Risk (Oct. 2021),

https://www.nist.gov/system/files/documents/2021/10/15/taxonomy_AI_risks.pdf; see GAO AI Report, *supra* note 2.

³⁸ The Department of Defense (DoD) adopted *5 Principles of Artificial Intelligence Ethics* that commits the Department to this principle of traceability. U.S. Department of Defense, *5 Principles of Artificial Intelligence Ethics*, <https://www.defense.gov/News/News-Stories/Article/Article/2094085/dod-adopts-5-principles-of-artificial-intelligence-ethics/>. See *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence* by the Department of Defense, Defense Innovation Board, available at https://media.defense.gov/2019/Oct/31/2002204458/-1/-/1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF.

Similarly, the *Principles of Artificial Intelligence Ethics for the Intelligence Community*³⁸ provide:

Transparent and Accountable – We will provide appropriate transparency to the public and our customers regarding our AI methods, applications, and uses within the bounds of security, technology, and releasability by law and policy, and consistent with the Principles of Intelligence Transparency for the IC. We will develop and employ mechanisms to identify responsibilities and provide accountability for the use of AI and its outcomes.

C. Documenting key decisions made with regard to the design and risk of the data sets, procedures, and outcomes.

As AI algorithms become more complex, the need for greater transparency grows. Experts are developing software tools that will address the “black box” problem³⁹ – not knowing how algorithms arrive at their final output – by analyzing complex AI systems and documenting how the system processes information, answers questions, and provides results.⁴⁰

Traceability is related to the need to maintain a complete account of the provenance of data, processes, and artifacts involved in the production of an AI model – and it should encompass all elements of an AI system, product or service, namely the data, the system, and the business model. It requires documentation of the data sets, procedures, and outcomes for the AI system or capability.⁴¹

Practical Considerations – In establishing traceability for AI products, services, systems, and capabilities, developers should create contemporaneous records that document key decisions made with regard to the design and risk of the AI data sets. This means using automated tools when appropriate and available, or otherwise using documentation techniques (online or manual) appropriate for the software development lifecycle and for conducting AI risk assessments. Computer scientists are developing data models and tools to fully document data, procedures and outcomes for AI systems. They enable some form of automated repetition of the construction of the artifacts.⁴²

Examples of the types of key decisions to be documented throughout the AI lifecycle include:

- *Business* – business-oriented requirements, expected uses and outcomes, key performance features (including when AI is used or relied upon in decision making). Human control over the selection of inputs and generation of outputs in order to reduce the risks of unintended adverse consequences.

³⁹ Cliff Kuang, *Can A.I. Be Taught to Explain Itself?* THE NEW YORK TIMES MAGAZINE (Nov. 21, 2017), <https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html>

⁴⁰ Neil Savage, *Breaking into the black box of artificial intelligence: Scientists are finding ways to explain the inner workings of complex machine-learning models*, NATURE (Mar. 29, 2022), <https://www.nature.com/articles/d41586-022-00858-1>.

⁴¹ The assessment for traceability includes:

- *Procedures*: Methods used for designing and developing the algorithmic system: how the algorithm was trained, which input data was gathered and selected, and how this occurred.
- *Data*: Methods used to test and validate the algorithmic system: information about the data used to test and validate.
- *Outcomes*: The outcomes of the algorithms or the subsequent decisions taken on the basis of these outcomes, as well as other potential decisions that would result from different cases (e.g., for other subgroups of users).

⁴² *Traceability for Trustworthy AI: A Review of Models and Tools*, <https://www.mdpi.com/2504-2289/5/2/20/htm>.

604

- *Data* – types, quantities, and sources of data to be used in training the AI systems and capabilities; modeling, analysis, evaluation.⁴³
- *AI risk assessment* – risks assessed, unintended bias, or hazardous use.
- *Cybersecurity risks* – risks of unauthorized access to, and compromise of the integrity of, the AI algorithms, software, training data, and/or model.
- *Design and development* – key design trade-offs, risks mitigated by the design. Review of algorithm(s), software code and the AI model.
- *Testing* – involvement of humans with detailed understanding of AI processes and industry domain issues. Testing of implementing software, model with data sets, and adjustments and correction of errors. Problems observed in generating desired outputs. Performance deficiencies, malfunctions, unintended outputs, and discovered risks observed.
- *Deployment*
- *Developers should respond promptly* to avert or mitigate AI risks that are identified at any point in the AI system/product life cycle.

In the event of a gap between actual and desired performance with an AI system, capability, product, or service, recurring errors or failures with specific processes and undesirable events reoccurring, traceability will enable root cause analysis, a process for understanding 'what happened' and solving a problem through looking back and drilling down to find out 'why it happened' in the first place. Then, looking to rectify the issue(s) so that it does not happen again, or reduce the likelihood that it will happen again.⁴⁴

The many benefits of root cause analysis include reducing risk and preventing recurring failures, improving performance, as well as the potential for cost reduction. It provides a logical approach to problem solving using data that already exist and a learning process for better understanding of relationships, causes and effect, and solutions. The process should lead to more robust AI systems and capabilities.

V. EXISTING ABA POLICY

The ABA House of Delegates passed two Resolutions that address AI. This Resolution builds on and is consistent with those existing ABA policies.

- ABA urges courts and lawyers to address the emerging ethical and legal issues related to the usage of artificial intelligence (“AI”) in the practice of law, including (1) bias, explainability, and transparency of automated decisions made by AI; (2) ethical and beneficial usage of AI; and (3) controls and oversight of AI and the vendors that provide AI. 19A112.

⁴³ The key is to fully understand the data’s behavior. Best practices include documenting assumptions around completeness of the data, addressing data biases, and reviewing new rules identified by the machine before implementing. If AI is being used to identify anomalies, companies can put checks and balances in place to manually test and determine if the results make sense.

⁴⁴ Chartered Institute of Internal Auditors, *Root Cause Analysis* (Sept. 22, 2020), <https://www.iaa.org.uk/resources/delivering-internal-audit/root-cause-analysis?downloadPdf=true>.

- ABA urges federal, state, local, territorial and tribal governments to:
 - Ensure due process and refrain from using pretrial risk assessment tools unless the data supporting the risk assessment is transparent, publicly disclosed, and validated; and
 - Recognize that an individual’s criminal history and other criteria may reflect structurally biased application of laws, policies or practices, as well as conscious or unconscious bias. 22M700.

V. CONCLUSION

This Resolution addresses important legal issues concerning AI by focusing on the principles of accountability, transparency and traceability. It states that in the context of AI, human and enterprise accountability and human authority, oversight, and control are required and it is not appropriate to shift legal responsibility to a computer or an “algorithm” rather than to responsible people and other legal entities.

It will ensure that courts and participants in the legal process have the capacity to evaluate and resolve legal questions and disputes by specifying the essential information that must be included in the design, development, deployment, and use of AI to ensure transparency and traceability. Passage of this Resolution will enhance AI by maximizing the benefits from the use of AI in a trustworthy and responsible manner and help to minimize the risks.

Respectfully Submitted,

Claudia Rast and Maureen Kelly, Co-Chairs
Cybersecurity Legal Task Force

February 2023

APPENDIXLAWS, COURT DECISIONS, AND LEADING REPORTS

An exhaustive analysis of federal, state, and international laws applicable to AI is outside the scope of this Report. Below are some of the highlights:

National Conference of State Legislatures (NCLS) *State AI Legislation*

<https://www.ncsl.org/research/telecommunications-and-information-technology/2020-legislation-related-to-artificial-intelligence.aspx>

General AI bills or resolutions were introduced in at least 17 states in 2021-22, and were enacted in Alabama, Colorado, Illinois, Mississippi, Vermont, and Washington.

General Data Protection Regulation (GDPR) Article 22 – AI Requirements⁴⁵

GDPR imposes legal requirements on whoever uses an AI system for profiling and/or automated decision-making (regardless of the *means* by which personal data are processed), even if they acquired the system from a third party. These requirements include Fairness; Transparency, including meaningful information about the logic involved in the AI system; and the right to human intervention, enabling the individual to challenge the automated decision.

European Commission, Proposal for a Regulation on Artificial Intelligence (April 2021)

COM/2021/206 final (Document 52021PC0206), <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>.

The Regulation introduces new obligations for vendors of AI systems, and includes requirements for high-risk AI systems and users.

European Parliament, The impact of the General Data Protection Regulation (GDPR) on artificial intelligence, PE 641.530 (June 2020),

[https://www.europarl.europa.eu/RegData/etudes/STUD/2020/641530/EPRS_STU\(2020\)641530_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2020/641530/EPRS_STU(2020)641530_EN.pdf).

***Holbrook v. Prodomax Automation Ltd.*, 2021 U.S. Dist. LEXIS 178325 (Sept. 20, 2021) U.S. Dist. Ct., W.D. Mich.**

⁴⁵ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance) (OJ L 119 04.05.2016, p. 1, CELEX: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679>).

Man Whose Wife Was Killed by Factory Robot Settles Mid-Trial, BLOOMBERG (Nov. 9, 2021), <https://news.bloomberglaw.com/product-liability-and-toxics-law/man-whose-wife-was-killed-by-factory-robot-settles-mid-trial>.

Eric L. Alexander, *Unintended Consequences for Software Liability?* REED SMITH (Nov. 26, 2021), <https://www.lexology.com/library/detail.aspx?g=54e4a579-500d-4db0-adc2-065bc9b06263>.

Leading Reports

White House Office of Science and Technology Policy (OSTP), *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (October 2022)

<https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

The Blueprint focuses on principles for automated decision-making systems: (1) Safe and effective systems; (2) Algorithmic discrimination protections; (3) Data privacy; (4) Notice and explanation; and (5) Human alternatives, consideration and fallback.

National Security Commission on Artificial Intelligence (NSCAI), *Final Report*

<https://www.nscai.gov/>.

Presents the strategy for the U.S. to win in the AI era by responsibly using AI for national security and defense, defending against AI threats, and promoting AI innovation. *Blueprints for Action* provide plans to implement the recommendations.

House Committee on Transportation and Infrastructure

Boeing 737 MAX Investigation, <https://transportation.house.gov/committee-activity/boeing-737-max-investigation>.

Final Committee Report on the Design, Development, and Certification of the Boeing 737 MAX (Sept. 2020).

NIST AI Risk Management Framework: Second Draft (August 2022)

https://www.nist.gov/system/files/documents/2022/08/18/AI_RMF_2nd_draft.pdf.

Intended for voluntary use “in addressing risks in the design, development, use, and evaluation of AI products, services, and systems.”

Artificial Intelligence and the Courts: Materials for Judges, American Association for the Advancement of Science (AAAS) (Sep. 2022)

<https://www.aaas.org/ai2/projects/law/judicialpapers>.

With the support of NIST, this AAAS project is developing resources to support judges as they address an increasing number of cases involving AI.

604

Stanford HAI, *Artificial Intelligence Index Report 2021*, Stanford Human-Centered Artificial Intelligence

https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf.

Presents unbiased, globally sourced data that will enable policy-makers, researchers, executives, and the public to develop intuitions about AI.

Industry IoT Consortium, *Industrial IoT Artificial Intelligence Framework* (Feb. 22, 2022)

<https://www.iiconsortium.org/pdf/Industrial-AI-Framework-Final-2022-02-21.pdf>.

Provides guidance in the development, training, documentation, communication, integration, deployment, and operation of AI-enabled industrial IoT systems.

OECD AI Principles (May 2019)

<https://oecd.ai/en/ai-principles>.

Promotes the use of innovative and trustworthy AI and respects human rights and democratic values.

European Commission, *European AI Alliance*

<https://futurium.ec.europa.eu/en/european-ai-alliance/pages/official-documents-and-reports>.

Council of Europe, Karen Yeung, *Responsibility and AI*, DGI(2019)05

<https://rm.coe.int/responsability-and-ai-en/168097d9c5>.

A study of the implications of advanced digital technologies (including AI systems) for the concept of responsibility within a human rights framework.

Katherine B. Forrest, *When Machines Can Be Judge, Jury, And Executioner: Justice In The Age Of Artificial Intelligence* (2021)

GENERAL INFORMATION FORM

Submitting Entity: Cybersecurity Legal Task Force

Submitted By: Claudia Rast and Maureen Kelly, Co-chairs

1. Summary of Resolution(s).

This Resolution presents guidance on how the legal system and its participants, including attorneys, regulators, and stakeholders – developers, integrators, suppliers, and operators (“developers”) of AI systems and capabilities – should assess three fundamental issues with AI: accountability, transparency and traceability.

The Resolution will ensure that courts and participants in the legal process have the capacity to evaluate and resolve legal questions and disputes by specifying the essential information that must be included in the design, development, deployment, and use of AI to ensure transparency and traceability.

2. Indicate which of the ABA’s four goals the resolution seeks to advance (1-Serve our Members; 2-Improve our Profession; 3-Eliminate Bias and Enhance Diversity; 4-Advance the Rule of Law) and provide an explanation on how it accomplishes this.

This Resolution meets Goal 4 – Advance the Rule of Law. The Resolution is designed to help mitigate the risks that can result through implementation of AI systems and capabilities and enhance the use of AI in a trustworthy and responsible manner.

3. Approval by Submitting and Co-sponsoring Entities.

The Cyberspace Legal Task Force voted to sponsor this Resolution on December 2, 2022.

The Antitrust Law Section voted to co-sponsor this Resolution on December 2, 2022.

The Tort, Trial & Insurance Practice (TIPS) Section voted to co-sponsor this Resolution on November 16, 2022.

The Science & Technology Law Section voted to co-sponsor this Resolution on December 20, 2022.

The Standing Committee on Law and National Security voted to co-sponsor this Resolution on November 19, 2022.

4. Has this or a similar resolution been submitted to the House or Board previously?
No.

5. What existing Association policies are relevant to this resolution and how would they be affected by its adoption?

604

The ABA House of Delegates has passed resolutions that address issues with AI. This Resolution builds on and is consistent with those ABA policies.

- ABA urges courts and lawyers to address the emerging ethical and legal issues related to the usage of artificial intelligence (“AI”) in the practice of law, including (1) bias, explainability, and transparency of automated decisions made by AI; (2) ethical and beneficial usage of AI; and (3) controls and oversight of AI and the vendors that provide AI. 19A112.
- ABA urges federal, state, local, territorial and tribal governments to:
 - Ensure due process and refrain from using pretrial risk assessment tools unless the data supporting the risk assessment is transparent, publicly disclosed, and validated to demonstrate the absence of conscious or unconscious racial, ethnic, or other demographic, geographic, or socioeconomic bias; and
 - Recognize that an individual’s criminal history and other criteria may reflect structurally biased application of laws, policies or practices, as well as conscious or unconscious bias. 22M700.

6. If this is a late report, what urgency exists which requires action at this meeting of the House?

This is not a late report. As private sector organizations and governments move rapidly to design, develop, deploy, and use AI systems and capabilities, now is a critical time for lawyers to articulate principles that are essential to ensuring that AI is developed and implemented in accordance with the law and well-accepted legal standards.

7. Status of Legislation. (If applicable)

S. 1605, FY 2022 National Defense Authorization Act – enacted

Legislation to strengthen the U.S. government’s artificial intelligence (AI) readiness, support long-term investments in AI ethics and safety research, and increase governmental AI transparency, were passed as part of the FY 2022 *National Defense Authorization Act (NDAA)*.

Artificial Intelligence Capabilities and Transparency (AICT) Act.

The AICT Act would implement recommendations of the National Security Commission on Artificial Intelligence’s (NSCAI) final report. Congress established the NSCAI through the FY 2019 *National Defense Authorization Act (NDAA)* in order to consider the methods and means necessary to advance the development and improve the government’s use of AI and related technology.

S. 2551 — Artificial Intelligence Training for the Acquisition Workforce Act or the AI Training Act

This bill requires the Office of Management and Budget (OMB) to establish or otherwise provide an AI training program for the acquisition workforce of executive agencies (e.g., those responsible for program management or logistics) to ensure that the workforce has knowledge of the capabilities and risks associated with AI.

U.S. States

General AI bills or resolutions were introduced in at least 17 states in 2021-22, and were enacted in Alabama, Colorado, Illinois, Mississippi, Vermont, and Washington.

National Conference of State Legislatures (NCLS), *State AI Legislation*, <https://www.ncsl.org/research/telecommunications-and-information-technology/2020-legislation-related-to-artificial-intelligence.aspx>.

8. Brief explanation regarding plans for implementation of the policy, if adopted by the House of Delegates.

This Resolution will be disseminated to members of Congress and State legislators in coordination and cooperation with the ABA Governmental Affairs Office, as well as executives of large and small companies that design, develop, deploy, and use AI systems, capabilities, products, and services.

It will alert them to the ABA's newly-adopted policy and encourage them to take action consistent with the ABA policy. We also encourage its use in Amicus Curiae briefs by the ABA.

9. Cost to the Association. (Both direct and indirect costs).
None.

10. Disclosure of Interest. (If applicable)
Not Applicable.

11. Referrals.

Sections:

Business Law
Civil Rights & Social Justice
Criminal Justice
Environment, Energy & Resources
Intellectual Property
International Law
Litigation
Public Contract Law
Science & Technology Law
State and Local Government Law
Tort, Trial & Insurance Practice

604

Standing Committees:

Cybersecurity Legal Task Force
Professional Responsibility

Divisions:

Young Lawyers
Senior Lawyers
Law Practice

12. Contact Name and Address Information. (Prior to the meeting)

Lucy L. Thomson, Delegate, District of Columbia Bar
Livingston PLLC, Washington, D.C.
lucythomson1@mindspring.com, (703) 798-1001

Roland Trope
Trope Law, New York, New York
rltrope@tropelaw.com, (917) 370-3705

13. Contact Name and Address Information. (Who will present the report to the House?)

Lucy L. Thomson, Delegate, District of Columbia Bar
Livingston PLLC, Washington, D.C.
lucythomson1@mindspring.com, (703) 798-1001

EXECUTIVE SUMMARY

1. Summary of the Resolution

This Resolution presents guidance on how the legal system and its participants, including attorneys, regulators, and stakeholders, such as developers, integrators, suppliers, and operators (“developers”) of AI systems and capabilities, should assess fundamental issues with AI by addressing the principles of accountability, transparency and traceability.

2. Summary of the Issues that the Resolution Addresses

This Resolution states that in the context of AI individual and enterprise accountability and human authority, oversight, and control is required and it is not appropriate to shift legal responsibility to a computer or an “algorithm” rather than to responsible people and other legal entities.

By focusing in the context of AI on the key issues accountability, transparency and traceability, passage of this Resolution will help mitigate the risks that can result through implementation of AI systems and capabilities and enhance the use of AI in a trustworthy and responsible manner.

3. Please Explain How the Proposed Policy Position Will Address the Issue

This Resolution presents guidance on how the legal system and its participants, including attorneys, regulators, and stakeholders, including developers, integrators, suppliers, and operators (“developers”) of AI systems and capabilities, should assess fundamental issues with AI by addressing the principles of accountability, transparency and traceability. It states that in the context of AI individual and enterprise accountability and human authority, oversight, and control is required and it is not appropriate to shift legal responsibility to a computer or an “algorithm” rather than to responsible people and other legal entities.

Further, this Resolution would ensure that courts and participants in the legal process will have the capacity to evaluate and resolve legal questions and disputes by specifying the essential information that must be included in the development, deployment and use of AI to ensure transparency and traceability.

4. Summary of Minority Views

None.



The State Bar of California

OPEN SESSION AGENDA ITEM 60-1 NOVEMBER 2023

DATE: November 16, 2023

TO: Members, Board of Trustees
Sitting as the Regulation and Discipline Committee

FROM: The Committee on Professional Responsibility and Conduct
Brandon Krueger, Chair, Committee on Professional Responsibility and Conduct
Erika Doherty, Program Director, Office of Professional Competence

SUBJECT: Recommendations from Committee on Professional Responsibility and
Conduct on Regulation of Use of Generative AI by Licensees

EXECUTIVE SUMMARY

This memorandum sets forth the Committee on Professional Responsibility and Conduct's (COPRAC) initial recommendations regarding lawyer use of generative AI. In short, COPRAC believes that the existing Rules of Professional Conduct are robust, and the standards of conduct cover the landscape of issues presented by generative AI in its current forms. However, COPRAC recognizes that generative AI is a rapidly evolving technology that presents novel issues that might necessitate new regulation and rules in the future.

As an initial step, COPRAC has developed, and recommends that the Board adopt *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* to assist lawyers in navigating their ethical obligations when using generative AI. COPRAC envisions that the Practical Guidance will be a living document that is periodically updated as the technology evolves and matures, and new issues are presented.

COPRAC also recommends that the Board direct State Bar staff to develop attorney education programs that assist lawyers to understand and gain competence regarding the potential risks, benefits and ethical implications of using generative AI; examine the potential impacts of generative AI on law students and bar applicants; and work with the Legislature and California

Supreme Court to consider new or revised regulations regarding the use of generative AI in the practice of law.

BACKGROUND

On May 18, 2023, the chair of the Board of Trustees directed COPRAC, which is charged with studying and providing consultation and assistance to the Board on matters involving professional responsibility, to explore potential regulation of the ethical use of generative AI in the legal profession. The chair directed that, by the Board's November 2023 meeting, COPRAC issue recommendations, which could include practical guidance, an advisory opinion or other resources, changes to the Rules of Professional Conduct or other rules or statutes, or other recommendations to ensure that AI is used competently and in compliance with the professional responsibility obligations of lawyers.

COPRAC undertook an effort to familiarize committee members with the current state of generative AI and to understand its potential implications for the legal profession prior to developing recommendations regarding lawyer use of this evolving technology. COPRAC accomplished this work by forming a working team on generative AI (that included experts in the field on an ad hoc basis) and discussions and considerations at four COPRAC meetings on June 23, July 28, September 15, and October 20, 2023. This work also included:

- Surveying lawyers regarding current and planned uses of generative AI in their practices;
- Researching generative AI capabilities, limitations, and risks, by reviewing various materials, including the principles and guidelines prepared by [MIT's Task Force on Responsible Use of Generative AI for Law](#), and consulting with experts in artificial intelligence and founders of generative AI products;
- Reviewing the current Rules of Professional Conduct, statutory authority, case law, and ethics opinions to evaluate whether these existing authorities address the use of generative AI and to identify potential new ethical issues raised by generative AI; and
- Examining approaches taken by other jurisdictions to regulate the use of generative AI, specifically any regulations directed toward lawyers.

DISCUSSION

The current Rules of Professional Conduct do not expressly address the use of generative AI, creating uncertainty about lawyers' ethical duties regarding such use.¹ However, the rules are intended to apply to lawyers engaged in a variety of practice areas and situations.

¹ Comment [1] to Rule 1.1 (Competence) is the only explicit reference to technology. The comment, adopted March 22, 2021, states, "[t]he duties set forth in this rule include the duty to keep abreast of the changes in the law and its practice, including the benefits and risks associated with relevant technology."

Historically, COPRAC has developed advisory ethics opinions that apply the rules and related authorities to certain situations. These opinions are issued for public input through a public comment process and ultimately approved by the Board of Trustees acting as the Regulation and Discipline Committee. After engaging in extensive study over the past several months, COPRAC believes that the existing rules can be applied to generative AI use at this time, and has prepared *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* (Practical Guidance), provided as Attachment A. This document is an interim step to provide guidance on this evolving technology while further rules and regulations are considered. The Practical Guidance sets forth the applicable Rules of Professional Conduct and statutory authority that would regulate the improper use of generative AI, and offers guidance for how a lawyer may comply with these ethics authorities.

The Practical Guidance is based, in part, on the principles and guidelines prepared by [MIT's Task Force on Responsible Use of Generative AI for Law](#), and addresses current concerns about lawyer use of generative AI, many of which apply in varying degrees to lawyer use of other technologies.

COPRAC recognizes that as the technology further develops, additional regulation, including amendments to the Rules of Professional Conduct, may be necessary. However, until there are issues presented by the use of generative AI that are not adequately addressed by existing rules and regulations, this Practical Guidance will remind lawyers of their existing professional responsibility obligations and assist lawyers with applying these obligations to new technology. In addition to recommending that the Board adopt the Practical Guidance, COPRAC intends to further study the following and, if necessary, return with further recommendations to the Board regarding:

- how to balance rules and guidance in the use of generative AI to protect clients and the public against its potential to facilitate efficiency and expanded access to justice;
- how to “supervise” non-human, nonlawyer assistance if the assistance allows for autonomous decision making by generative AI;
- whether the duty of competency should specifically require competency in generative AI (i.e., requirement more than what exists in Rule 1.1, Comment [1]); and
- whether a lawyer should be required to communicate to their client the use of generative AI and in what contexts.

The impact of generative AI on the profession extends well beyond a lawyer’s professional responsibility obligations. In addition to publishing and maintaining the Practical Guidance, COPRAC recommends that the Board take other action regarding generative AI:

Develop Attorney Education Addressing Generative AI

COPRAC recommends that the Board direct the Office of Professional Competence (OPC) to develop a one-hour minimum continuing legal education (MCLE) course that would satisfy the new, one-hour requirement for continuing legal education on technology in the practice of law and that addresses the competent use of generative AI (State Bar rule 2.72(C)(2)(a)(iv)).

COPRAC further recommends that the Board direct OPC to update the mandatory New Attorney Training, which new licensees must complete within their first year of practice, to include technological competence training for lawyers using generative AI. COPRAC believes that education in this area will allow lawyers to utilize generative AI for the benefit of their clients and to expand access to legal services while upholding professional ethics without harm to the public while the technology continues to develop.

Explore Regulatory Changes to Protect the Public

Generative AI products are being developed for a multitude of uses and for a variety of professions. They are also being developed to provide legal assistance to unrepresented persons. While generative AI may be of great benefit in minimizing the justice gap, it could also create harm if self-represented individuals are relying on generative AI outputs that provide false information. COPRAC recommends that the Board take action to:

- Work with the Legislature and the California Supreme Court to determine whether the unauthorized practice of law should be more clearly defined or articulated through statutory or rule changes; and
- Work with the Legislature to determine whether legal generative AI products should be licensed or regulated and, if so, how.

Consider the Impact of Generative AI on Law Students and Bar Applicants

Additionally, COPRAC recommends that the Board consider taking action to address generative AI use by law students by:

- Directing the Committee of Bar Examiners to explore requirements for California-accredited law schools to require courses regarding the competent use of generative AI; and
- Directing the Committee of Bar Examiners to explore regulations or rules related to the bar exam and generative AI.

COPRAC recognizes that the Practical Guidance document and other recommendations are a first step in the regulation of generative AI use by California lawyers, and that the State Bar is one of the first attorney regulatory agencies to address this technology. Through these initial recommendations, COPRAC believes that the State Bar will allow for attorneys and consumers

to gain the benefits of this transformative technology, while promoting responsible use of generative AI in a manner that will prevent public harm.

FISCAL/PERSONNEL IMPACT

None

AMENDMENTS TO RULES

None

AMENDMENTS TO BOARD OF TRUSTEES POLICY MANUAL

None

STRATEGIC PLAN GOALS & IMPLEMENTATION STEPS

Goal 3. Protect the Public by Regulating the Legal Profession

RECOMMENDATIONS

Should the Board of Trustees, sitting as the Regulation and Discipline Committee, concur in COPRAC's proposed Practical Guidance and further recommendations, passage of the following resolutions is recommended:

RESOLVED, that the Board of Trustees sitting as the Regulation and Discipline Committee, upon recommendation of the State Bar Committee on Professional Responsibility and Conduct, approves the publication of the *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law*, provided as Attachment A; and it is

FURTHER RESOLVED, that the Board of Trustees sitting as the Regulation and Discipline Committee, upon recommendation of the State Bar Committee on Professional Responsibility and Conduct, directs the State Bar Office of Professional Competence to (1) develop a one-hour minimum continuing legal education (MCLE) course that would satisfy the new, one-hour requirement for continuing legal education on technology in the practice of law and that addresses the competent use of generative AI; and (2) update the New Attorney Training to include technological competence training for lawyers using generative AI; and it is

FURTHER RESOLVED, that the Board of Trustees sitting as the Regulation and Discipline Committee, upon recommendation of the State Bar Committee on Professional Responsibility and Conduct, directs State Bar staff to

work with the Legislature and the California Supreme Court to determine whether (1) the unauthorized practice of law should be more clearly defined or articulated through statutory or rule changes; and (2) legal generative AI products should be licensed or regulated and, if so, how; and it is

FURTHER RESOLVED, that the Board of Trustees sitting as the Regulation and Discipline Committee, upon recommendation of the State Bar Committee on Professional Responsibility and Conduct, directs the State Bar Office of Admissions and the Committee of Bar Examiners to explore (1) requirements for California-accredited law schools to require courses regarding the competent use of generative AI; and (2) regulations or rules related to the bar exam and generative AI.

ATTACHMENT LIST

- A.** Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law

**THE STATE BAR OF CALIFORNIA
STANDING COMMITTEE ON
PROFESSIONAL RESPONSIBILITY AND CONDUCT
PRACTICAL GUIDANCE FOR THE USE OF GENERATIVE ARTIFICIAL INTELLIGENCE IN THE
PRACTICE OF LAW**

EXECUTIVE SUMMARY

Generative AI is a tool that has wide-ranging application for the practice of law and administrative functions of the legal practice for all licensees, regardless of firm size, and all practice areas. Like any technology, generative AI must be used in a manner that conforms to a lawyer’s professional responsibility obligations, including those set forth in the Rules of Professional Conduct and the State Bar Act. A lawyer should understand the risks and benefits of the technology used in connection with providing legal services. How these obligations apply will depend on a host of factors, including the client, the matter, the practice area, the firm size, and the tools themselves, ranging from free and readily available to custom-built, proprietary formats.

Generative AI use presents unique challenges; it uses large volumes of data, there are many competing AI models and products, and, even for those who create generative AI products, there is a lack of clarity as to how it works. In addition, generative AI poses the risk of encouraging greater reliance and trust on its outputs because of its purpose to generate responses and its ability to do so in a manner that projects confidence and effectively emulates human responses. A lawyer should consider these and other risks before using generative AI in providing legal services.

The following Practical Guidance is based on current professional responsibility obligations for lawyers and demonstrates how to behave consistently with such obligations. While this guidance is intended to address issues and concerns with the use of generative AI and products that use generative AI as a component of a larger product, it may apply to other technologies, including more established applications of AI. This Practical Guidance should be read as guiding principles rather than as “best practices.”

PRACTICAL GUIDANCE

Applicable Authorities	Practical Guidance
<p>Duty of Confidentiality</p> <p>Bus. & Prof. Code, § 6068, subd. (e)</p> <p>Rule 1.6</p> <p>Rule 1.8.2</p>	<p>Generative AI products are able to utilize the information that is input, including prompts and uploaded documents or resources, to train the AI, and might also share the query with third parties or use it for other purposes. Even if the product does not utilize or share inputted information, it may lack reasonable or adequate security.</p> <p>A lawyer must not input any confidential information of the client into any generative AI solution that lacks adequate confidentiality and security protections. A lawyer must anonymize client information and avoid entering details that can be used to identify the client.</p> <p>A lawyer or law firm should consult with IT professionals or cybersecurity experts to ensure that any AI system in which a lawyer would input confidential client information adheres to stringent security, confidentiality, and data retention protocols.</p> <p>A lawyer should review the Terms of Use or other information to determine how the product utilizes inputs. A lawyer who intends to use confidential information in a generative AI product should ensure that the provider does not share inputted information with third parties or utilize the information for its own use in any manner, including to train or improve its product.</p>
<p>Duties of Competence and Diligence</p> <p>Rule 1.1</p> <p>Rule 1.3</p>	<p>It is possible that generative AI outputs could include information that is false, inaccurate, or biased.</p> <p>A lawyer must ensure competent use of the technology, including the associated benefits and risks, and apply diligence and prudence with respect to facts and law.</p>

Applicable Authorities	Practical Guidance
	<p>Before using generative AI, a lawyer should understand to a reasonable degree how the technology works, its limitations, and the applicable terms of use and other policies governing the use and exploitation of client data by the product.</p> <p>Overreliance on AI tools is inconsistent with the active practice of law and application of trained judgment by the lawyer.</p> <p>AI-generated outputs can be used as a starting point but must be carefully scrutinized. They should be critically analyzed for accuracy and bias, supplemented, and improved, if necessary. A lawyer must critically review, validate, and correct both the input and the output of generative AI to ensure the content accurately reflects and supports the interests and priorities of the client in the matter at hand, including as part of advocacy for the client. The duty of competence requires more than the mere detection and elimination of false AI-generated results.</p> <p>A lawyer’s professional judgment cannot be delegated to generative AI and remains the lawyer’s responsibility at all times. A lawyer should take steps to avoid over-reliance on generative AI to such a degree that it hinders critical attorney analysis fostered by traditional research and writing. For example, a lawyer may supplement any AI-generated research with human-performed research and supplement any AI-generated argument with critical, human-performed analysis and review of authorities.</p>
<p>Duty to Comply with the Law</p> <p>Bus. & Prof. Code, § 6068(a)</p> <p>Rule 8.4</p> <p>Rule 1.2.1</p>	<p>A lawyer must comply with the law and cannot counsel a client to engage, or assist a client in conduct that the lawyer knows is a violation of any law, rule, or ruling of a tribunal when using generative AI tools.</p> <p>There are many relevant and applicable legal issues surrounding generative AI, including but not limited to compliance with AI-specific laws, privacy laws, cross-border data transfer laws, intellectual property laws, and cybersecurity concerns. A lawyer should analyze the relevant laws and regulations applicable to the attorney or the client.</p>

Applicable Authorities	Practical Guidance
<p>Duty to Supervise Lawyers and Nonlawyers, Responsibilities of Subordinate Lawyers</p> <p>Rule 5.1 Rule 5.2 Rule 5.3</p>	<p>Managerial and supervisory lawyers should establish clear policies regarding the permissible uses of generative AI and make reasonable efforts to ensure that the firm adopts measures that give reasonable assurance that the firm’s lawyers and non lawyers’ conduct complies with their professional obligations when using generative AI. This includes providing training on the ethical and practical aspects, and pitfalls, of any generative AI use.</p> <p>A subordinate lawyer must not use generative AI at the direction of a supervisory lawyer in a manner that violates the subordinate lawyer’s professional responsibility and obligations.</p>
<p>Communication Regarding Generative AI Use</p> <p>Rule 1.4 Rule 1.2</p>	<p>A lawyer should evaluate their communication obligations throughout the representation based on the facts and circumstances, including the novelty of the technology, risks associated with generative AI use, scope of the representation, and sophistication of the client.</p> <p>The lawyer should consider disclosure to their client that they intend to use generative AI in the representation, including how the technology will be used, and the benefits and risks of such use.</p> <p>A lawyer should review any applicable client instructions or guidelines that may restrict or limit the use of generative AI.</p>
<p>Charging for Work Produced by Generative AI and Generative AI Costs</p> <p>Rule 1.5 Bus. & Prof. Code, §§ 6147–6148</p>	<p>A lawyer may use generative AI to more efficiently create work product and may charge for actual time spent (e.g., crafting or refining generative AI inputs and prompts, or reviewing and editing generative AI outputs). A lawyer must not charge hourly fees for the time saved by using generative AI.</p> <p>Costs associated with generative AI may be charged to the clients in compliance with applicable law.</p> <p>A fee agreement should explain the basis for all fees and costs, including those associated with the use of generative AI.</p>

Applicable Authorities	Practical Guidance
<p>Candor to the Tribunal; and Meritorious Claims and Contentions</p> <p>Rule 3.1</p> <p>Rule 3.3</p>	<p>A lawyer must review all generative AI outputs, including, but not limited to, analysis and citations to authority for accuracy before submission to the court, and correct any errors or misleading statements made to the court.</p> <p>A lawyer should also check for any rules, orders, or other requirements in the relevant jurisdiction that may necessitate the disclosure of the use of generative AI.</p>
<p>Prohibition on Discrimination, Harassment, and Retaliation</p> <p>Rule 8.4.1</p>	<p>Some generative AI is trained on biased information, and a lawyer should be aware of possible biases and the risks they may create when using generative AI (e.g., to screen potential clients or employees).</p> <p>Lawyers should engage in continuous learning about AI biases and their implications in legal practice, and firms should establish policies and mechanisms to identify, report, and address potential AI biases.</p>
<p>Professional Responsibilities Owed to Other Jurisdictions</p> <p>Rule 8.5</p>	<p>A lawyer should analyze the relevant laws and regulations of each jurisdiction in which a lawyer is licensed to ensure compliance with such rules.</p>

AMERICAN BAR ASSOCIATION

HOUSE OF DELEGATES

ADOPTED AUGUST 12-13, 2019

RESOLUTION

RESOLVED, That the American Bar Association urges courts and lawyers to address the emerging ethical and legal issues related to the usage of artificial intelligence (“AI”) in the practice of law including: (1) bias, explainability, and transparency of automated decisions made by AI; (2) ethical and beneficial usage of AI; and (3) controls and oversight of AI and the vendors that provide AI.

REPORT¹

I. PURPOSE OF THIS RESOLUTION AND REPORT

Lawyers increasingly are using artificial intelligence (“AI”) in their practices to improve the efficiency and accuracy of legal services offered to their clients. But while AI offers cutting-edge advantages and benefits, it also raises complicated questions implicating professional ethics.

The purpose of this resolution and report is to urge courts and lawyers to address the emerging legal and ethical issues related to the usage of AI in the practice of law.

Courts and lawyers must be aware of the issues involved in using (and not using) AI, and they should address situations where their usage of AI may be flawed or biased.

In order to assist courts and lawyers in addressing these AI issues, we will be exploring the establishment of a working group to, in part, define guidelines for legal and ethical AI usage, and potentially develop a model standard that could come to the American Bar Association House of Delegates for adoption. We acknowledge that there are many AI principles being developed by organizations and governments, including the OECD Principles on Artificial Intelligence², the Universal Guidelines for AI³, the IEEE’s Ethically Aligned Design⁴, and California’s ACR-215 23 Asilomar AI Principles (2017-2018)⁵. As part of the working group, we intend to study such principles to recommend an ABA specific AI principle. While this report focuses on AI usage by courts and lawyers in the practice of law, the concerns set forth in this report - AI bias, explainability, transparency, ethical and beneficial uses of AI, monitoring, accountability, controls and oversight, can apply broadly. In the future, it might be appropriate for the ABA and the proposed working group to focus on the broader ethical usage of AI by courts, lawyers, federal, state, local, territorial and tribal governments and the private sector beyond the practice of law.

Section II of this report provides an overview of AI and the different AI tools used in the practice of law. Section III, in turn, analyzes a lawyer’s ethical duties in connection with AI technology. Section IV explores how bias can affect AI and the importance of using diverse teams when developing AI technology. Section V discusses questions to ask when adopting an AI solution or engaging an AI vendor. And finally, the report concludes with Section VI.

II. OVERVIEW OF HOW ARTIFICIAL INTELLIGENCE IS CHANGING THE LAW

Artificial intelligence promises to change not only the practice of law but our economy as a whole. We clearly are on the cusp of an AI revolution. But what does all this mean, as

¹ This report is based on the article “Legal Ethics in the Use of Artificial Intelligence” by Janine Cerny, Steve Delchin, and Huu Nguyen, https://download.pli.edu/WebContent/pm/249218/pdf/02-22-19_1600_115843_LegalEthics.pdf with full permission

² <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>

³ <https://thepublicvoice.org/ai-universal-guidelines/>

⁴ <https://ethicsinaction.ieee.org/>

⁵ https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201720180ACR215

a practical matter, for lawyers? What is AI? And how is it being used in the practice of law?

A. Defining AI.

Artificial intelligence has been defined as “the capability of a machine to imitate intelligent human behavior.”⁶ Others have defined it as “cognitive computing” or “machine learning.”⁷ Although there are many descriptive terms used, AI at its core encompasses tools that are trained rather than programmed. It involves teaching computers how to perform tasks that typically require human intelligence such as perception, pattern recognition, and decision-making.⁸

B. How AI Is Being Used In The Practice Of Law

There are many different ways that lawyers today are using AI to improve productivity and provide better legal services to their clients. Below are several of the main examples. As AI becomes even more advanced in the coming years, it fundamentally will transform the practice of law. Lawyers who do not adopt AI will be left behind.

1. Electronic Discovery/Predictive Coding.

Lawyers, predictably, use AI for electronic discovery. The process involves an attorney training the computer how to categorize documents in a case. Through a method of predictive coding, AI technology is able to classify documents as relevant or irrelevant, among other classifications, after extrapolating data gathered from a sample of documents classified by the attorney.⁹

2. Litigation Analysis/Predictive Analysis.

AI also is being used to predict the outcome of litigation through the method of predictive analytics. AI tools utilize case law, public records, dockets, and jury verdicts to identify patterns in past and current data.¹⁰ AI then analyzes the facts of a lawyer’s case to provide an intelligent prediction of the outcome.¹¹

⁶ *Artificial Intelligence*, MERRIAM-WEBSTER (April 6, 2017), available at <https://www.merriam-webster.com/dictionary/artificial%20intelligence>.

⁷ Lisa Morgan, *4 Types of Machine Intelligence You Should Know*, Information Week (Apr. 10, 2018) <https://www.informationweek.com/big-data/ai-machine-learning/4-types-of-machine-intelligence-you-should-know/a/d-id/1331480>.

⁸ Sterling Miller, *Artificial Intelligence – What Every Legal Department Really Needs To Know*, Ten Things You Need to Know as In-House Counsel (Aug. 15, 2017), <https://hilgersgraben.com/blogs/blogs-hidden.html/article/2017/08/15/ten-things-artificial-intelligence-what-every-legal-department-really-needs-to-know>

⁹ David L. Gordon & Rebecca L. Ambrose, *The Ethics of Artificial Intelligence*, The Jackson Lewis Corporate Counsel Conference (2017), https://www.jacksonlewis.com/sites/default/files/docs/Final_The%20Ethics%20of%20Artificial%20Intelligence_Gordon%20and%20Ambrose.pdf.

¹⁰ *Supra*, note 3

¹¹ *Id.*

3. Contract Management.

AI tools are being used by lawyers to assist with contract management. This is particularly valuable to inside counsel who quickly need to identify important information in contracts. For example, AI tools can flag termination dates and alert the lawyer about deadlines for sending a notice of renewal. AI tools also can identify important provisions in contracts, such as most favored nation clauses, indemnification obligations, and choice of law provisions, among others.¹²

4. Due Diligence Reviews.

AI is being used to assist in automated due diligence review for corporate transactions to reduce the burden of reviewing large numbers of documents.¹³ Similar to contract management, due diligence review involves the computer identifying and summarizing key clauses from contracts.¹⁴

5. “Wrong Doing” Detection.

AI is being used to search company records to detect bad behavior preemptively. AI is able to see beyond attempts to disguise wrongdoing and identify code words.¹⁵ AI can also review employee emails to determine morale, which may lead to identification of wrongdoing.¹⁶ For example, in one test using emails of Enron executives, AI was able to detect tension amongst employees that was correlated with a questionable business deal.¹⁷

6. Legal Research.

AI traditionally has been used to assist with legal research, but it increasingly is becoming more sophisticated. With AI, lawyers can rely on natural language queries—rather than simple Boolean queries—to return more meaningful and more insightful results.¹⁸ AI also can be used to produce basic legal memos. One AI program called Ross Intelligence, which uses IBM’s Watson AI technology, can produce a brief legal memo in response to

¹² *Id.*

¹³ *Id.*

¹⁴ Lauri Donahue, *A Primer on Using Artificial Intelligence in the Legal Profession*, Harvard Journal of Law and Technology (Jan. 3, 2018) <https://jolt.law.harvard.edu/digest/a-primer-on-using-artificial-intelligence-in-the-legal-profession>.

¹⁵ Sterling Miller, *Artificial Intelligence and its Impact on Legal Technology: To Boldly Go Where No Legal Department Has Gone Before*, Thomson Reuters, <https://legal.thomsonreuters.com/en/insights/articles/AI-and-its-impact-on-legal-technology>.

¹⁶ Frank Partnoy, *What Your Boss Could Learn by Reading the Whole Company’s Emails*, The Atlantic (Sep. 2018) <https://www.theatlantic.com/magazine/archive/2018/09/the-secrets-in-your-inbox/565745/>.

¹⁷ *Id.*

¹⁸ *Supra*, note 3.

a lawyer's legal question.¹⁹ Over time, such AI technology will become more and more powerful.

7. AI to Detect Deception.²⁰

Finally, as AI becomes more advanced, it will be used by lawyers to detect deception. Researchers, for example, are working on developing AI that can detect deception in the courtroom. In one test run, an AI system performed with 92 percent accuracy, which the researchers described as “significantly better” than humans.²¹ While AI is still being developed for use in courtrooms, it already is being deployed outside the practice of law. For example, the United States, Canada, and European Union have run pilot programs using deception-detecting kiosks for border security.²²

C. It is Essential for Lawyers to be Aware of AI.

The bottom line is that it is essential for lawyers to be aware of how AI can be used in their practices to the extent they have not done so yet. AI allows lawyers to provide better, faster, and more efficient legal services to companies and organizations. The end result is that lawyers using AI are better counselors for their clients. In the next few years, the use of AI by lawyers will be no different than the use of email by lawyers—an indispensable part of the practice law.²³

Not surprisingly, given its benefits, more and more business leaders are embracing AI, and they naturally will expect both their in-house lawyers and outside counsel to embrace it as well. Lawyers who already are experienced users of AI technology will have an advantage and will be viewed as more valuable to their organizations and clients. From a professional development standpoint, lawyers need to stay ahead of the curve when it comes to AI. But even apart from the business dynamics, professional ethics requires lawyers to be aware of AI and how it can be used to deliver client services. As explored next, a number of ethical rules apply to lawyers' use and non-use of AI.

III. THE LEGAL ETHICS OF AI.

Given the transformative nature of AI, it is important for courts and lawyers to understand how existing and well established ethical rules may apply to the use of AI.

A. Several Ethics Rules Apply To Lawyer's Use (And Non-Use) of AI.

¹⁹ Steve Lohr, *A.I. Is Doing Legal Work. But It Won't Replace Lawyers, Yet*, New York Times (Mar. 9, 2017) <https://www.nytimes.com/2017/03/19/technology/lawyers-artificial-intelligence.html>.

²⁰ Shivali Best, *The Robot That Knows When You're Lying*, DailyMail (Dec. 20, 2017), <http://www.dailymail.co.uk/sciencetech/article-5197747/AI-detects-expressions-tell-people-lie-court.html>.

²¹ *Id.*

²² Jeff Daniels, *Lie-detecting Computer Kiosks Equipped with Artificial Intelligence Look Like the Future of Border Security*, CNBC (May 15, 2018) <https://www.cnbc.com/2018/05/15/lie-detectors-with-artificial-intelligence-are-future-of-border-security.html>.

²³ *Supra*, note 3.

There are a number of ethical duties that apply to the use of (and non-use of) AI by lawyers, including the duties of: (1) competence (and diligence), (2) communication, (3) confidentiality, and (4) supervision. These duties as applied to AI technology are discussed below.

1. Duty of Competence

Under Rule 1.1 of the ABA Model Rules, a lawyer must provide competent representation to his or her client. The rule states that “[c]ompetent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.”²⁴ The duty of competence requires lawyers to be informed, and up to date, on current technology. In 2012, this was made clear when the ABA adopted Comment 8 to Rule 1.1 which states that “[t]o maintain the requisite knowledge and skill, lawyers should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology”²⁵

As one author points out, there does not appear to be any instance “in which AI represents the standard of care in an area of legal practice, such that its use is necessary.”²⁶ Nonetheless, lawyers generally must understand the technology available to improve the legal services they provide to clients. Lawyers have a duty to identify the technology that is needed to effectively represent the client, as well as determine if the use of such technology will improve service to the client.²⁷

Under Rule 1.1, lawyers also must have a basic understanding of how AI tools operate. While lawyers cannot be expected to know all the technical intricacies of AI systems, they are required to understand how AI technology produces results. As one legal commentator notes, “[i]f a lawyer uses a tool that suggests answers to legal questions, he must understand the capabilities and limitations of the tool, and the risks and benefits of those answers.”²⁸

2. Duty to Communicate

ABA Model Rule 1.4 governs a lawyer’s duty to communicate with clients and requires a lawyer to “reasonably consult with the client about the means by which the client’s objectives are to be accomplished.”²⁹ A lawyer’s duty of communication under Rule 1.4 includes discussing with his or her client the decision to use AI in providing legal services.

²⁴ ABA Model Rule 1.1

²⁵ Hedda Litwin, *The Ethical Duty of Technology Competence: What Does it Mean for You?*, National Association of Attorneys General, <https://www.naag.org/publications/nagtri-journal/volume-2-issue-4/the-ethical-duty-of-technology-competence-what-does-it-mean-for-you.php>.

²⁶ James Q. Walker, *What’s Artificial About Intelligence? The Ethical and Practical Considerations When Lawyers Use AI Technology*, Bloomberg Law (2018), <https://www.rkollp.com/newsroom-publications-443.html>.

²⁷ *Supra*, note 4.

²⁸ David Lat, *The Ethical Implications of Artificial Intelligence*, Above the Law: Law2020, <https://abovethelaw.com/law2020/the-ethical-implications-of-artificial-intelligence/>.

²⁹ ABA Model Rule 1.4.

A lawyer should obtain approval from the client before using AI, and this consent must be informed. The discussion should include the risks and limitations of the AI tool.³⁰ In certain circumstances, a lawyer's decision *not* to use AI also may need to be communicated to the client if using AI would benefit the client.³¹ Indeed, the lawyer's failure to use AI could implicate ABA Model Rule 1.5, which requires lawyer's fees to be reasonable. Failing to use AI technology that materially reduces the costs of providing legal services arguably could result in a lawyer charging an unreasonable fee to a client.³²

3. Duty of Confidentiality

Under ABA Model Rule 1.6, lawyers owe their clients a generally duty of confidentiality. This duty specifically requires a lawyer to "make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client."³³ The use of some AI tools may require client confidences to be "shared" with third-party vendors. As a result, lawyers must take appropriate steps to ensure that their clients' information appropriately is safeguarded.³⁴ Appropriate communication with the client also is necessary.

To minimize the risks of using AI, a lawyer should discuss with third-party AI providers the confidentiality safeguards in place. A lawyer should inquire about "what type of information is going to be provided, how the information will be stored, what security measures are in place with respect to the storage of the information, and who is going to have access to the information."³⁵ AI should not be used in the representation unless the lawyer is confident that the client's confidential information will be secure.

4. Duty to Supervise

Under ABA Model Rules 5.1 and 5.3, lawyers have an ethical obligation to supervise lawyers and nonlawyers who are assisting lawyers in the provision of legal services to ensure that their conduct complies with the Rules of Professional Conduct.³⁶ In 2012, the title of Model Rule 5.3 was changed from "Responsibilities Regarding Nonlawyer Assistants" to "Responsibilities Regarding Nonlawyer Assistance."³⁷ The change clarified that the scope of Rule 5.3 encompasses nonlawyers whether human or not. Under Rules 5.1 and 5.3, lawyers are obligated to supervise the work of AI utilized in the provision of legal services, and understand the technology well enough to ensure compliance with the

³⁰ *Supra*, note 4.

³¹ *Id.*

³² *Ethical Use of Artificial Intelligence in the Legal Industry: The Rules of Professional Conduct*, Emerging Industries and Technology Committee Newsletter, (March 2018), <https://insolvencyintel.abi.org/bankruptcyarticles/ethical-use-of-artificial-intelligence-in-the-legal-industry-the-rules-of-professional-conduct>.

³³ ABA Model Rule 1.6.

³⁴ *Supra*, note 4.

³⁵ *Id.*

³⁶ *Id.*

³⁷ *Variations of the ABA Model Rules of Professional Conduct*, ABA CPR Policy Implementation Committee (Sep. 29, 2017)

https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/mrpc_5_3.pdf.

lawyer's ethical duties. This includes making sure that the work product produced by AI is accurate and complete and does not create a risk of disclosing client confidential information.³⁸

There are some tasks that should not be handled by today's AI technology, and a lawyer must know where to draw the line. At the same time, lawyers should avoid underutilizing AI, which could cause them to serve their clients less efficiently.³⁹ Ultimately, it's a balancing act. Given that many lawyers are focused on detail and control over their matter, it is easy to see why "the greater danger might very well be underutilization of, rather than overreliance upon, artificial intelligence."⁴⁰

B. Key Practical Takeaways Relating to The Ethics of AI.

There clearly are a number of ethical rules that apply to lawyers' use and non-use of AI technology, and they have real-world application. Lawyers must be informed about AI's ability to deliver efficient and accurate legal services to clients while keeping in mind the ethical requirements and limitations. Ultimately, lawyers must exercise independent judgment, communicate with clients, and supervise the work performed by AI. In many ways, the ethical issues raised by AI are simply a permutation of ethical issues that lawyers have faced before with regard to other technology. It shows that the legal ethics rules are adaptable to new technologies, and AI is no exception.

IV. BIAS AND TRANSPARENCY IN THE AI CONTEXT.

There is a final, often overlooked consideration in a lawyer's use of AI technology, and that is the problem of bias. For all the advantages that AI offers for lawyers, there also is a genuine concern that AI technology may reflect the biases and prejudices of its developers and trainers, which in turn may lead to skewed results. It is critical for lawyers using AI to understand and address how bias can impact AI results.

The problem of bias in the development and use of AI potentially implicates professional ethics. In August 2016, the ABA adopted Model Rule 8.4(g), which prohibits harassment and discrimination by lawyers against eleven protected classes.⁴¹ Rule 8.4(g) states that it is professional misconduct for a lawyer to "engage in conduct that the lawyer knows or reasonably should know is harassment or discrimination on the basis of race, sex, religion, national origin, ethnicity, disability, age, sexual orientation, gender identity, marital status or socioeconomic status in conduct related to the practice of law."⁴² About 20 states already have some variation of ABA Model Rule 8.4 on the books, and several other states are considering whether to adopt ABA's new expansive rule. Lawyers in jurisdictions that have adopted some form of Rule 8.4 must consider whether their use of

³⁸ *Supra*, note 4.

³⁹ *Supra*, note 22.

⁴⁰ *Id.*

⁴¹ *ABA Rule 8.4 Finding Few Followers, but Sparking Lots of Encouraging Discussion*, ABA (Aug. 3, 2018) https://www.americanbar.org/news/abanews/aba-news-archives/2018/08/aba_rule_8_4_finding/.

⁴² ABA Model Rule 8.4(g).

AI is consistent with the rule. Moreover, even in jurisdictions that have not adopted some form of Rule 8.4, lawyers must consider how bias in the use of AI could create risks for clients.

Bias in AI technology stems from the nature of AI tools, which involve machine training rather than programming. If the data used for training is biased, the AI tool will produce a biased result. For example, one major company recently launched an AI tool that could have text-based conversations with individuals.⁴³ The tool continuously learned how to respond in conversations based on previous conversations. Unfortunately, the tool began to mimic the discriminatory viewpoints of the people it previously engaged in conversation.⁴⁴

As yet another example, the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) software used by some courts to predict the likelihood of recidivism in criminal defendants has been shown by studies to be biased against African-Americans.⁴⁵ For these reasons, it is important to have diverse teams developing AI to ensure that biases are minimized. The data used for training AI should also be carefully reviewed in order to prevent bias.

In the AI world, there has been a movement away from “black box” AI, in which an AI model is not able to explain how it generated its output based on the input.⁴⁶ The preferred model is now “explainable AI,”⁴⁷ which is able to provide the reasoning for how decisions are reached. The importance of transparency in the use of AI is being recognized by governments. New York City, for example, recently passed a law that requires creation of a task force that monitors algorithms used by its government, such as those used to assign children to public schools.⁴⁸ One of the task force’s responsibilities is to determine how to share with the public the factors that go into the algorithms.⁴⁹

There are also industry specific laws that prohibit bias and require transparency which may cover AI decision making. Competent counsel should understand these laws and their AI context as well. For example, under the Fair Credit Reporting Act (“FCRA”), 15 U.S.C. § 1681 et seq., among other requirements, any financial institution that uses a

⁴³ Jonathon Vanian, *Unmasking A.I.’s Bias Problem*, Fortune (June 25, 2018)

<https://www.fortune.com/longform/ai-bias-problem/>.

⁴⁴ *Id.*

⁴⁵ Julia Angwin, et. al., *Machine Bias*, ProPublica (May 23, 2016)

<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

⁴⁶ Jason Bloomberg, *Don’t Trust Artificial Intelligence? Time to Open the AI ‘Black Box’*, Forbes (Sep. 16, 2018) <https://www.forbes.com/sites/jasonbloomberg/2018/09/16/dont-trust-artificial-intelligence-time-to-open-the-ai-black-box/#56c1d9a3b4a7>.

⁴⁷ *Opening AI’s Black Box Will Become a Priority*, PwC,

<https://www.pwc.com/us/en/services/consulting/library/artificial-intelligence-predictions/explainable-ai.html>.

⁴⁸ Elizabeth Zima, *Could New York City’s AI Transparency Bill Be a Model for the Country?*, Government Technology (Jan. 4, 2018), <https://www.govtech.com/policy/Could-New-York-Citys-AI-Transparency-Bill-Be-A-Model-for-the-Country.html>.

⁴⁹ *Id.*

credit report or another type of consumer report to deny a consumer's application for credit, insurance, or employment – or to take another adverse action against the consumer – must tell the consumer, and must give the consumer the name, address, and phone number of the agency that provided the information. Upon the request of a consumer for a credit score, a consumer reporting agency shall supply to the consumer a statement and notice that includes “all of the key factors that adversely affected the credit score of the consumer in the model used,” and any consumer reporting agency shall provide trained personnel to explain to the consumer any information required to be furnished to the consumer under the Act (15 U.S.C. §1681g (f) and (g); see *also* 15 U.S.C. §1681m for requirements of adverse action notices). And the Equal Credit Opportunity Act (“ECOA”), 15 U.S.C. § 1691 et seq. states:

(a) **ACTIVITIES CONSTITUTING DISCRIMINATION** It shall be unlawful for any creditor to discriminate against any applicant, with respect to any aspect of a credit transaction—

(1) on the basis of race, color, religion, national origin, sex or marital status, or age (provided the applicant has the capacity to contract);

(2) because all or part of the applicant's income derives from any public assistance program; or

(3) because the applicant has in good faith exercised any right under this chapter.

Ultimately, the need for lawyers to understand how AI generates outputs is important for combatting bias and providing good counsel to clients. And it may be required by legal ethics. As detailed above, lawyers have a duty to communicate with clients, and explaining why AI generates a particular outcome may be included as part of that duty. The good news is that while AI has the potential to be biased, AI is much more predictable than humans. It is easier to remedy bias in machines than it is in humans. Given their role as officers of the court, it is critical for lawyers to be on the forefront of understanding how bias in the use of AI can impact outcomes achieved by the legal profession and society as a whole.

V. QUESTIONS TO ASK WHEN ADOPTING AN AI SOLUTION OR ENGAGING AN AI VENDOR

Lawyers and courts will most likely adopt AI through their third party vendors. Before adopting such solutions, lawyers and courts should ask their vendors the following questions and ensure the vendors understand the following issues:

i. AI Bias, Explainability, and Transparency

- Before using AI, the technology should be determined not to have built-in bias due to its programming or its data.
- The lawyer and court should ensure that AI vendors providing the tool to the lawyer and court are aware of and take into account the potential for bias, including disparate impact.

112

Questions to ask:

- Can the result of the AI's decision be explained in a meaningful and lawful way to affected stakeholders, where appropriate?
- Is the training set examined to minimize potential of data bias?
- Do the AI's data and machine-learning operations reinforce bias? Do the operations fail to or give poor performance for certain segments of the population due to age, gender, race, ethnicity, etc.?
- Does the AI identify itself as AI where appropriate or required by law?

ii. Ethical and Beneficial

- AI, its production, and deployment should be beneficial (or at least not detrimental) to the lawyer, the court, clients, and society in general.
- Deployment of AI should take into account the needs and viewpoints of the lawyer's and court's various stakeholders (e.g., clients, plaintiffs, defendants).
- The use of AI should take into account accessibility for those with disabilities, both enhancing access where possible and minimizing impacts on the disabled (for example, an online chatbot provided by a court might also have a voice interface, or vice versa).
- The use of AI should align with the ethical codes and principles.

Questions to ask:

- Does AI promote civil activities, where appropriate (e.g., AI tools that do not hinder freedom of speech or assembly)?
- Depending on the industry, does AI accommodate diverse populations?

iii. Monitoring, Accountability, Controls, and Oversight

The lawyer and court should have control and oversight of AI vendors and what AI does and how it operates.

- The use of AI should be monitored for potential legal and ethical issues.
- AI should be designed to retain records and to allow for the re-creation of decision-making steps or processes, especially when accidents might occur.
- Legal counsel should be part of the process of accountability, controls, and oversight in order to protect the attorney-client privilege as well as to ensure legal compliance.
- AI and its usage should be audited and auditable.

Questions to ask:

- Is there a single lawyer, staff person, or officer, such as a Chief Artificial Intelligence Officer, who oversees the AI program?

- Does the lawyer or court understand AI and its risks?
 - Is the AI semi-autonomous or fully autonomous?
 - Does the AI incorporate machine learning or is it static?
 - Are people interacting directly with AI, and how?
- How does the lawyer or court know if the AI is operating properly?
- Is the keeping of AI data and decisions part of the lawyer's or court's records retention policy and obligations?

iv. Privacy

Because AI can often be used in monitoring people (such as workplace monitoring), and making decisions about people based on their personal information, it is important that the courts and lawyers address the privacy impact in using the AI. To the extent that lawyers and law firms are subject to privacy laws, an AI impact analysis may need to assess such usage's compliance with such laws, such as the GDPR.

VI. CONCLUSION

This resolution, if adopted, will urge lawyers and courts to address the emerging ethical and legal issues related to the usage of artificial intelligence in the practice of law as described in this report.

Respectfully submitted,

William B. Baker

Chair, Science & Technology Law Section
August 2019

GENERAL INFORMATION FORM

Submitting Entities: Science & Technology Law Section

Submitted By: William B. Baker, Chair, Science & Technology Law Section

1. Summary of Recommendation(s).
The American Bar Association urges courts and lawyers to address the emerging ethical and legal issues related to the usage of artificial intelligence (“AI”) in the practice of law including: (1) bias, explainability, and transparency of automated decisions made by AI; (2) ethical and beneficial usage of AI; and (3) controls and oversight of AI and the vendors that provide AI.
2. Approval by Submitting Entity.
Approved by Science & Technology Law Section on May 6, 2019.
3. Has this or a similar recommendation been submitted to the ABA House of Delegates or Board of Governors previously?
No
4. What existing Association policies are relevant to this recommendation and how would they be affected by its adoption?
In August 2012, the ABA amended Model Rule 1.1 of the Model Rules of Professional Conduct to add Comment 6, which states that a lawyer has a responsibility to keep abreast of the benefits and risks associated with using relevant technology.⁵⁰ This resolution urges action related to a specific type of technology, AI, that is or will become increasingly used in business and by lawyers.⁵¹
5. If this is a late Report, what urgency exists which requires action at this meeting of the House?
N/A
6. Status of Legislation. (If applicable.)
N/A
7. Brief explanation regarding plans for implementation of the policy, if adopted by the House of Delegates.

⁵⁰

https://www.americanbar.org/groups/professional_responsibility/publications/model_rules_of_professional_conduct/rule_1_1_competence/comment_on_rule_1_1/;

⁵¹ See generally <https://biglawbusiness.com/artificial-intelligence-creeps-into-big-law-endangers-some-jobs>

The Section of Science & Technology Law intends to study with interested ABA entities a possible model standard for legal and ethical usage of AI by courts and lawyers. This resolution could also be used by the ABA, as well as by ABA members to promote continuing legal education related to AI.

8. Cost to the Association. (Both direct and indirect costs.)
Adoption of this proposed resolution would result in only minor indirect costs associated with staff time devoted to arranging teleconference calls for Section members and other interested persons, as part of the staff members' overall substantive responsibilities.
9. Disclosure of Interest. (If applicable.)
None
10. Referrals. (List entities to which the recommendation has been referred, the date of referral and the response of each entity if known.)
This Report with Recommendations was circulated to the leadership of the ABA Section of Civil Rights and Social Justice, Innovation Center, Litigation, CPR, Judicial Division, GP Solo, and Law and National Security, and the Cyber Legal Task Force.
11. Contact Person. (Prior to the meeting. Please include name, address, telephone number and email address.)

William Baker
1300 Pennsylvania Ave, Suite 700,
Washington DC 20004,
571-317-1922,
wbaker@potomacclaw.com

Richard Field
755 Anderson Avenue
Cliffside Park, NJ 07010
201-941-8015
field@pipeline.com

Bonnie Fought
55 Tiptoe Lane Hillsborough,
CA 94010,
650-218-6248,
aba@garber-fought.net

Huu Nguyen
30 Rockefeller Plaza,
New York, NY 10112,
212-872-9802,

huu.nguyen@squirepb.com

Caryn Cross Hawk
Section Director
American Bar Association
321 North Clark
Chicago, IL 60654
312-988-5601
caryn.hawk@americanbar.org

12. Contact Person. (Who will present the report to the House. Please include email address and cell phone number.)

Richard Field
755 Anderson Avenue
Cliffside Park, NJ 07010
201-941-8015
field@pipeline.com

Bonnie Fought
55 Tiptoe Lane
Hillsborough, CA 94010
650-218-6248,
aba@garber-fought.net

EXECUTIVE SUMMARY1. Summary of the Recommendation.

The American Bar Association urges courts and lawyers to address the emerging ethical and legal issues related to the usage of artificial intelligence (“AI”) in the practice of law including: (1) bias, explainability, and transparency of automated decisions made by AI; (2) ethical and beneficial usage of AI; and (3) controls and oversight of AI and the vendors that provide AI.

2. Summary of the issue which the Recommendation addresses.

Artificial intelligence promises to change the practice of law. There are many different ways that lawyers today are using AI to improve productivity and provide better legal services to their clients, and the usage of AI tools in the legal profession will only increase. It is essential for lawyers to be aware of (a) how AI can be used in their practices, including who their ethical duties apply to the use of AI, (b) the problem of bias in the development and use of AI, and (c) proper control and oversight of the use of AI by lawyers and their vendors.

3. An explanation of how the proposed policy position will address the issue.

The proposed policy position will increase understanding in the legal profession of the legal and ethical issues posed by the usage of AI.

4. A summary of any minority views or opposition which have been identified.

N/A

Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools

Varun Magesh*
Stanford University

Faiz Surani*
Stanford University

Matthew Dahl
Yale University

Mirac Suzgun
Stanford University

Christopher D. Manning
Stanford University

Daniel E. Ho[†]
Stanford University

Abstract

Legal practice has witnessed a sharp rise in products incorporating artificial intelligence (AI). Such tools are designed to assist with a wide range of core legal tasks, from search and summarization of caselaw to document drafting. But the large language models used in these tools are prone to “hallucinate,” or make up false information, making their use risky in high-stakes domains. Recently, certain legal research providers have touted methods such as retrieval-augmented generation (RAG) as “eliminating” (Casetext, 2023) or “avoid[ing]” hallucinations (Thomson Reuters, 2023), or guaranteeing “hallucination-free” legal citations (LexisNexis, 2023). Because of the closed nature of these systems, systematically assessing these claims is challenging. In this article, we design and report on the first pre-registered empirical evaluation of AI-driven legal research tools. We demonstrate that the providers’ claims are overstated. While hallucinations are reduced relative to general-purpose chatbots (GPT-4), we find that the AI research tools made by LexisNexis (Lexis+ AI) and Thomson Reuters (Westlaw AI-Assisted Research and Ask Practical Law AI) each hallucinate between 17% and 33% of the time. We also document substantial differences between systems in responsiveness and accuracy. Our article makes four key contributions. It is the first to assess and report the performance of RAG-based proprietary legal AI tools. Second, it introduces a comprehensive, preregistered dataset for identifying and understanding vulnerabilities in these systems. Third, it proposes a clear typology for differentiating between hallucinations and accurate legal responses. Last, it provides evidence to inform the responsibilities of legal professionals in supervising and verifying AI outputs, which remains a central open question for the responsible integration of AI into law.¹

1 Introduction

In the legal profession, the recent integration of large language models (LLMs) into research and writing tools presents both unprecedented opportunities and significant challenges (Kite-Jackson, 2023). These systems promise to perform complex legal tasks, but their adoption remains hindered by a critical flaw: their tendency to generate incorrect or misleading information, a phenomenon generally known as “hallucination” (Dahl et al., 2024).

*Equal contribution.

[†]Corresponding author: deho@stanford.edu.

¹Our dataset, tool outputs, and labels will be made available upon publication. This version of the manuscript (June 6, 2024) is updated to reflect an evaluation of Westlaw’s AI-Assisted Research.

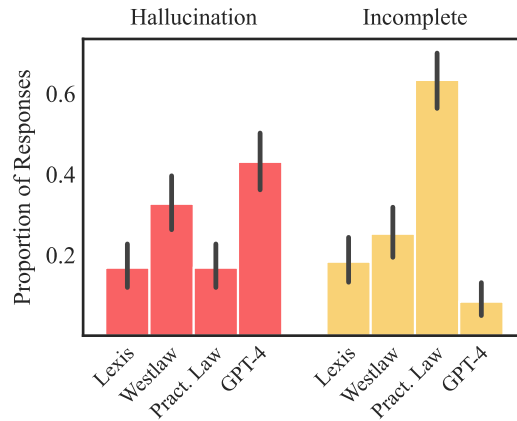


Figure 1: Comparison of hallucinated and incomplete answers across generative legal research tools. Hallucinated responses are those that include false statements or falsely assert a source supports a statement. Incomplete responses are those that fail to either address the user’s query or provide proper citations for factual claims.

As some lawyers have learned the hard way, hallucinations are not merely a theoretical concern (Weiser and Bromwich, 2023). In one highly-publicized case, a New York lawyer faced sanctions for citing ChatGPT-invented fictional cases in a legal brief (Weiser, 2023); many similar incidents have since been documented (Weiser and Bromwich, 2023). In his 2023 annual report on the judiciary, Chief Justice John Roberts specifically noted the risk of “hallucinations” as a barrier to the use of AI in legal practice (Roberts, 2023).

Recently, however, legal technology providers such as LexisNexis and Thomson Reuters (parent company of Westlaw) have claimed to mitigate, if not entirely solve, hallucination risk (LexisNexis, 2023; Casetext, 2023; Thomson Reuters, 2023, *inter alia*). They say their use of sophisticated techniques such as retrieval-augmented generation (RAG) largely prevents hallucination in legal research tasks.² (We provide details on RAG systems in Section 3.1 below.)

But none of these bold proclamations have been accompanied by empirical evidence. Moreover, the term “hallucination” itself is often left undefined in marketing materials, leading to confusion about which risks these tools genuinely mitigate. This study seeks to address these gaps by evaluating the performance of AI-driven legal research tools offered by LexisNexis (Lexis+ AI) and Thomson Reuters (Westlaw AI-Assisted Research and Ask Practical Law AI), and, for comparison, GPT-4.

Our findings, summarized in Figure 1, reveal a more nuanced reality than the one presented by these providers: while RAG appears to improve the performance of language models in answering legal queries, the hallucination problem persists at significant levels. To offer one simple example, shown in the top left panel of Figure 2, the Westlaw system claims that a paragraph in the Federal Rules of Bankruptcy Procedure (FRBP) states that deadlines are jurisdictional. But no such paragraph

²The following are official statements from Lexis, Casetext, and Thomson Reuters; however, none of them has provided any clear evidence so far to support their claims about the capabilities of their AI-based legal research tools:

Lexis: “Unlike other vendors, however, *Lexis+ AI delivers 100% hallucination-free linked legal citations* connected to source documents, grounding those responses in authoritative resources that can be relied upon with confidence.” (Wellen, 2024a) (emphasis added).

Casetext: “Unlike even the most advanced LLMs, *CoCounsel does not make up facts, or ‘hallucinate,’* because we’ve implemented controls to limit CoCounsel to answering from known, reliable data sources—such as our comprehensive, up-to-date database of case law, statutes, regulations, and codes—or not to answer at all.” (Casetext, 2023) (emphasis added).

Thomson Reuters: “*We avoid [hallucinations] by relying on the trusted content within Westlaw* and building in checks and balances that ensure our answers are grounded in good law.” (Thomson Reuters, 2023) (emphasis added). “We’ve all heard horror stories where generative AI just makes things up. That doesn’t work for the legal industry. They have to trust the content that AI serves up. With Ask Practical Law AI, *all the responses are based on the expert resources of Practical Law.*” (Thomson Reuters, 2024b) (emphasis added)

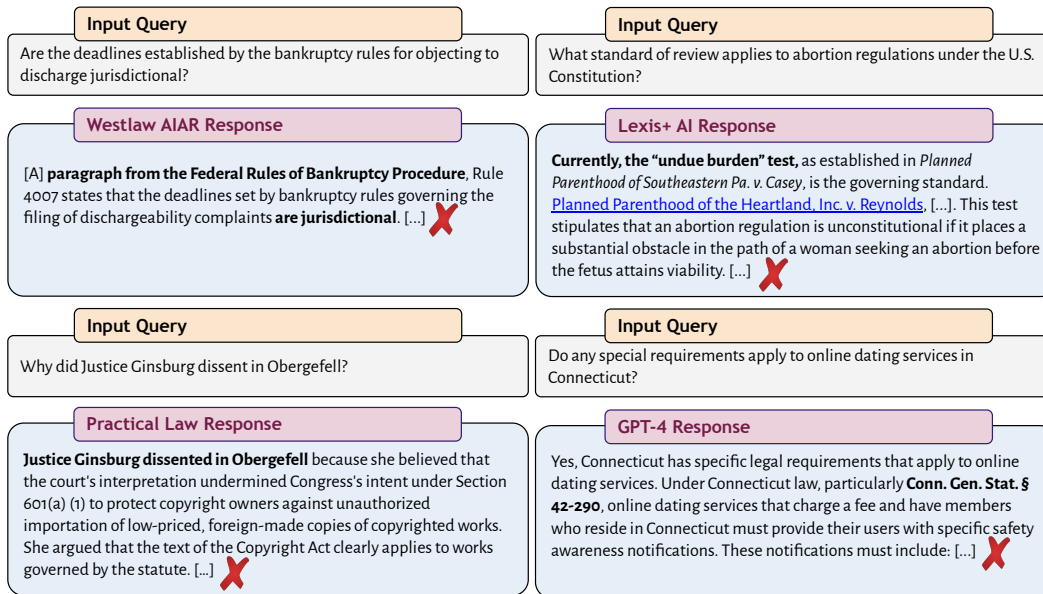


Figure 2: *Top left:* Example of a hallucinated response by Westlaw’s AI-Assisted Research product. The system makes up a statement in the Federal Rules of Bankruptcy Procedure that does not exist. *Top right:* Example of a hallucinated response by LexisNexis’s Lexis+ AI. *Casey* and its undue burden standard were overruled by the Supreme Court in *Dobbs v. Jackson Women’s Health Organization*, 597 U.S. 215 (2022); the correct answer is rational basis review. *Bottom left:* Example of a hallucinated response by Thomson Reuters’s Ask Practical Law AI. The system fails to correct the user’s mistaken premise—in reality, Justice Ginsburg joined the Court’s landmark decision legalizing same-sex marriage—and instead provides additional false information about the case. *Bottom right:* Example of a hallucinated response from GPT-4, which generates a statutory provision that does not exist.

exists, and the underlying claim is itself unlikely to be true in light of the Supreme Court’s holding in *Kontrick v. Ryan*, 540 U.S. 443, 447-48 & 448 n.3 (2004), which held that FRBP deadlines under a related provision were not jurisdictional.³

We also document substantial variation in system performance. LexisNexis’s Lexis+ AI is the highest-performing system we test, answering 65% of our queries accurately. Westlaw’s AI-Assisted Research is accurate 42% of the time, but hallucinates nearly twice as often as the other legal tools we test. And Thomson Reuters’s Ask Practical Law AI provides incomplete answers (refusals or ungrounded responses; see Section 4.3) on more than 60% of our queries, the highest rate among the systems we tested.

Our article makes four key contributions. First, we conduct the first systematic assessment of leading AI tools for real-world legal research tasks. Second, we manually construct a preregistered dataset of over 200 legal queries for identifying and understanding vulnerabilities in legal AI tools. We run these queries on LexisNexis (Lexis+ AI), Thomson Reuters (Ask Practical Law AI), Westlaw (AI-Assisted Research), and GPT-4 and manually review their outputs for accuracy and fidelity to authority. Third, we offer a detailed typology to refine the understanding of “hallucinations,” which enables us to rigorously assess the claims made by AI service providers. Last, we not only uncover limitations of current technologies, but also characterize the reasons that they fail. These results inform the responsibilities of legal professionals in supervising and verifying AI outputs, which remains an important open question for the responsible integration of AI into law.

The rest of this work is organized as follows. Section 2 provides an overview of the rise of AI in law and discusses the central challenge of hallucinations. Section 3 describes the potential and limitations of RAG systems to reduce hallucinations. Section 4 proposes a framework for evaluating

³We ran the queries for Lexis+ AI and Thomson Reuters Ask Practical Law AI in Figure 2 as a test prior to the creation of our benchmark dataset; because our queries for the evaluation presented in this article were preregistered, these two examples are not included in our results below.

hallucinations in a legal RAG system. Because legal research commonly requires the inclusion of citations, we define a *hallucination* as a response that contains either incorrect information or a false assertion that a source supports a proposition. Section 5 details our methodology to evaluate the performance of AI-based legal research tools (legal AI tools). Section 6 presents our results. We find that legal RAG can reduce hallucinations compared to general-purpose AI systems (here, GPT-4), but hallucinations remain substantial, wide-ranging, and potentially insidious. Section 7 discusses the limitations of our study and the challenges of evaluating proprietary legal AI systems, which have far more restrictive conditions of use than AI systems available in other domains. Section 8 discusses the implications for legal practice and legal AI companies. Section 9 concludes with implications of our findings for legal practice.

2 Background

2.1 The Rise and Risks of Legal AI

Lawyers are increasingly using AI to augment their legal practice, and with good reason: from drafting contracts, to analyzing discovery productions, to conducting legal research, these tools promise significant efficiency gains over traditional methods. As of January 2024, at least 41 of the top 100 largest law firms in the United States have begun to use some form of AI in their practice (Henry, 2024); among a broader sample of 384 firms, 35% now report working with at least one generative AI provider (Collens et al., 2024). And in a recent survey of 1,200 lawyers practicing in the United Kingdom, 14% say that they are using generative AI tools weekly or more often (Greenhill, 2024).

However, adoption of these tools is not without risk. Legal AI tools present unprecedented ethical challenges for lawyers, including concerns about client confidentiality, data protection, the introduction of new forms of bias, and lawyers’ ultimate duty of supervision over their work product (Avery et al., 2023; Cyphert, 2021; Walters, 2019; Yamane, 2020). Recognizing this, the bar associations of California (2023), New York (2024), and Florida (2024) have all recently published guidance on how AI should be safely and ethically integrated into their members’ legal practices. Courts have weighed in as well: as of May 2024, more than 25 federal judges have issued standing orders instructing attorneys to disclose or limit the use of AI in their courtrooms (Law360, 2024).

In order for these guidelines to be effective, however, lawyers need to first understand what exactly an AI tool is, how it works, and the ways in which it might expose them to liability. Do different tools have different error rates—and what kinds of errors are likely to manifest? What training do lawyers need in order to spot these errors—and can they do anything as users to mitigate them? Are there particular tasks that current AI tools are particularly adept at—and are there any that lawyers should stay away from?

This paper moves beyond previous work on general-purpose AI tools (Choi et al., 2024; Dahl et al., 2024; Schwarcz and Choi, 2023) by answering these questions specifically for *legal* AI tools—namely, the tools that have been carefully developed by leading legal technology companies and that are currently being marketed to lawyers as avoiding many of the risks known to exist in off-the-shelf offerings. In doing so, we aim to provide the concrete empirical information that lawyers need in order to assess the ethical and practical dangers of relying on these new commercial AI products.

2.2 The Hallucination Problem

We focus on one problem of AI that has received considerable attention in the legal community: “hallucination,” or the tendency of AI tools to produce outputs that are demonstrably false.⁴ In multiple high-profile cases, lawyers have been reprimanded for submitting filings to courts citing nonexistent case law hallucinated by an AI service (Weiser, 2023; Weiser and Bromwich, 2023). Previous work has found that general-purpose LLMs hallucinate on legal queries on average between 58% and 82% of the time (Dahl et al., 2024). Yet this prior work did not examine tools specifically developed for the legal setting, such as tools that use LLMs with auxiliary legal databases and RAG.

⁴Theoretical work has shown that hallucinations must occur at a certain rate for calibrated generative language models, regardless of their architecture, training data quality, or size (Kalai and Vempala, 2023).

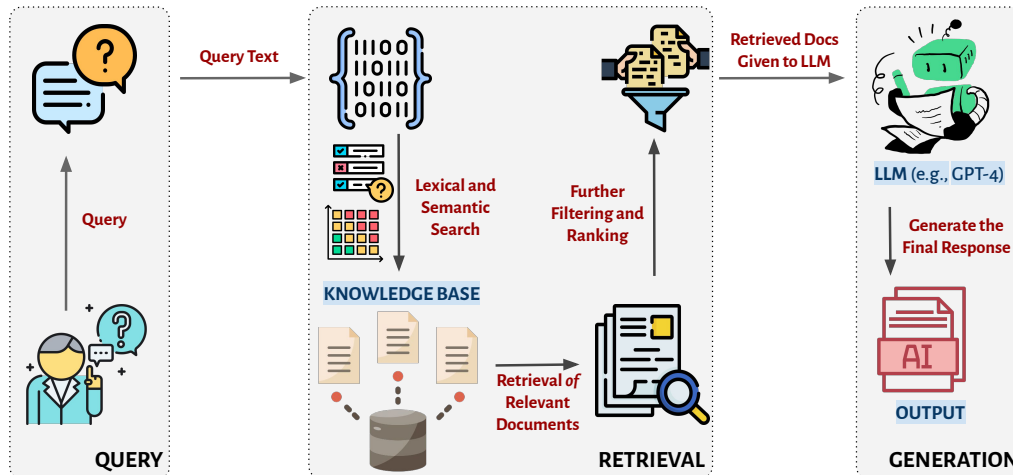


Figure 3: Schematic diagram of a retrieval-augmented generation (RAG) system. Given a user query (left), the typical process consists of two steps: (1) retrieval (middle), where the query is embedded with natural language processing and a retrieval system takes embeddings and retrieves the relevant documents (e.g., Supreme Court cases); and (2) generation (right), where the retrieved texts are fed to the language model to generate the response to the user query. Any of the subsidiary steps may introduce error and hallucinations into the generated response. (Icons are credited to FlatIcon.)

And because these tools are placed prominently before lawyers on leading legal research platforms (i.e., LexisNexis and Thomson Reuters / Westlaw), a systematic examination is sorely needed.

In this article, we focus on *factual* hallucinations. In the legal setting, there are three primary ways that a model can be said to hallucinate: it can be unfaithful to its training data, unfaithful to its prompt input, or unfaithful to the true facts of the world (Dahl et al., 2024). Because we are interested in legal research tools that are meant to help lawyers understand legal facts, we focus on the third category: factual hallucinations.⁵ However, in Section 4.3 below, we also expand on this definition by decomposing factual hallucinations into two dimensions: *correctness* and *groundedness*. We hope that this distinction will provide useful guidance for users seeking to understand the precise way that these tools can be helpful or harmful.

3 Retrieval-Augmented Generation (RAG)

3.1 The Promise of RAG

Across many domains, the fairly new technique of retrieval-augmented generation (RAG) is being seen and heavily promoted as the key technology for making LLMs effective in domain-specific contexts. It allows general LLMs to make effective use of company- or domain-specific data and to produce more detailed and accurate answers by drawing directly from retrieved text. In particular, RAG is commonly touted as the solution for legal hallucinations. In a February 2024 interview, a Thomson Reuters executive asserted that, within Westlaw AI-Assisted Research, RAG “dramatically reduces hallucinations to nearly zero” (Ambrogio, 2024). Similarly, LexisNexis has said that RAG enables it to “deliver accurate and authoritative answers that are grounded in the closed universe of authoritative content” (Wellen, 2024b).⁶

As depicted in Figure 3, RAG comprises two primary steps to transform a query into a response: (1) retrieval and (2) generation (Lewis et al., 2020; Gao et al., 2024). Retrieval is the process of selecting

⁵Other definitions of hallucination could be more relevant in other contexts. For example, future research should examine AI tools for contract analysis or document summarization. For that analysis, it would be more important to study hallucinations with respect to the tool’s input prompt, rather than with respect to the general facts of the world. Evaluation standards for such generative AI output, however, are still in flux.

⁶In Section 4.3 below, we discuss how different companies may be using definitions of “hallucination” different from the ones more commonly accepted in the literature or in popular discourse.

relevant documents from a large universe of documents. This process is familiar to anyone who uses a search engine: using keywords, user information, and other context, a search engine quickly identifies a handful of relevant web pages out of the millions available on the internet. Retrieval systems can be simple, like a keyword search, or complex, involving machine learning techniques to capture the semantic meaning of a query (such as neural text embeddings).

With the retrieved documents in hand, the second step of generation involves providing those documents to a LLM along with the text of the original query, allowing the LLM to use *both* to generate a response. Many RAG systems involve additional pre- and post-processing of their inputs and outputs (e.g., filtering and extraction depicted in the middle panel of Figure 3), but retrieval and generation are the hallmarks of a RAG pipeline.

The advantage of RAG is obvious: including retrieved information in the prompt allows the model to respond in an “open-book” setting rather than in “closed-book” one. The LLM can use the information in the retrieved documents to inform its response, rather than its hazy internal knowledge. Instead of generating text that conforms to the general trends of a highly compressed representation of its training data, the LLM can rely on the full text of the relevant information that is injected directly into its prompt.

For example, suppose that an LLM is asked to state the year that *Brown v. Board of Education* was decided. In a closed-book setting, the LLM, without access to an external knowledge base, would generate an answer purely based on its internal knowledge learned during training—but a more obscure case might have little or no information present in the training data, and the model could generate a realistic-sounding year that may or may not be accurate. In a RAG system, by contrast, the retriever would first look up the case name in a legal database, retrieve the relevant metadata, and then provide that to the LLM, which would use the result to provide the user a response to their query.

On paper, RAG has the potential to substantially mitigate many of the kinds of legal hallucinations that are known to afflict off-the-shelf LLMs (Dahl et al., 2024)—the technique performs well in many general question-answering situations (Guu et al., 2020; Lewis et al., 2020; Siriwardhana et al., 2023). However, as we show in the next section, RAG systems are no panacea.

3.2 Limitations of RAG

There are several reasons that RAG is unlikely to fully solve the hallucination problem (Barnett et al., 2024). Here, we highlight some that are unique to the legal domain.

First, retrieval is particularly challenging in law. Many popular LLM benchmarking datasets (Rajpurkar et al., 2016; Yang et al., 2018) contain questions with clear, unambiguous references that address the question in the source database. Legal queries, however, often do not admit a single, clear-cut answer (Mik, 2024). In a common law system, case law is created over time by judges writing opinions; this precedent then builds on precedent in the way that a chain novel might be written in *seriatim* (Dworkin, 1986). By construction, these legal opinions are not atomic *facts*; indeed, on some views, the law is an “essentially contested” concept (Waldron, 2002). Thus, deciding what to retrieve can be challenging in a legal setting. At best, a RAG system must be able to locate information from multiple sources across time and place in order to properly answer a query. And at worst, there may be no set of available documents that definitively answers the query, if the question presented is novel or indeterminate.

Second, document relevance in the legal context is not based on text alone. Most retrieval systems identify relevant documents based on some kind of text similarity (Karpukhin et al., 2020). But the retrieval of documents that only *seem* textually relevant—and are ultimately irrelevant, or “distracting”—negatively affects performance on general question-answering tasks (Cuconasu et al., 2024; Chen et al., 2024). Problems of this type are likely to compound in the legal domain. In different jurisdictions and in different time periods, the applicable rule or the relevant jurisprudence may differ. Even similar-sounding text in the correct time and place may not apply if special conditions are not met. The problem may be worse if a rule that applies in a special condition conflicts with a more broadly applicable rule. The LLM may have been trained on a much greater volume of text supporting the broadly applicable rule, and may be more faithful to its training data than to the retrieval context. Consequently, designing a high-quality research tool that deals with this problem requires careful attention to non-textual elements of retrieval and the deference of the model to different sources of

information. These challenges are familiar to lawyers using “natural language” searches, as opposed to more deterministic boolean searches, on legal research platforms.

Third, the generation of meaningful legal text is also far from straightforward. Legal documents are generally written for other lawyers immersed in the same issue, and they rely on an immense amount of background knowledge to properly understand and apply. A helpful generative legal research tool would have to do far more than simple document summarization; it would need to synthesize facts, holdings, and rules from different pieces of text while keeping the appropriate legal context in mind. For example, consider this Lexis+ AI exchange:

Prompt: What are some notable opinions written by Judge Luther A. Wilgarten?

Lexis+ AI Response: One of the notable opinions written by Judge Luther A. Wilgarten is the case of Luther v. Locke. The case citation is 2010 U.S. Dist. LEXIS 144868 and it was decided in the Eastern District of Virginia on January 11, 2010. In this case, the plaintiff had filed a Notice of Appeal of Judge Ellis’s decision, but failed to properly prosecute the appeal. [...]

While the retrieved citation offered is a real case and hence “hallucination-free” in a narrow sense, it was not written by Judge Wilgarten, a fictional judge who never served on the bench (Miner, 1989).⁷ And while the generated passages are based on the actual case, the second sentence contradicts the premise, suggesting Judge *Ellis* wrote the opinion, but the opinion was actually written by Judge Brinkema (and involved a prior decision by Judge Ellis, which forms the basis for the RAG response). Nor is the decision notable, as it was an unpublished opinion cited only once outside of its direct history. Hallucinations are compounded by poor retrieval and erroneous generation.

Conceptualizing the potential failure modes of legal RAG systems requires domain expertise in both computer science *and* law. As is apparent once we examine the component parts of a RAG system in Figure 3, each of the subsidiary steps (the embedding, the design of lexical and semantic search, the number of documents retrieved, and filtering and extraction) involves design choices that can affect the quality of output (Barnett et al., 2024), each with potentially subtle trade-offs (Belkin, 2008). In the next section, we devise a new task suite specifically designed to probe the prevalence of RAG-resistant hallucinations, complementing existing benchmarking efforts that target AI’s legal knowledge in general (Dahl et al., 2024) and its capacity for legal reasoning (Guha et al., 2023).

4 Conceptualizing Legal Hallucinations

The binary notion of hallucination developed in Dahl et al. (2024) does not fully capture the behavior of RAG systems, which are intended to generate information that is both accurate and grounded in retrieved documents. We expand the framework of legal hallucinations to *two* primary dimensions: correctness and groundedness. Correctness refers to the factual accuracy of the tool’s response (Section 4.1). Groundedness refers to the relationship between the model’s response and its cited sources (Section 4.2).

Decomposing factual hallucinations in this way enables a more nuanced analysis and understanding of how exactly legal AI tools fail in practice. For example, a response could be correct but improperly grounded. This might happen when retrieval results are poor or irrelevant, but the model happens to produce the correct answer, falsely asserting that an unrelated source supports its conclusion. This can mislead the user in potentially dangerous ways.

4.1 Correctness

We say that a response is *correct* if it is both factually correct and relevant to the query. A response is *incorrect* if it contains any factually inaccurate information. For the purposes of this analysis, we label an answer that is partially correct—that is, one that contains correct information that does not fully address the question—as correct. If a response is neither correct nor incorrect, because

⁷This retrieval error likely reflects the similarity in the embedding space between “Judge Luther A. Wilgarten” and the terms “judge” (mentioned 9 times in the 900-some word order) and “William Luther,” the plaintiff in the case.

	Description	Example
<i>Correctness</i>		
Correct	Response is factually correct and relevant	The right to same sex marriage is protected under the U.S. Constitution. <i>Obergefell v. Hodges</i> , 576 U.S. 644 (2015).
Incorrect	Response contains factually inaccurate information	There is no right to same sex marriage in the United States.
Refusal	Model refuses to provide any answer or provides an irrelevant answer	I’m sorry, but I cannot answer that question. Please try a different query.
<i>Groundedness</i>		
Grounded	Key factual propositions make valid references to relevant legal documents	The right to same sex marriage is protected under the U.S. Constitution. <i>Obergefell v. Hodges</i> , 576 U.S. 644 (2015).
Misgrounded	Key factual propositions are cited but the source does not support the claim	The right to same sex marriage is protected under the U.S. Constitution. <i>Miranda v. Arizona</i> , 384 U.S. 436 (1966).
Ungrounded	Key factual propositions are not cited	The right to same sex marriage is protected under the U.S. Constitution.

Table 1: A summary of our coding criteria for correctness and groundedness, along with hypothetical responses to the query “Does the Constitution protect a right to same sex marriage?” that would fall under each of the categories. Groundedness is only applicable for correct responses. The categories which qualify as a “hallucination” are highlighted in red.

the model simply declines to respond, we label that as a *refusal*. See the top panel of Table 1 for examples of each of these three codings of correctness.⁸

4.2 Groundedness

For correct responses, we additionally evaluate each response’s groundedness. A response is *grounded* if the key factual propositions in its response make valid references to relevant legal documents. A response is *ungrounded* if key factual propositions are not cited. A response is *misgrounded* if key factual propositions are cited but misinterpret the source or reference an inapplicable source. See the bottom panel of Table 1 for examples illustrating groundedness.

Note that our use of the term *grounded* deviates somewhat from the notion in computer science. In the computer science literature, groundedness refers to adherence to the source documents provided, regardless of the relevance or accuracy of the provided documents (Agrawal et al., 2023). In this paper, by contrast, we evaluate the quality of the retrieval system and the generation model together in the legal context. Therefore, when we say *grounded*, we mean it in the legal sense—that is, responses that are correctly grounded in actual governing caselaw. If the retrieval system provides documents that are inappropriate to the jurisdiction of interest, and the model cites them in its response, we call that *misgrounded*, even though this might be a technically “grounded” response in the computer-science sense.

4.3 Hallucination

We now adopt a precise definition of a hallucination in terms of the above variables. A response is considered *hallucinated* if it is either incorrect or misgrounded. In other words, if a model makes a false statement or falsely asserts that a source supports a statement, that constitutes a hallucination.

⁸Note that for our false premise questions, the desired behavior is for the model to refute and state the false assumption in the user’s prompt. A gold-standard response to such a question would therefore be a statement that the assumption may be incorrect, with a case law citation to the opposite proposition. However, for these false premise questions alone, we also label a refusal which mentions the fact that no pertinent sources were found as correct.

This definition provides technical clarity to the popular concept of hallucination, which is a term that is currently being used inconsistently by different industry actors. For example, in one interview, one Thomson Reuters executive appeared to refer to hallucinations as exclusively instances when an AI system fabricates the *existence* of a case, statute, or regulation, distinct from more general problems of accuracy (Ambrogi, 2024). Yet, in a December 2023 press release, another Thomson Reuters executive defined hallucinations differently, as “responses that sound plausible but are completely false” (Thomson Reuters, 2023).

LexisNexis, by contrast, uses the term hallucination in yet a different way. LexisNexis claims that its AI tool provides “linked hallucination-free legal citations” (LexisNexis, 2023), but, as we demonstrate below, this claim can only be true in the most narrow sense of “hallucination,” in that their tool does indeed *link* to real legal documents.⁹ If those linked sources are irrelevant, or even contradict the AI tool’s claims, the tool has, in our sense, engaged in a hallucination. Failing to capture that dimension of hallucination would require us to conclude that a tool that links only to *Brown v. Board of Education* on every query (or provides cases for fictional judges as in the instance of Luther A. Wilgarten) has provided “hallucination-free” citations, a plainly irrational result.

More concretely, consider the *Casey* example in Figure 2, where the linked citation *Planned Parenthood v. Reynolds* is a real case that has not been overturned.¹⁰ However, the model’s answer relies on *Reynolds*’ description of *Planned Parenthood v. Casey*, a case that has been overturned. The model’s response is incorrect, and its citation serves only to mislead the user about the reliability of its answer (Goddard et al., 2012).

These errors are potentially more dangerous than fabricating a case outright, because they are subtler and more difficult to spot.¹¹ Checking for these kinds of hallucinations requires users to click through to cited references, read and understand the relevant sources, assess their authority, and compare them to the propositions the model seeks to support. Our definition reflects this more complete understanding of “hallucination.”

4.4 Accuracy and Incompleteness

Alongside *hallucinations*, we also define two other top-level labels in terms of our correctness and groundedness variables: *accurate responses*, which are those that are both correct and grounded, and *incomplete responses*, which are those that are either refusals or ungrounded.

We code correct but ungrounded responses as incomplete because, unlike a misgrounded response, an ungrounded response does not actually make any false assertions. Because an ungrounded response does not provide key information (supporting authorities) that the user needs, it is marked incomplete.

5 Methodology

5.1 AI-Driven Legal Research Tools

We study the hallucination rate and response quality of three available RAG-based AI research tools: LexisNexis’s Lexis+ AI, Thomson Reuters’s Ask Practical Law AI, and Westlaw’s AI-Assisted Research. As nearly every practicing U.S. lawyer knows, Thomson Reuters (the parent company of Westlaw) and LexisNexis¹² have historically enjoyed a virtual duopoly over the legal research market (Arewa, 2006) and continue to be two of the largest incumbents now selling legal AI products (Ma et al., 2024).

Lexis+ AI functions as a standard chatbot interface, like ChatGPT, with a text area for the user to enter an open-ended inquiry. In contrast to traditional forms of legal search, “boolean” connectors and search functions like AND, OR, and W/n are neither required nor supported. Instead, the user simply formulates their query in natural language, and the model responds in kind. The user then has

⁹Of course, there is some evidence that Lexis+ AI does not succeed even by this metric. McGreel (2024) reports instances of Lexis+ AI citing cases decided in 2025.

¹⁰*Reynolds* even appears in the citation list with a positive Shepardization symbol.

¹¹As Gottlieb (2024) reports in one of the assessments by law firms of generative AI products, “The importance of reviewing and verifying the accuracy of the output, including checking the AI’s answers against other sources, makes any efficiency gains difficult to measure.”

¹²LexisNexis is owned by the RELX Group.

the option to continue the chat by asking another question, which the tool will respond to with the complete context of both questions. Introduced in October 2023, Lexis+ AI states that it has access to LexisNexis’s entire repository of case law, codes, rules, constitution, agency decisions, treatises, and practical guidance, all of which it presumably uses to craft its responses. While not much technical detail is published, it is known that Lexis+ AI implements a proprietary RAG system that ensures that every prompt “undergoes a minimum of five crucial checkpoints . . . to produce the highest quality answer” (Wellen, 2024a).¹³

Ask Practical Law AI, introduced in January 2024 and offered on the Westlaw platform, is a more limited product, but it operates in a similar way. Like Lexis+ AI, Ask Practical Law AI also functions as a chatbot, allowing the user to input their queries in natural language and responding to them in the same format. However, instead of accessing all the primary sources that Lexis+ AI uses, Ask Practical Law AI only retrieves information from Thomson Reuters’s database of “practical law” documents—“expert resources . . . that have been created and curated by more than 650 bar-admitted attorney editors” (Thomson Reuters, 2024b) promising “90,000+ total resources across 17 practice areas” (Thomson Reuters, 2024a). Thomson Reuters markets this database for general legal research: “Practical Law provides trusted, up-to-date legal know-how across all major practice areas to help attorneys deliver accurate answers quickly and confidently.” Performing RAG on these materials, Thomson Reuters claims, ensures that its system “only returns information from [this] universe” (Thomson Reuters, 2024b).

Westlaw’s AI-Assisted Research (AI-AR), introduced in November 2023, is also a standard chatbot interface, promising “answers to a far broader array of questions than what we could anticipate with human power alone” (Thomson Reuters, 2023). The RAG system retrieves information from Westlaw’s databases of cases, statutes, regulations, West Key Numbers, headnotes, and KeyCite markers (Thomson Reuters, 2023). While not much technical detail is provided, AI-AR appears to rely on OpenAI’s GPT-4 system (Ambrogi, 2023). This system was built out after a \$650 million acquisition of Casetext, which had developed legal research systems on top of GPT-4 (Ambrogi, 2023). RAG is prominently touted as addressing hallucinations: one Thomson Reuters official stated, “We avoid [hallucinations] by relying on the trusted content within Westlaw and building in checks and balances that ensure our answers are grounded in good law” (Thomson Reuters, 2023). While AI-AR has been sold to law firms, it has not been made available generally for educational and research purposes.¹⁴

Both AI-AR and Ask Practical Law AI are made available via the Westlaw platform and are commonly referred to as AI products within Westlaw.¹⁵ For shorthand, we will refer to Ask Practical Law AI as a Thomson Reuters system and AI-AR as a Westlaw system, as this appears to track the internal company product distinctions.

To provide a point of reference for the quality of these bespoke legal research tools—and because AI-AR appears to be built on top of GPT-4—we also evaluate the hallucination rate and response quality of GPT-4, a widely available LLM that has been adopted as a knowledge-work assistant (Dell’Acqua et al., 2023; Collens et al., 2024). GPT-4’s responses are produced in a “closed-book” setting; that is, produced without access to an external knowledge base.

¹³Since the completion of our evaluation for this paper in April 2024, LexisNexis has released a “second generation” version of its tool. Our results do not speak to the performance of this second generation product, if different. Accompanying this release, LexisNexis noted, “our promise is not perfection, but that all linked legal citations are hallucination-free” (LexisNexis, 2024).

¹⁴Thomson Reuters denied three requests for access by our team at the time we conducted our initial evaluation. The company provided access after the initial release of our results.

¹⁵The home page of Practical Law is titled “Practical Law US - Westlaw” and is located on a subdomain of westlaw.com (Google, 2024). See also, e.g., Berkeley Law School (2024) (noting that “Ask Practical Law AI” is now available on Westlaw”); Yale Law School (2024) (describing “Ask Practical Law AI” as a Westlaw product); University of Washington (2024) (describing “Practic[al] Law [a]s a database within Westlaw”); Suffolk University (2023) (noting “Ask Practical Law AI (Westlaw)”); Campbell (2024) (writing that “Westlaw released Ask Practical Law AI to academic accounts”).

Category	Count	Perc.	Description	Example Query
General legal research	80	39.6%	Common-law doctrine questions, previously published practice bar exam questions, holding questions	Has a habeas petitioner’s claim been “adjudicated on the merits” for purposes of 28 U.S.C. § 2254(d) where the state court denied relief in an explained decision but did not expressly acknowledge a federal-law basis for the claim?
Jurisdiction or time-specific	70	34.7%	Questions about circuit splits, overturned cases, or new developments	In the Sixth Circuit, does the Americans with Disabilities Act require employers to accommodate an employee’s disability that creates difficulties commuting to work?
False premise	22	10.9%	Questions where the user has a mistaken understanding of the law	I’m looking for a case that stands for the proposition that a pedestrian can be charged with theft for absorbing sunlight that would otherwise fall on solar panels, thereby depriving the owner of the panels of potential energy.
Factual recall questions	30	14.9%	Basic queries about facts not requiring interpretation, like the year a case was decided.	Who wrote the majority opinion in <i>Candela Laser Corp. v. Cynosure, Inc.</i> , 862 F. Supp. 632 (D. Mass. 1994)?

Table 2: The high-level categories of the query dataset, with counts and percentages (Perc.) of queries, descriptions, and sample queries.

5.2 Query Construction

We design a diverse set of legal queries to probe different aspects of a legal RAG system’s performance. We develop this benchmark dataset to represent real-life legal research scenarios, without prior knowledge of whether they would succeed or fail.

For ease of interpretation, we group our queries into four broad categories:

1. **General legal research questions:** common-law doctrine questions, holding questions, or bar exam questions
2. **Jurisdiction or time-specific questions:** questions about circuit splits, overturned cases, or new developments
3. **False premise questions:** questions where the user has a mistaken understanding of the law
4. **Factual recall questions:** queries about facts of cases not requiring interpretation, such as the author of an opinion, and matters of legal citation

Queries in the first category ($n = 80$) are the paradigmatic use case for these tools, asking general questions of law. For instance, such queries pose bar exam questions that have ground-truth answers, but in contrast to assessments that focus only on the accuracy of the multiple choice answer (e.g., [Martínez, 2024](#)), we assess hallucinations in the fully generated response. Queries in the second category ($n = 70$) probe for jurisdictional differences or developing areas in the law, which represent precisely the kinds of active legal questions requiring up-to-date legal research. Queries in the third category ($n = 22$) probe for the tendency of LLMs to assume that premises in the query are true, even when flatly false. The last category ($n = 30$) probes the extent to which RAG systems are able to overcome known vulnerabilities about how general LLMs encode legal knowledge ([Dahl et al., 2024](#)).

Table 2 describes these categories in more depth and provides an example of a question that falls within each category. We used 20 queries from LegalBench’s Rule QA task verbatim ([Guha et al., 2023](#)), and 20 BARBRI bar exam prep questions verbatim ([BARBRI, Inc., 2013](#)). Each of the 162

other queries were hand-written or adapted for use in our benchmark. Appendix A provides a more granular list of the types of queries and descriptive information.

Our dataset advances AI benchmarking in five respects. First, it is expressly designed to move the evaluation of AI systems from standard question-answer settings with a discrete and known answer (e.g., multiple choice) to the generative (e.g., open-ended) setting (Raji et al., 2021; Li and Flanigan, 2024; McIntosh et al., 2024). Prior work has evaluated the amount of legal information that LLMs can produce (Dahl et al., 2024), but this kind of benchmark does not capture the practical benefits and risks of everyday use cases. Legal practice is more than answering multiple choice questions. Of course, because these are not simple queries, their design and evaluation is time-intensive—all queries must be written based on external legal knowledge and submitted by hand through the providers’ web interfaces, and evaluation of answers requires careful assessment of the tool’s legal analysis and citations, which can be voluminous.

Second, our queries are specifically tailored to RAG-based, open-ended legal research tools. This differentiates our dataset from previously released legal benchmarks, like LegalBench (Guha et al., 2023). Most LegalBench tasks are tailored towards legal analysis of information given to the model in the prompt; tasks like contract analysis or issue spotting. Our queries are written specifically for RAG-based legal research tools; each query is an open-ended legal question that requires legal analysis supported by relevant legal documents that the model must retrieve. This provides a more realistic representation of the way that lawyers are intended to use these tools. Our goal with our dataset is to move beyond anecdotal accounts and offer a systematic investigation of the potential strengths and weaknesses of these tools, responding to documented challenges in evaluating AI in law (Kapoor et al., 2024; Guha et al., 2023).

Third, these queries are designed to represent the temporal and jurisdictional variation (e.g., overruled precedents, circuit splits) that is often the subject of live legal research (Beim and Rader, 2019). We hypothesize that AI systems are not able to encode this type of multifaceted and dynamic knowledge at the moment, but these are precisely the kinds of inquiries requiring legal research. Due to the nature of legal authority, attorneys will inevitably have questions specific to their time, place, and facts, and even the most experienced lawyers will need to ground their understanding of the legal landscape when facing issues of first impression.

Fourth, the queries probe for “contrafactual bias,” or the tendency of chat systems to assume the veracity of a premise even when false (Dahl et al., 2024). Many claim that AI systems will help to address longstanding access to justice issues (Bommasani et al., 2022; Chien et al., 2024; Chien and Kim, 2024; Perlman, 2023; Tan et al., 2023), but contrafactual bias poses particular risk for *pro se* litigants and lay parties.

Last, to guard against selection bias in our results (i.e., choosing queries based on hallucination results), we modeled best practices with our dataset by preregistering our study and associated queries with the Open Science Foundation prior to performing our evaluation (Surani et al., 2024).¹⁶

5.3 Query Execution

For Lexis+ AI, Thomson Reuters’s Ask Practical Law AI, and Westlaw’s AI-AR, we executed each query by copying and pasting it into the chat window of each product.¹⁷ For GPT-4, we prompted the LLM via the OpenAI API (model gpt-4-turbo-2024-04-09) with the following instruction, appending the query afterwards:

You are a helpful assistant that answers legal questions. Do not hedge unless absolutely necessary, and be sure to answer questions precisely and cite caselaw for propositions.

This prompt aims to ensure comparability with legal AI tools, particularly by prompting for legal citations and concrete factual assertions. We recorded the complete response that each tool gave,

¹⁶We did not run any preregistered query against any tool prior to registration, with one exception, changes-in-law-73 (“When does the undue burden standard apply in abortion cases?”). Some queries were slightly rephrased during evaluation to better elicit an answer with factual content (a prospect explicitly contemplated by the pre-registration); those queries are marked as such in our released dataset and documented in Appendix B.1.

¹⁷We created a new “conversation” for each query.

along with any references to case law or documents. The dataset was preregistered on March 22, 2024 and all queries on Lexis+ AI, Ask Practical Law AI, and GPT-4 were run between March 22 and April 22, 2024. Queries on Westlaw’s AI-AR system were run between May 23–27, 2024.

5.4 Inter-Rater Reliability

To code each response according to the concepts of correctness, groundedness, and hallucination, we relied on our expert domain knowledge in law to hand-score each model response according to the rubric developed in Section 4. As noted above, efficiently evaluating AI-generated text remains an unsolved problem with inevitable trade-offs between internal validity, external validity, replicability, and speed (Liu et al., 2016; Hashimoto et al., 2019; Smith et al., 2022). These problems are particularly pronounced in our legal setting, where our queries represent real legal tasks. Accordingly, techniques of letting these legal AI tools “check themselves”—which have become popular in other AI evaluation pipelines (Manakul et al., 2023; Mündler et al., 2023; Zheng et al., 2023)—are not suitable for this application. Precisely because adherence to authority is so important in legal writing and research, our tasks must be qualitatively evaluated by hand according to the definitions of correctness and groundedness that we have carefully constructed. This makes studying these legal AI tools expensive and time-consuming: this is a cost that must be reflected in future conversations about how to responsibly integrate these AI products into legal workflows.

To ensure that our queries are sufficiently well-defined and that our coding definitions are sufficiently precise, we evaluated the inter-rater reliability of different labelers on our data. Task responses were first graded by one of three different labelers. A fourth labeler then labeled a random sample of 48 responses, stratified by model and task type. We oversampled *The Bluebook* citation task slightly because it is particularly technical. The fourth labeler did not discuss anything with the first three labelers and did not have access to the initial labels. Their knowledge of the labeling process came only from our written documentation of labeling criteria, fully described in Appendix D.

With this protocol, we find a Cohen’s kappa (Cohen, 1960) of 0.77 and an inter-rater agreement of 85.4% on the final outcome label (correct, incomplete, or hallucinated) between the evaluation labeler and the initial labels. This is a substantial degree of agreement that suggests that our task and taxonomy of labels are well defined. Our results are comparable to similar evaluations for complex, hand-graded legal tasks (Dahl et al., 2024).¹⁸

6 Results

Section 6.1 describes our findings on hallucinations and responsiveness. Section 6.2 examines the varied and sometimes insidious nature of hallucinations. Section 6.3 provides a typology of the potential causes of inaccuracies we encountered.

6.1 Hallucinations Persist Across Query Types

Commercially-available RAG-based legal research tools still hallucinate. Over 1 in 6 of our queries caused Lexis+ AI and Ask Practical Law AI to respond with misleading or false information. And Westlaw hallucinated substantially more—*one-third of its responses* contained a hallucination.

On the positive side, these systems are less prone to hallucination than GPT-4, but users of these products must remain cautious about relying on their outputs.

The left panel of Figure 4 provides a breakdown of response types across the four products. Lexis+ AI’s answers are accurate (i.e., correct and grounded) for 65% of queries, compared to much lower accuracy rates of 41% and 19% by Westlaw and Practical Law AI, respectively. The right panel of Figure 4 also provides the hallucination rate when an answer is responsive, showing that Lexis+ AI appears to have a statistically significantly lower hallucination rate than Westlaw and Thomson Reuters, even conditional on a response.

¹⁸In updating results to include AI-AR, we also conducted another round of validation of every hallucination coding. This validation led to nearly identical results—for instance, the accuracy rate of Ask Practical Law AI in Figure 1 increased from 19% to 20%, which is of course within the bounds of inter-rater reliability.

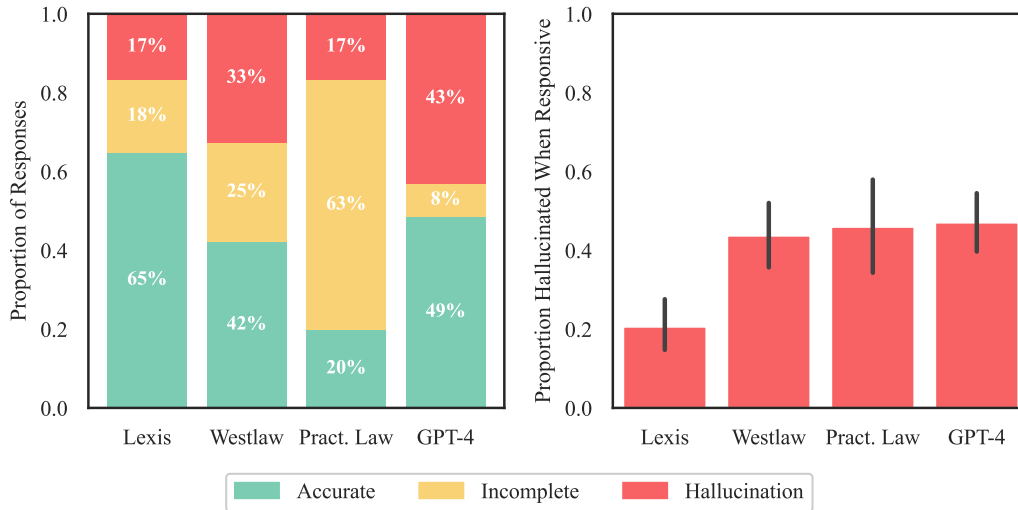


Figure 4: *Left panel:* overall percentages of accurate, incomplete, and hallucinated responses. *Right panel:* the percentage of answers that are hallucinated when a direct response is given. Westlaw AI-AR and Ask Practical Law AI respond to fewer queries than GPT-4, but the responses that they do produce are not significantly more trustworthy. Vertical bars denote 95% confidence intervals.

Figure 5 also breaks down these statistics by query type. We observe that, while hallucination rates are slightly higher for jurisdiction and time specific questions, they remain high for general legal research questions, such as questions posed on the bar exam. Accuracy rates are highest on “false premise” questions—in which the query contains a mistaken understanding of law—and lower on the categories which represent real-world use by attorneys.

Westlaw’s high hallucination rate is driven by several kinds of errors (as discussed further in Section 6.2), but we note that it is also the system which tends to generate the longest answers. Excluding refusals to answer, Westlaw has an average word length of 350 (SD = 120), compared to 219 (SD = 114) by Lexis+ AI and 175 (SD = 67) by Ask Practical Law AI.¹⁹ With longer answers, Westlaw contains more falsifiable propositions and therefore has a greater chance of containing at least one hallucination. Lengthier answers also require substantially more time to check, verify, and validate, as every proposition and citation has to be independently evaluated.

Responsiveness differs dramatically across systems. As shown in Figure 4, Lexis+ AI, Westlaw AI-AR, and Ask Practical Law AI provide incomplete answers 18%, 25% and 62% of the time, respectively. The low responsiveness of Ask Practical Law AI can be explained by its more limited universe of documents. Rather than connecting its retrieval system to the general body of law (including cases, statutes, and regulations), Ask Practical Law AI draws solely from articles about legal practice written by its in-house team of lawyers.

On the other hand, the Westlaw and Lexis retrieval systems are connected to a wider body of case law and primary sources. This means that they have access to all the documents that are in principle necessary to answer any of our questions. Both systems often offer high-quality responses. In one instance, Lexis+ AI pointed to a false premise in one of our questions. The question `scalr-19` asked whether the six year statute of limitation applied to retaliatory discharge actions under the False Claims Act. The question was drawn from *Graham County Soil & Water Conservation District v. U.S.*, 559 U.S. 280 (2010), where the Court held that there was ambiguity. Congress moved thereafter to amend the statute to clarify the statute of limitations. Lexis+ AI explained the mistaken premise, and cited the relevant, updated code section. Similarly, when prompted about the need for specific, proven “teaching, suggestion, or motivation” (TSM) that would have led a person of ordinary skill in the art to combine the relevant prior art for a finding of obviousness, AI-AR correctly responded by discussing the Supreme Court’s decision in *KSR v. Teleflex*, 550 U.S. 398 (2007), which rejected a rigid notion of the Federal Circuit’s TSM test.

¹⁹This is based on a simple word count separating based on space.

Westlaw AI-AR Hallucination	Explanation
1. “[W]hen the lender receives the collateral that secured the fraudulent loan, this is considered a return of ‘any part’ of the loan money . . . This was established in the Supreme Court case <i>Robers v. U.S.</i> ”	<i>Robers</i> held precisely the opposite: “the phrase ‘any part of the property . . . returned’ refers to the property the banks lost . . . and not to the collateral.” 572 U.S. 639, 642 (2014).
2. “ <i>D.M. v. State</i> . . . has been overruled by <i>Davis v. State</i> . Also, the case <i>Millbrook v. U.S.</i> was reversed by the same case at a later date.”	<i>Millbrook v. United States</i> is a U.S. Supreme Court decision that controls on federal questions. 569 U.S. 50 (2013). The Nebraska Supreme Court did not cite, much less ‘reverse,’ it in <i>Davis v. State</i> . 297 Neb. 955 (2017).
3. “[A] paragraph from the Federal Rules of Bankruptcy Procedure [FRBP], Rule 4007 states that the deadlines set by bankruptcy rules governing the filing of dischargeability complaints are jurisdictional .”	There is no such paragraph in the FRBP and 4007’s deadlines are unlikely to be jurisdictional under <i>Kontrick v. Ryan</i> , 540 U.S. 443, 447–48, 448 n.3 (2004).
4. “The Supreme Court of Alaska . . . noted . . . in the absence of physical injury, the losses for claimed emotional damage [from a misdiagnosis] should not be permitted for legal and public policy reasons (<i>Chizmar v. Mackie</i> , 896 P.2d 196 (1995)).”	The Supreme Court of Alaska in fact reversed this trial court holding (that losses are not permitted on public policy grounds) and allowed the question to proceed to the jury. 896 P.2d at 204-05.
5. “[I]n the <i>Doo v. Packwood</i> case, a restrictive covenant precluding the operation of a grocery store in a building conveyed was found to be potentially inequitable and unenforceable . . . (<i>Doo v. Packwood</i> , 265 Cal.App.2d 752 (1968)).”	The <i>Doo</i> court in fact held the opposite. It reversed the trial court’s finding that the covenant was unenforceable. 265 Cal.App.2d at 756, 759.
6. “[T]he FCC’s new policy regarding ‘ fleeting expletives ’ is arbitrary and capricious under the Administrative Procedure Act.”	The Supreme Court in fact found the opposite: the fleeting expletives policy was not arbitrary and capricious. <i>FCC v. Fox</i> , 556 U.S. 502 (2009).
7. “In the <i>Reid v. Fargo</i> case, there is no mention or reference made to the <i>Irvine v. The Hesper</i> case. ”	The <i>Reid</i> court expressly relied on the <i>Hesper</i> case. 241 U.S. 544, 548 (1916).
8. “In the case of <i>Bragg v. Linden Research, Inc.</i> , the court recognized participants’ full intellectual property protection for the digital content they created or owned in the game <i>Second Life</i> . . . <i>Bragg v. Linden Research, Inc.</i> , 487 F.Supp.2d 593 (2007).”	The court took no such position and was only describing actions by the defendant (<i>Linden</i>): “ <i>Linden</i> announced that it would recognize participants’ full intellectual property protection.” 487 F.Supp.2d at 595.
9. “In <i>Washington Legal Foundation v. Legal Foundation of Washington</i> , the Court of Appeals for the Ninth Circuit held that the interest generated by IOLTA pooled trust accounts is the property of the clients, and the appropriation of that interest for charitable purposes was a per se ‘taking’ within the meaning of the Fifth Amendment.”	This statement falsely credits the holding of a Ninth Circuit panel instead of the superseding holding of the Ninth Circuit en banc court, which overruled the panel. 271 F.3d 835, 857-58 (applying ad hoc, not per se takings analysis).
10. “ 18 U.S.C.A. § 3583 . . . was recognized as unconstitutional by <i>United States v. Winter Rose Old Rock</i> .	The <i>Winter Rose</i> court in fact rejected the defendant’s claim that § 3583 was unconstitutional, 76 F.4th 1314, 1316 (10th Cir. 2023), and only discussed and distinguished <i>Haymond</i> , 588 U.S. 634 (2019) (4-1-4 decision with judgment that specific application of § 3583(k) was unconstitutional).

Table 3: Ten examples of hallucinations in Westlaw’s AI-Assisted Research responses, with explanations for why they are coded as hallucinations.

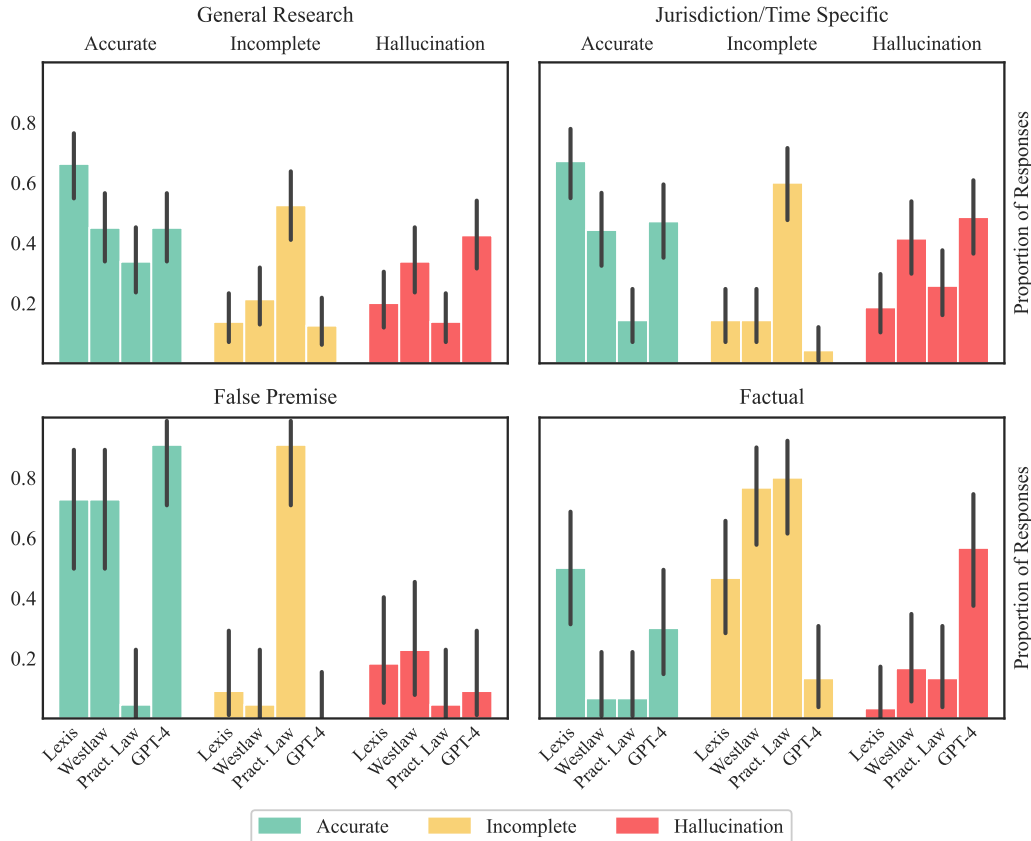


Figure 5: Response evaluations broken down by question category. We show the accuracy (green), incompleteness (yellow), and hallucination (red) rate for each question category. Vertical bars denote 95% confidence intervals. This figure shows that hallucinations are not driven by an isolated category and persist across task types and questions, such as bar exam and appellate litigation issues.

6.2 Hallucinations Can Be Insidious

These systems can be quite helpful when they work. But as we now illustrate in detail, their answers are often significantly flawed. We find that these systems continue to struggle with elementary legal comprehension: describing the holding of a case (Zheng et al., 2021), distinguishing between legal actors (e.g., between the arguments of a litigant and the holding of the court), and respecting the hierarchy of legal authority. Identifying these misunderstandings often requires close analysis of cited sources. These vulnerabilities remain problematic for AI adoption in a profession that requires precision, clarity, and fidelity.

Tables 3, 4, and 5 provide examples of hallucinations in the Westlaw, Lexis, and Practical Law systems, respectively.²⁰ In each example, our detailed analysis of responses and cited cases reveals a serious inaccuracy and hallucination in the system response. The following sections refer to examples in these tables to illustrate different failure modes in legal RAG systems.

Misunderstanding Holdings. Systems do not seem capable of consistently making out the holding of a case. This is a serious issue, as legal research relies centrally on distinguishing the holding from other parts of the case. Table 3 rows 1, 4, 5, and 6 provide examples of when Westlaw states a summary that is the direct opposite of the actual holding of a case, including case by the U.S. Supreme Court. For instance, Westlaw states that collateral is considered a return of “any part” of the loan, indicating that this was established by the Supreme Court in *Roberts v. U.S.*, but *Roberts* held the exact opposite (Table 3 row 1). In another response, Lexis+ AI recites Missouri legislation

²⁰The number of examples reported are roughly proportional to the relative hallucination rates between tools.

criminalizing unauthorized camping on state-owned lands. But that legislation comes from the statement of facts and analysis, and in the cited case, the Missouri Supreme Court actually held that legislation unconstitutional (Table 4, row 3).

Distinguishing Between Legal Actors. Systems can fail to distinguish between arguments made by litigants and statements by the court. In one example, Westlaw attributes an action of the defendant to the court (Table 3 row 8) and in another it stated that a provision of the U.S. Code was found unconstitutional by the 10th Circuit, when in fact the 10th Circuit rejected that argument by the defendant (Table 3, row 10).

Respecting the Order of Authority. All models strain in grasping hierarchies of legal authority. This is crucial, as courts often discuss similar propositions that may be in tension. When sources conflict, a complex system of precedence and hierarchy determines governing law. Sorting through different sources to find the authoritative ones requires legal “background knowledge” about the way that different courts interact in different jurisdictions, and even systems with direct access to case law can fail to adhere to these legal hierarchies. For example:

- Westlaw asserts that a U.S. Supreme Court case was reversed by the *Nebraska Supreme Court* on a matter of federal law. That is not possible in the U.S. legal system, and in fact the Nebraska Supreme Court did not so much as cite the Supreme Court case in question (Table 3 row 2).
- Westlaw confuses holdings between different levels of courts (Table 3 rows 5, 9). In row 9, for instance, Westlaw properly states the holding of the Ninth Circuit panel, but improperly attributes it to the Ninth Circuit sitting en banc, which actually overruled the panel on that issue.
- Lexis+ AI fails to distinguish between district and appellate courts. In Table 4 row 1, Lexis+ AI transmogrifies a district court recitation of the trial court standard for awarding attorney’s fees into a patently incorrect standard of *appellate review* of attorney’s fees (incorrectly stating that an appeals court may disturb attorney’s fees “as long as they provide reasoning”).
- In Table 4 row 2, Lexis+ AI describes a rule established in *Arturo D.* as good law, with citation to the case that actually overrules *Arturo D.*

We note one additional area where systems struggle with orders of authority. In numerous instances, we observed the Westlaw system stating a proposition based on an overruled or reversed case, *without* citing the case. These errors may stem from design choices: Westlaw may be adding citations in a second pass, after generating the statement, while suppressing the citation of cases that receive a “red flag” under its KeyCite system.²¹ For instance, when prompted about the equity clean-up doctrine, which allows courts of equity to decide legal and equity issues when it has jurisdiction over the equity issues, AI-AR properly cites the rule, but then notes, “However, this general rule does not apply when the facts relied on to sustain the equity jurisdiction fail of establishment.” This statement is unaccompanied by an in-text citation; the language appears only in a search result below the response, in a Missouri case²² that was overruled on that issue by the Missouri Supreme Court.²³ We believe this suppression behavior can be dangerous—it impedes verification of the claims most likely to be false.

Fabrications. The systems we test occasionally generate text that is unrelated or deviates materially from retrieved documents.

- Westlaw generates provisions of law that do not exist. For instance, it asserts that the Federal Rules of Bankruptcy (FRBP) state that deadlines are jurisdictional, which is not a statement contained in the FRBP text at all (Table 3 row 3). (The hallucination seems to emanate

²¹Per Westlaw, a red flag indicates that a case “is no longer good law for at least one of the points of law it contains”(Thomson Reuters, 2019). In our labeled sample, we were not able to observe such cases being cited, though they were sometimes discussed without citation.

²²State ex rel. Leonardi v. Sherry, No. ED 82789, 2003 WL 21384384, at *1 (Mo. Ct. App. June 17, 2003).

²³See State ex rel. Leonardi v. Sherry, 137 S.W.3d 462, 472 (Mo. 2004) (“The dissenting opinion apparently would cling to the inefficient and wasteful need for a second trial at law if equity ‘fails of establishment’ in the initial request for equitable relief.”).

Lexis+ AI Hallucination	Explanation
1. “In <i>Vadatech v. Overnight</i> [2016 Nev. Dist. LEXIS 1619], the court awarded attorney’s fees and costs to the prevailing party. It cited Nevada Supreme Court precedent that appellate courts can modify fee awards as long as they provide reasoning , without needing to find an abuse of discretion.”	<i>Vadatech</i> , a district court opinion, said no such thing and instead recited the trial court standard for granting attorney’s fees—set by the Nevada Supreme Court—not the standard of appellate modification of such fees.
2. “If a driver is unable to produce identification documents or produces documents that appear to be false, a limited search of places in the vehicle where the driver may have hidden a wallet is reasonable. In the case of Arturo D., the court upheld the warrantless vehicle search when the driver failed to produce the required documentation <i>People v. Lopez</i> , 8 Cal. 5th 353.”	<i>Lopez</i> actually overturned <i>Arturo D.</i> . Lexis+ AI describes <i>Arturo D.</i> as if it is good law and supported by <i>Lopez</i> .
3. “The new section 67.2300 of TAFP HB 1606 , which was signed into law by Governor Parson . . . makes the act of unauthorized sleeping and camping on state-owned lands a class C misdemeanor . <i>Byrd v. State</i> , 679 S.W.3d 492.”	HB 1606 was in fact struck down by the Missouri Supreme Court in the cited case <i>Byrd v. State</i> , 679 S.W.3d 492 (Mo. 2023). This detail is not mentioned in Lexis+ AI’s response.
4. “The equity clean-up doctrine is a self-imposed ordinance that closes the doors of a court of equity to one tainted with an inequity or bad faith relative to the matter in which he seeks relief. . . . <i>EEOC v. Recruit U.S.A., Inc.</i> , 939 F.2d 746”	This statement confuses the equity clean-up doctrine with the “clean hands” doctrine . <i>EEOC v. Recruit U.S.A.</i> , 939 F.2d 746, 752 (9th Cir. 1991).
5. “The Telephone Consumer Protection Act of 1991 (TCPA) grants federal courts exclusive jurisdiction over actions brought by the state attorneys general, but places exclusive jurisdiction over private civil actions in state courts .”	The Supreme Court in fact found concurrent state and federal court jurisdiction over private suits under the TCPA. <i>Mims v. Arrow Financial Services</i> , 565 U.S. 368 (2012).
6. “The Eleventh Circuit has not definitively ruled on the causation standard for Family and Medical Leave Act (FMLA) retaliation claims . Some courts have applied the “but-for” causation standard, which requires a showing that the adverse action would not have occurred but for the protected activity <i>Lapham v. Walgreen Co.</i> , 88 F.4th 879 . . . However, other courts have applied a motivating factor or negative factor standard. . . . <i>Lee v. Yellowpages.Com LLC</i> , 2019 U.S. Dist. LEXIS 250502, <i>Nelson v. Americold Logistics, LLC</i> , 2020 U.S. Dist. LEXIS 63611.”	Lexis claims the standard is not definitive, but <i>Lapham</i> is the most recent controlling case, and the other cases contesting the issue are older inferior court rulings.

Table 4: Six examples of hallucinations in Lexis+ AI responses with explanations for why they are coded as hallucinations.

from a retrieved 1996 bankruptcy court case, which is also likely invalid under the Supreme Court’s *Kontrick* decision, which found that bankruptcy deadlines are not jurisdictional.)

- Westlaw misinterprets the Supreme Court’s specific holding on a statutory *subsection* as the 10th Circuit finding the entire statutory section unconstitutional, when in fact the 10th Circuit rejected the defendant’s claim of unconstitutionality (Table 3 row 10).
- Lexis+ AI attributes a description of the equity clean-up doctrine to a case that only discusses the “clean hands” doctrine (Table 4 row 4).

Practical Law Hallucination	Explanation
1. “Based on 28 U.S.C § 2680(h), the United States does waive its sovereign immunity for certain intentional torts committed by federal investigative or law enforcement officers acting within the scope of their employment. However, this waiver applies only if these officers are acting in their capacity to search, seize evidence, or make arrests for federal crimes [Federal Tort Claims Act (FTCA): Exceptions and Defenses]”	The cited Practical Law resource states: “ However, the tortious act need not occur while the officer is searching, seizing, or arresting for the proviso to apply (Millbrook v. United States, 569 U.S. 50, 55-57 (2013))”
2. “Yes, a person subject to domestic-violence-related restraining orders can be prohibited from purchasing a firearm in the Fifth Circuit . [State and Local Regulation of Firearms Under the Second Amendment]”	The Practical Law article actually states: “Federal law prohibits a person who is subject to a domestic violence restraining order from possessing a firearm. . . . However, the US Court of Appeals for the Fifth Circuit has held that this law violates the Second Amendment (United States v. Rahimi, 61 F.4th 443 (5th Cir. 2023))”
3. “A dying declaration is admissible in federal court under the Federal Rules of Evidence (FRE) 804(b)(2) . . . This provision is commonly invoked in homicide prosecutions, where the victim’s last words about the identity or description of their killer are introduced into evidence. However, it could theoretically apply to any case where these specific conditions are met. ”	The dying declaration exception is not just “commonly” invoked in homicide prosecutions; FRE 804(b)(2) specifically limits the admissibility of dying declarations to cases involving “a prosecution for homicide or in a civil case,” not any case.
4. “The undue burden standard applies in abortion cases when assessing the constitutionality of state regulations on abortion before a fetus becomes viable. . . . [Health Plan Coverage of Surgical and Medication Abortion and Related Services, Group Health Plan Coverage Considerations for Employers After the Overturning of Roe v. Wade, Employee Assistance Program (EAP) Compliance]”	Dobbs v. Jackson Women’s Health, 597 U.S. 215 (2022) overturns Casey and the undue burden standard. All three Practical Law articles discuss <i>Roe</i> and <i>Casey</i> ’s reversal, with one title including the phrase “Overturning of <i>Roe</i> .”

Table 5: Four examples of hallucinations in Thomson Reuters’s Ask Practical Law AI response, with explanations of why they are coded as hallucinations. The Practical Law documents cited are named in square brackets.

6.3 A Typology of Legal RAG Errors

Interpreting why an LLM hallucinates is an open problem (Ji et al., 2023; Zou et al., 2023a). While it is possible to identify correlates of hallucination (Dahl et al., 2024), it is hard to conclusively explain why a model hallucinates on one question but not another, or why one model hallucinates where another does not.

RAG systems, however, are composed of multiple discrete components (Gao et al., 2024). While each piece may be a black box, due to the lack of documentation by providers, we can partially observe the way that information moves between them. Lexis+ AI, Ask Practical Law AI, and AI-AR each show the list of documents which were retrieved and given to the model (though not exactly which pieces of text are passed in). Consequently, comparing the retrieved documents and the written response allows us to develop likely explanations for the reasons for hallucination.

In this section, we present a typology of different causes of RAG-related hallucination that we observe in our dataset. Other analyses of RAG failure points identify a larger number of distinct failure points (Barnett et al., 2024; Chen et al., 2024). Our typology collapses some of these, since we focus on broader causes that can be identified using the limited information we have about the systems we test. Our typology also introduces new failure points unique to the legal context that have not previously been considered in analyses of general-purpose RAG systems. Evaluations of general purpose RAG systems often assume that all retrievable documents (1) contain true information and (2)

are authoritative and applicable, an assumption that is not true in the legal setting (Barnett et al., 2024; Chen et al., 2024).²⁴ Legal documents often contain outdated information, and their relevance varies by jurisdiction, time period, statute, and procedural posture. Determining whether a document is binding or persuasive often requires non-trivial reasoning about its content, metadata, and relationship with the user query.

This typology is intended to be useful to both legal researchers and AI developers. For legal researchers, it illustrates some pathways to incorrect outputs, and highlights specific areas of caution. For developers, it highlights areas for improvement in these tools. The categories that we present are not mutually exclusive; the failures we observe are often driven by multiple causes or have unclear causes. Table 6 compares the prevalence of different hallucination causes in our typology. Because these are closed systems, we are not able to clearly identify a single point of failure for each hallucination.

Contributing Cause	Lexis	Westlaw	Pract. Law
Naive Retrieval	0.47	0.20	0.34
Inapplicable Authority	0.38	0.23	0.34
Reasoning Error	0.28	0.61	0.49
Sycophancy	0.06	0.00	0.03

Table 6: This table shows prevalence of different contributing causes among all hallucinated responses for each model. Because the types are not mutually exclusive, the proportions do not sum to 1.

Naive retrieval. Many failures in the three systems stem from poor retrieval—failing to find the most relevant sources available to address the user’s query. For instance, when asked to define the “moral wrong doctrine,” a doctrine pertaining to mistake-of-fact instructions in criminal prosecutions for morally wrongful acts (doctrine-test-177), Lexis+ AI relies on a source which defines moral *turpitude*, a legal term of art with a seemingly similar but actually unrelated meaning.

Part of the challenge is that retrieval itself often requires legal reasoning. As Section 3.2 discusses, legal sources are not composed of unambiguous facts. Lawyers are often taught to analyze situations with an IRAC framework—first identify the issue (I) and governing legal rule (R), then analyze (A) the facts with that rule to arrive at a conclusion (C) (Guha et al., 2023). For example, bar-exam-96 asks whether an airline’s motion to dismiss should be granted in a wrongful death suit arising out of a plane crash. Ask Practical Law AI retrieves sources discussing motions to dismiss in various contexts such as bankruptcy and patent litigation. But correctly answering this question requires identifying the true underlying issue as being one about *tort negligence*, not general procedures for motions to dismiss. Thomson Reuters’s tool likely errs because it fails to perform this analytical step prior to querying its database, thereby ending up with sources pertaining to the wrong issue.

Inapplicable authority. An inapplicable authority error occurs when a model cites or discusses a document that is not legally applicable to the query. This can be because the authority is for the wrong jurisdiction, wrong statute, wrong court, or has been overruled. This kind of error is uniquely important and prevalent in the legal setting, and has not been explored as thoroughly in prior literature (Barnett et al., 2024; Gao et al., 2024). One example is Lexis+ AI’s response to scalar-15. This question asks about certain deadlines under Bankruptcy Rule 4004, but the model describes and cites a case about tax court deadlines under 26 U.S.C.S. § 6213(a) instead. This could be because the excerpt of the case that is given to the model does not include key information, or because the model was given that information and ignored it. Because it is not possible to see exactly what information is available to the model, it is not possible to say precisely where the error occurs.

Sycophancy. LLM assistants have been found to display “sycophancy,” a tendency to agree with the user even when the user is mistaken (Sharma et al., 2023). While sycophancy can cause hallucinations (Dahl et al., 2024), we found that Lexis+ AI, AI-AR, and GPT-4 were quite capable at navigating our false premise queries, and often corrected the false premise without hallucination. For example,

²⁴Chen et al. (2024) consider the possibility of retrievable documents that contain false information. However, its evaluation focuses on a significantly simplified setting that is not applicable to the complexity of legal use cases.

false-holding-statements-108 asks for a case showing that due process rights can be violated by negligent government action. Lexis+ AI steers the user towards the correct answer, stating that intentional interference can violate due process, and that negligent interference cannot, supporting these propositions with case law. Ask Practical Law AI also seldom hallucinated in this category, but refused to answer at all in the overwhelming majority of queries.

Reasoning errors. In addition to the more complex behaviors described above, LLM-based systems also tend to make elementary errors of reasoning and fact. The legal research systems we test are no exception. We observe such errors most frequently in Westlaw; though retrieved results often seemed relevant and helpful, the model would not always correctly reason through the text to arrive at the correct conclusion. In one instance (Table 3 row 8), AI-AR describes a district court decision as “recogniz[ing] participant’s full intellectual property protection for the digital content they created or owned in the game Second Life. . .” But as the passage cited by the model makes clear, the court held no such thing. It was describing the statements of the *defendant*, and the language model made a simple factual error in describing the passage given to it.

7 Limitations

While our study provides critical information about widely deployed AI tools in legal practice, it comes with certain limitations.

First, our evaluation is limited to three specific products by LexisNexis, Thomson Reuters, and Westlaw. The legal AI product space is growing rapidly with many startups (e.g., Harvey, Vincent AI) (Ma et al., 2024). Access to these emerging systems is even more restricted than to the services offered in LexisNexis and Westlaw, making evaluation exceptionally challenging.²⁵ That said, our approach provides a common benchmark that can be deployed for similar systems as they become available.

Second, our evaluation only captures a point in time. Even over the course of our study, we noticed the responses of these systems—particularly Lexis+ AI—evolve over time. While these changes may improve responses, we note that benchmarking, evaluation, and supervision remain difficult when a model changes over time (Chen et al., 2023).²⁶ This is compounded by uncertainty over whether such differences are driven by changes in the base model (e.g., GPT-4) or by engineering by the legal technology provider. More generally, a fundamental concern for the evaluation of LLMs lies in test leakage—because language models are trained on all available data, they may memorize data that is used for evaluation (Li and Flanigan, 2024; Oren et al., 2023; Deng et al., 2023). That is a particularly challenging concern when the only mechanism for accessing legal AI tools is by sending test prompts to providers themselves. Even if providers fix the discrete errors noted above, that may not mean that the problems we identify have been solved in general.²⁷

Third, while we have been able to design an effective evaluation framework for chat-based interfaces, the evaluation for more specified generative tasks is still evolving. LegalBench (Guha et al., 2023), for instance, still requires manual evaluation of certain generative outputs, and we do not here assess Casetext CoCounsel’s effectiveness at drafting open-ended legal memoranda. Developing benchmarks for the full range of legal tasks—e.g., deposition summaries, legal memoranda, contract review—remains an important open challenge for the field (Kapoor et al., 2024).

Fourth, although we designed the first benchmark dataset, the sample size of 202 queries remains small in comparison to other evaluations such as Dahl et al. (2024). There are two reasons for this. In contrast to general-purpose LLMs, which have open models or API access, LexisNexis, Thomson

²⁵Even AI-Assisted Research was exclusively available to law firms when we initially conducted the evaluation of Lexis+ AI and Ask Practical Law AI (Thomson Reuters, 2023).

²⁶Indeed, even presenting the same query to these models may yield different answers each time, as the text decoding process may not be set to be deterministic (e.g., via the temperature parameter). GPT-4, for instance, is known not to be deterministic. It is also not clear what retrieval parameters (e.g., similarity threshold or top-*k* value) are used, impeding consistent analysis of the model.

²⁷For instance, OpenAI appeared to patch its system to prevent adversarial attacks with specific suffixes discovered in Zou et al. (2023b), but the underlying vulnerability may still persist. As one of the authors of that study noted, “Companies like OpenAI have just patched the suffixes in the paper, but numerous other prompts acquired during training remain effective. Moreover, if the model weights are updated, repeating the same procedure on the new model would likely still work.”

Reuters, and Westlaw restrict access to their interfaces.²⁸ In addition, extensive *manual* work is required to evaluate the results of each query, making it harder to scale automated evaluations. The trend toward LLM-based evaluations may address the latter obstacle, but the fact remains that the legal AI product space remains quite closed.

Fifth, while we managed to develop a measurement protocol that yielded substantial agreement between human raters, we acknowledge that groundedness may exist on a spectrum. A citation, for instance, might point to a case that has been overruled, but that case might still be helpful to an attorney in starting the research process. In our setting, we coded such instances as misgrounded, but whether the model is helpful will still fundamentally have to be determined by use cases and evaluations that involve human interactions with the system. The range of failure points documented in Section 6.3 provides a more granular sense of the limitations of current AI systems.

Sixth, some might argue that our benchmark dataset does not represent the natural distribution of queries. We designed our benchmark to reflect a wide range of query types and to constitute a challenging real-world dataset. Questions are ones that arise on the bar exam, that arise in appellate litigation, that present circuit splits, that present issues that are dynamically changing, and that were contributed by the legal community (Guha et al., 2023). The benchmark was designed to be challenging precisely because (a) those are the settings where legal research is needed the most, and (b) it responds to the marketing claims by providers. It is true that these may not represent all tasks for which lawyers turn to generative AI. Our estimate of the hallucination rate is not meant to be an unbiased estimate of the (unknown) population-level rate of hallucinations in legal AI queries, but rather to assess whether hallucinations have in fact been solved by RAG, as claimed. We show that hallucinations persist across the wide range of task types (see Figure 1) and the full natural distribution of such queries is (a) only known to legal technology providers, (b) highly in flux given uncertainty about the appropriate use of AI in law, and (c) itself endogenous to assessments of reliability and marketing claims.

Last, our primary goal is limited to assessing the hallucination rate, accuracy, and groundedness on emerging legal technology. These are central concepts to the trustworthiness of AI tools, but they are not the sole criteria for the quality and value of a legal research system. For instance, notwithstanding the many hidden hallucinations, the overall output of Lexis+ AI and AI-AR may still be quite valuable for distinct use cases (e.g., starting on a research thread). But evaluations like the one we designed here are critical to understanding these appropriate use cases.

8 Implications

Excitement over the potential for AI to transform the practice of law is at an all-time high. On the demand side, lawyers fear missing out on the real gains in efficiency and thoroughness that new AI tools can offer. On the supply side, the companies developing these tools continue to market them as more and more powerful (Markelius et al., 2024). We agree that these tools are hugely promising (Chien and Kim, 2024; Choi et al., 2024), but our research has important implications for both the lawyers using these products and the myriad of companies now marketing them.

8.1 Implications for Legal Practice

In the United States, all lawyers are required to abide by certain professional and ethical rules. Most jurisdictions have adopted a version of the Model Rules of Professional Conduct, which are issued by the American Bar Association (American Bar Association, 2018). Two of these rules bear directly on the integration of AI into law: Rule 1.1’s duty of competence and Rule 5.3’s duty of supervision (Cyphert, 2021; Walters, 2019; Yamane, 2020). Competence requires “legal knowledge, skill, thoroughness and preparation” (Rule 1.1); supervision requires “reasonable efforts to ensure that the [non-lawyer’s] conduct is compatible with the professional obligations of the lawyer” (Rule 5.3).

In addition to these rules, the bar associations of New York (2024), California (2023), and Florida (2024) have all recently published more detailed guidance on how lawyers’ ethical responsibilities intersect with their use of AI. For example, the New York State Bar Association’s AI Task Force

²⁸See, for example, § 2.2 of the LexisNexis Terms of Service (LexisNexis, 2023), which prohibits programmatic access.

states that lawyers “have a duty to understand the benefits, risks and ethical implications” associated with the tools that they use (2024, 57); similarly, the State Bar of California’s Standing Committee on Professional Responsibility and Conduct implores lawyers to “understand the risks and benefits of the technology used in connection with providing legal services” (2023, 1).

In other words, lawyers’ ability to comply with their professional duties in both of these jurisdictions is contingent on access to *specific* information about empirical risks and benefits of legal AI. Yet, so far, no legal AI company has provided this information. The New York State Bar Association points its members to a list of publications and fora that discuss matters related to AI in general (2024, 76-77), but general knowledge is not the same as understanding the trade-offs of specific tools.

Indeed, our work shows that the risks and benefits associated with AI-driven legal research tools are different from those associated with general-purpose chatbots like GPT-4. As we discuss in Section 6, the tools we study in this article differ in responsiveness and accuracy, and these differences may even change over time within the same tool. The closed nature of these tools, however, makes it difficult for lawyers to assess when it is safe to trust them. Official documentation does not clearly illustrate what they can do for lawyers and in which areas lawyers should exercise caution. Thus, given the high rate of hallucinations that we uncover in this article, lawyers are faced with a difficult choice: either verify by hand each and every proposition and citation produced by these tools (thereby undercutting the efficiency gains that AI is promised to provide), or risk using these tools without full information about their specific risks and benefits (thereby neglecting their core duties of competency and supervision).

8.2 Implications for Legal AI Companies

Legal AI developers face dilemmas as well. On the one hand, these companies are subject to economic pressures to compete in an increasingly crowded market (Ma et al., 2024), pressures made more acute by the recent entry of previously copyrighted and proprietary data into the public domain (Henderson et al., 2022; Östling et al., 2024; The Library Innovation Lab, 2024). On the other hand, like all businesses, they are also constrained by laws and regulations limiting the products they can lawfully offer and advertise. We flag two of these potential restrictions here.

First, companies must be careful not to overclaim or misrepresent the abilities of their AI products. As we discuss in Section 1, a number of legal AI providers are currently making claims about their products’ ability to “eliminat[e]” (Casetext, 2023) or “avoid” hallucinations (Thomson Reuters, 2023), yet, as we note in Section 4.3, these same companies are inconsistently using the term “hallucination” in ways that may not conform to users’ expectations. Without additional precision about the exact mistakes that their tools purportedly avoid, companies may find themselves exposed to civil liability for unfair competition or false, misleading, or unsubstantiated claims. For instance, under Section 43(a) of the Lanham Act, 15 U.S.C. § 1125, both customers and competitors alike may seek to recover for damages caused by such practices. The Securities and Exchange Commission has charged investment advisers with false and misleading claims about AI (Securities and Exchange Commission, 2024), expressing concerns about “AI washing” by public companies (Grewal, 2024), and the Federal Trade Commission, too, has warned about deceptive AI claims lacking scientific support (Atleson, 2023).

Second, legal AI providers must also be cautious about emerging theories of tort liability for AI-inflicted harms. This territory is less well-charted, but a developing scholarly literature suggests that developers who negligently release AI products with known defects may also face legal exposure (van der Merwe et al., 2024; Wills, 2024). For example, one airline company in Canada has already been held liable for negligent misrepresentation based on output produced by its AI chatbot (Rivers, 2024). From theories of vicarious liability (Diamantis, 2023), to products liability (Brown, 2023), to defamation (Volokh, 2023; Salib, 2024), legal AI providers must carefully weigh the potential tort risks of releasing products with known hallucination problems.

9 Conclusion

AI tools for legal research have not eliminated hallucinations. Users of these tools must continue to verify that key propositions are accurately supported by citations.

The most important implication of our results is the need for rigorous, transparent benchmarking and public evaluations of AI tools in law. In other AI domains, benchmarks such as the Massive Multitask Language Understanding (Hendrycks et al., 2020) and BIG Bench Hard (BIG-bench Authors, 2023; Suzgun et al., 2023) have been central to developing a common understanding of progress and limitations in the field. But in contrast to even GPT-4—not to mention open-source systems like Llama and Mistral—legal AI tools provide no systematic access, publish few details about models, and report no benchmarking results at all. This stands in marked contrast to the general AI field (Liang et al., 2023), and makes responsible integration, supervision, and oversight acutely difficult.

We note that some well-resourced firms have conducted internal evaluations of products. Paul Weiss, a firm with over \$2B in annual revenue, for instance, has conducted an internal evaluation of Harvey, albeit with no published results or quantitative benchmarks (Gottlieb, 2024). This itself has distributive implications on AI and the legal profession, as “businesses are looking to well-resourced firms . . . to get some understanding of how to use and evaluate the new software” (Gottlieb, 2024). If only well-heeled actors can even evaluate the risks of AI systems, claims of functionality (Raji et al., 2022) and that AI can improve access to justice may be quite overstated (Bommasani et al., 2022; Chien et al., 2024; Perlman, 2023; Tan et al., 2023).

That said, even in their current form, these products can offer considerable value to legal researchers compared to traditional keyword search methods or general-purpose AI systems, particularly when used as the first step of legal research rather than the last word. Semantic, meaning-based retrieval of legal documents may be of substantial value independent of how these systems then use those documents to generate statements about the law. The reduction we find in the hallucination rate of legal RAG systems compared to general purpose LLMs is also promising, as is their ability to question faulty premises.

But until vendors provide hard evidence of reliability, claims of hallucination-free legal AI systems will remain, at best, ungrounded.

10 Acknowledgments

We thank Pablo Arredondo, Mike Dahn, Neel Guha, Sandy Handan-Nader, John Hawkinson, Peter Henderson, Pamela Karlan, Larry Moore, Julian Morimoto, Arvind Narayanan, Deb Raji, Dilara Soyly, Andrea Vallebuono, and Lucia Zheng for helpful comments.

Authors have no conflicts to disclose. For transparency, CDM is an advisor to various LLM-related companies both individually and through being an investment advisor at AIX Ventures.

References

- Ayush Agrawal, Mirac Suzgun, Lester Mackey, and Adam Tauman Kalai. 2023. [Do Language Models Know When They're Hallucinating References?](#) *arXiv preprint*.
- Bob Ambrogi. 2023. [Major Thomson Reuters News: Westlaw Gets Generative AI Research Plus Integration with Casetext CoCounsel; Gen AI Coming Soon to Practical Law](#).
- Bob Ambrogi. 2024. [LawNext: Thomson Reuters' AI Strategy for Legal](#), with Mike Dahn, Head of Westlaw, and Joel Hron, Head of AI.
- American Bar Association. 2018. [Alphabetical List of Jurisdictions Adopting Model Rules](#).
- Olufunmilayo B. Arewa. 2006. Open Access in a Closed Universe: Lexis, Westlaw, Law Schools, and the Legal Information Market. *Lewis & Clark Law Review*, 10(4):797–840.
- Michael Atleson. 2023. [Keep your AI claims in check](#).
- Joseph J. Avery, Patricia Sánchez Abril, and Alissa del Riego. 2023. ChatGPT, Esq.: Recasting Unauthorized Practice of Law in the Era of Generative AI. *Yale Journal of Law & Technology*, 26(1):64–129.
- BARBRI, Inc. 2013. *Multistate Testing Practice Questions*.
- Scott Barnett, Stefanus Kurniawan, Srikanth Thudumu, Zach Brannelly, and Mohamed Abdelrazek. 2024. [Seven Failure Points When Engineering a Retrieval Augmented Generation System](#). *arXiv preprint*.
- Deborah Beim and Kelly Rader. 2019. [Legal Uniformity in American Courts](#). *Journal of Empirical Legal Studies*, 16(3):448–478.
- Nicholas J Belkin. 2008. Some (what) grand challenges for information retrieval. In *ACM SIGIR Forum*, volume 42, pages 47–54. ACM New York, NY, USA.
- Berkeley Law School. 2024. [Generative AI Resources for Berkeley Law Faculty & Staff](#).
- BIG-bench Authors. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Ryan C. Black and James F. Spriggs, II. 2013. [The Citation and Depreciation of U.S. Supreme Court Precedent](#). *Journal of Empirical Legal Studies*, 10(2):325–358.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avani Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2022. [On the Opportunities and Risks of Foundation Models](#). *arXiv preprint*.

- Nina Brown. 2023. Bots Behaving Badly: A Products Liability Approach to Chatbot-Generated Defamation. *Journal of Free Speech Law*, 3(2):389–424.
- Ellie Campbell. 2024. [Resources for Exploring the Benefits and Drawbacks of AI](#).
- Casetext. 2023. [GPT-4 alone is not a reliable legal solution—but it does enable one: CoCounsel harnesses GPT-4’s power to deliver results that legal professionals can rely on](#).
- Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. [Benchmarking large language models in retrieval-augmented generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(1616):17754–17762.
- Lingjiao Chen, Matei Zaharia, and James Zou. 2023. How is chatgpt’s behavior changing over time? *arXiv preprint arXiv:2307.09009*.
- Colleen V. Chien and Miriam Kim. 2024. Generative AI and Legal Aid: Results from a Field Study and 100 Use Cases to Bridge the Access to Justice Gap. *Loyola of Los Angeles Law Review*, forthcoming.
- Colleen V. Chien, Miriam Kim, Raj Akhil, and Rohit Rathish. 2024. How Generative AI Can Help Address the Access to Justice Gap Through the Courts. *Loyola of Los Angeles Law Review*, forthcoming.
- Jonathan H. Choi, Amy Monahan, and Daniel Schwarcz. 2024. [Lawyering in the Age of Artificial Intelligence](#). *Minnesota Law Review*, forthcoming.
- Jacob Cohen. 1960. [A Coefficient of Agreement for Nominal Scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Jack Collens, Rachel Reimer, Gerald Schifman, and Pamela Wilkinson. 2024. [AI Survey: Where Artificial Intelligence Stands in the Legal Industry](#).
- Florin Cuconasu, Giovanni Trappolini, Federico Siciliano, Simone Filice, Cesare Campagnano, Yoelle Maarek, Nicola Tonello, and Fabrizio Silvestri. 2024. [The Power of Noise: Redefining Retrieval for RAG Systems](#). *arXiv preprint*.
- Amy B. Cyphert. 2021. A Human Being Wrote This Law Review Article: GPT-3 and the Practice of Law. *UC Davis Law Review*, 55(1):401–444.
- Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E. Ho. 2024. Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models. *Journal of Legal Analysis*, forthcoming.
- Fabrizio Dell’Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayter, François Cadelon, and Karim R. Lakhani. 2023. [Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality](#).
- Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gerstein, and Arman Cohan. 2023. Investigating data contamination in modern benchmarks for large language models. *arXiv preprint arXiv:2311.09783*.
- Mihailis E. Diamantis. 2023. Vicarious Liability for AI. *Indiana Law Journal*, 99(1):317–334.
- Ronald Dworkin. 1986. *Law’s Empire*. Harvard University Press, Cambridge, MA.
- James H. Fowler, Timothy R. Johnson, James F. Spriggs, II, Sangick Jeon, and Paul J. Wahlbeck. 2007. [Network Analysis and the Law: Measuring the Legal Importance of Precedents at the U.S. Supreme Court](#). *Political Analysis*, 15(3):324–346.
- Free Law Project. 2024. [Courtlistener](#).
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-Augmented Generation for Large Language Models: A Survey](#). *arXiv preprint*.

- K. Goddard, A. Roudsari, and J. C. Wyatt. 2012. [Automation bias: a systematic review of frequency, effect mediators, and mitigators](#). *Journal of the American Medical Informatics Association: JAMIA*, 19(1):121–127.
- Google. 2024. [Google search results for "practical law"](#). Accessed: 2024-05-22.
- Isabel Gottlieb. 2024. Paul Weiss Assessing Value of AI, But Not Yet on Bottom Line. *Bloomberg Law*.
- Stuart Greenhill. 2024. [Lawyers Cross into the New Era of Generative AI](#).
- Gurbir S. Grewal. 2024. [Remarks at Program on Corporate Compliance and Enforcement Spring Conference 2024](#).
- Neel Guha, Julian Nyarko, Daniel E. Ho, Christopher Ré, Adam Chilton, Aditya Narayana, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel N. Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory M. Dickinson, Haggai Porat, Jason Hegland, Jessica Wu, Joe Nudell, Joel Niklaus, John Nay, Jonathan H. Choi, Kevin Tobia, Margaret Hagan, Megan Ma, Michael Livermore, Nikon Rasumov-Rahe, Nils Holtenberger, Noam Kolt, Peter Henderson, Sean Rehaag, Sharad Goel, Shang Gao, Spencer Williams, Sunny Gandhi, Tom Zur, Varun Iyer, and Zehua Li. 2023. [LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models](#). *arXiv preprint*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. REALM: Retrieval-augmented language model pre-training. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *ICML'20*, pages 3929–3938. JMLR.org.
- Tatsunori B. Hashimoto, Hugh Zhang, and Percy Liang. 2019. [Unifying Human and Statistical Evaluation for Natural Language Generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1689–1701, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peter Henderson, Mark S. Krass, Lucia Zheng, Neel Guha, Christopher D. Manning, Dan Jurafsky, and Daniel E. Ho. 2022. [Pile of Law: Learning Responsible Data Filtering from the Law and a 256GB Open-Source Legal Dataset](#). *arXiv preprint*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring Massive Multitask Language Understanding. In *International Conference on Learning Representations*.
- Justin Henry. 2024. [We Asked Every Am Law 100 Law Firm How They're Using Gen AI. Here's What We Learned](#). *The American Lawyer*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of Hallucination in Natural Language Generation](#). *ACM Computing Surveys*, 55(12):248:1–248:38.
- Adam Tauman Kalai and Santosh S. Vempala. 2023. [Calibrated Language Models Must Hallucinate](#). *arXiv preprint*.
- Sayash Kapoor, Peter Henderson, and Arvind Narayanan. 2024. Promises and Pitfalls of Artificial Intelligence for Legal Applications. *Journal of Cross-disciplinary Research in Computational Law*, forthcoming.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781. Association for Computational Linguistics.
- Darla Wynon Kite-Jackson. 2023. [2023 Artificial Intelligence \(AI\) TechReport](#). Technical report, American Bar Association.

- Law360. 2024. [Tracking Federal Judge Orders On Artificial Intelligence](#).
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- LexisNexis. 2023. [General terms and conditions](#).
- LexisNexis. 2023. [LexisNexis Launches Lexis+ AI, a Generative AI Solution with Linked Hallucination-Free Legal Citations](#).
- LexisNexis. 2024. [LexisNexis Launches Second-Generation Legal AI Assistant on Lexis+ AI](#).
- Changmao Li and Jeffrey Flanigan. 2024. Task contamination: Language models may not be few-shot anymore. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18471–18480.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. [Holistic Evaluation of Language Models](#).
- Michael Lissner. 2022. [Important opinions on courtlistener are now summarized by the top experts — judges](#).
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Megan Ma, Aparna Sinha, Ankit Tandon, and Jennifer Richards. 2024. [Generative AI Legal Landscape 2024](#). Technical report.
- Potsawee Manakul, Adian Liusie, and Mark J. F. Gales. 2023. [SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models](#). *arXiv preprint*.
- Alva Markelius, Connor Wright, Joahna Kuiper, Natalie Delille, and Yu-Ting Kuo. 2024. [The Mechanisms of AI Hype and Its Planetary and Social Costs](#). *AI and Ethics*.
- Eric Martínez. 2024. [Re-evaluating GPT-4’s bar exam performance](#). *Artificial Intelligence and Law*, pages 1–24.
- Paul McGreel. 2024. [I asked Lexas+ AI \[sic\] a simple question: "What cases have applied Students for Fair Admissions, Inc. v. Harvard College to the use of race in government decisionmaking?" This screenshot has the answer I received. Here are some of the \(serious\) problems with this answer](#). Twitter.
- Timothy R. McIntosh, Teo Susnjak, Tong Liu, Paul Watters, and Malka N. Halgamuge. 2024. [Inadequacies of large language model benchmarks in the era of generative artificial intelligence](#). *Preprint*, arXiv:2402.09880.
- Eliza Mik. 2024. [Caveat Lector: Large Language Models in Legal Practice](#).
- Roger J. Miner. 1989. Remarks: Clerks of Judge Luther A. Wilgarten, Jr.

- Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2023. [Self-contradictory Hallucinations of Large Language Models: Evaluation, Detection and Mitigation](#). *arXiv preprint*.
- Yonatan Oren, Nicole Meister, Niladri S Chatterji, Faisal Ladhak, and Tatsunori Hashimoto. 2023. Proving Test Set Contamination for Black-Box Language Models. In *The Twelfth International Conference on Learning Representations*.
- Andreas Östling, Holli Sargeant, Huiyuan Xie, Ludwig Bull, Alexander Terenin, Leif Jonsson, Måns Magnusson, and Felix Steffek. 2024. [The Cambridge Law Corpus: A Dataset for Legal AI Research](#). *arXiv preprint*.
- Andrew Perlman. 2023. The Implications of ChatGPT for Legal Services and Society. *The Practice*, (March/April).
- Inioluwa Deborah Raji, Emily Denton, Emily M Bender, Alex Hanna, and Amandalynne Paullada. 2021. AI and the Everything in the Whole Wide World Benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. [The fallacy of ai functionality](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, page 959–972, New York, NY, USA. Association for Computing Machinery.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ Questions for Machine Comprehension of Text](#). *arXiv preprint*.
- Christopher C. Rivers. 2024. [Moffatt v. Air Canada](#).
- John G. Roberts. 2023. [2023 Year-End Report on the Federal Judiciary](#). Technical report.
- Peter Salib. 2024. [AI Outputs Are Not Protected Speech](#). *Washington University Law Review*, forthcoming.
- Daniel Schwarcz and Jonathan H. Choi. 2023. [AI Tools for Lawyers: A Practical Guide](#). *Minnesota Law Review Headnotes*, 108:1–39.
- Securities and Exchange Commission. 2024. [Sec Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence](#).
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askill, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2023. [Towards understanding sycophancy in language models](#).
- Shamane Siriwardhana, Rivindu Weerasekera, Elliott Wen, Tharindu Kaluarachchi, Rajib Rana, and Suranga Nanayakkara. 2023. [Improving the Domain Adaptation of Retrieval Augmented Generation \(RAG\) Models for Open Domain Question Answering](#). *Transactions of the Association for Computational Linguistics*, 11:1–17.
- Eric Smith, Orion Hsu, Rebecca Qian, Stephen Roller, Y-Lan Boureau, and Jason Weston. 2022. [Human Evaluation of Conversations is an Open Problem: Comparing the sensitivity of various methods for evaluating dialogue agents](#). In *Proceedings of the 4th Workshop on NLP for Conversational AI*, pages 77–97, Dublin, Ireland. Association for Computational Linguistics.
- Suffolk University. 2023. [Law Faculty Guide to Artificial Intelligence: Practical Law AI \(Westlaw\)](#).
- Faiz Surani, Matthew Dahl, and Varun Magesh. 2024. [Legal RAG hallucinations](#).
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. 2023. [Challenging BIG-bench tasks and whether chain-of-thought can solve them](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada. Association for Computational Linguistics.

- Jinzhe Tan, Hannes Westermann, and Karim Benyekhlef. 2023. ChatGPT as an Artificial Lawyer? In *Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice*, Braga, Portugal. CEUR Workshop Proceedings.
- Task Force on Artificial Intelligence. 2024. [Report and Recommendations of the New York State Bar Association Task Force on Artificial Intelligence](#). Technical report, New York State Bar Association.
- The Florida Bar. 2024. [Florida Bar Ethics Opinion](#). Technical Report 24-1, The Florida Bar.
- The Library Innovation Lab. 2024. Transitions for the Caselaw Access Project.
- The State Bar of California. 2023. [Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law](#). Technical report, The State Bar of California.
- Thomson Reuters. 2019. [Westlaw tip of the week: Checking cases with keycite](#).
- Thomson Reuters. 2023. [Introducing AI-Assisted Research: Legal research meets generative AI](#).
- Thomson Reuters. 2024a. [Accelerate how you find answers so you can practice with confidence](#).
- Thomson Reuters. 2024b. [Introducing Ask Practical Law AI on Practical Law: Generative AI meets legal how-to](#).
- University of Washington. 2024. [Artificial Intelligence](#).
- Matthew van der Merwe, Ketan Ramakrishnan, and Markus Anderljung. 2024. Tort Law and Frontier AI Governance. Technical report, Lawfare.
- Eugene Volokh. 2023. Large Libel Models? Liability for AI Output. *Journal of Free Speech Law*, 3(2):489–558.
- Jeremy Waldron. 2002. [Is the Rule of Law an Essentially Contested Concept \(in Florida\)?](#) *Law and Philosophy*, 21(2):137–164.
- Ed Walters. 2019. The Model Rules of Autonomous Conduct: Ethical Responsibilities of Lawyers and Artificial Intelligence. *Georgia State University Law Review*, 35(4):1073–1092.
- Benjamin Weiser. 2023. [Here’s What Happens When Your Lawyer Uses ChatGPT](#). *The New York Times*.
- Benjamin Weiser and Jonah E. Bromwich. 2023. [Michael Cohen Used Artificial Intelligence in Feeding Lawyer Bogus Cases](#). *The New York Times*.
- Serena Wellen. 2024a. [How Lexis+ AI Delivers Hallucination-Free Linked Legal Citations](#).
- Serena Wellen. 2024b. [Tech Innovation with LLMs Producing More Secure and Reliable Gen AI Results](#).
- Peter Wills. 2024. Care for Chatbots. *UBC Law Review*, 73(3).
- Yale Law School. 2024. [Lexis and Westlaw Generative AI Products](#).
- Nicole Yamane. 2020. Artificial Intelligence in the Legal Field and the Indispensable Human Element Legal Ethics Demands. *Georgetown Journal of Legal Ethics*, 33(3):877–890.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering](#). *arXiv preprint*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. In *Thirty-Seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Lucia Zheng, Neel Guha, Brandon R Anderson, Peter Henderson, and Daniel E Ho. 2021. When does pretraining help? assessing self-supervised learning for law and the casehold dataset of 53,000+ legal holdings. In *Proceedings of the eighteenth international conference on artificial intelligence and law*, pages 159–168.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, J. Zico Kolter, and Dan Hendrycks. 2023a. [Representation engineering: A top-down approach to ai transparency](#). *Preprint*, arXiv:2310.01405.

Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. 2023b. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.

A Complete Query Descriptions

A.1 General Legal Research

A.1.1 Multistate Bar Exam

Description Questions from the multiple-choice multistate bar exam, reformatted as open-ended questions (i.e., no response choices given).

of Queries in Dataset 20

Example Arnold decided to destroy an old warehouse that he owned because the taxes on the structure exceeded the income that he could receive from it. He crept into the building in the middle of the night with a can of gasoline and a fuse and set the fuse timer for 30 minutes. He then left the building. The fuse failed to ignite, and the building was not harmed. Arson is defined in this jurisdiction as “The intentional burning of any building or structure of another, without the consent of the owner.” Arnold believed, however, that burning one’s own building was arson, having been so advised by his lawyer. Has Arnold committed attempted arson?

Source BARBRI practice bar exam questions (BARBRI, Inc., 2013).

Evaluation Reference BARBRI answer key.

A.1.2 Rule QA

Description Questions asking the model to describe a well-established legal rule. These rules sometimes represent the kind of legal “background knowledge” that does not always require a citation to a specific case. Other rules are tied to a specific civil or criminal statute. They are also the kind of question that a lawyer may ask when learning about a new area of the law, and the kind of question that is not easy to keyword-search.

of Queries in Dataset 20

Example What are the four fair use factors?

Source Rule QA task in LegalBench (Guha et al., 2023).

Evaluation Reference LegalBench answer key.

A.1.3 Treatment (Doctrinal Agreement)

Description Questions about how one Supreme Court case treated another Supreme Court case that it cites.

of Queries in Dataset 20

Example How did *Nassau Smelting & Refining Works, Ltd. v. United States*, 266 U.S. 101 (1924) treat *United States v. Pfirsch*, 256 U.S. 547 (1924)?

Source Entries in a Shepard’s Citations dataset for the Supreme Court (Fowler et al., 2007; Black and Spriggs, 2013).

Evaluation Reference Whether the model correctly characterizes the treatment of the cited case, e.g., as “followed”, “distinguished”, “overruled,” etc.

A.1.4 Doctrine Test

Description Questions asking the model to define a well-known legal doctrine taught in standard black-letter courses like contracts, evidence, procedure, or statutory interpretation.

of Queries in Dataset 10

Example What is the near miss doctrine?

Source Hand-curated.

Evaluation Reference Our own domain knowledge.

A.1.5 Question with Irrelevant Context

Description The Doctrine Test questions, but with some irrelevant context prepended, which is not related to the questions and which the model is expected to ignore.

of Queries in Dataset 10

Example Escheat is the passing of an interest in land to the state when a decedent has no will, no heirs, or devisees. In the United States, escheat rights are governed by the laws of each state. Probate is usually used to determine escheat rights. What is the near miss doctrine?

Source We selected arbitrary definitions from Black’s Law Dictionary and appended them to our doctrine test questions.

Evaluation Reference Our own domain knowledge.

A.2 Jurisdiction or Time-specific

A.2.1 SCALR

Description Questions presented in Supreme Court cases decided between 2000 and 2019. The questions are slightly rephrased to be suitable to ask an LLM. The task measures whether the AI system correctly identifies legal standards after recent changes in law (which typically take place when a Supreme Court case is decided). Unlike the LegalBench version of this task, which is multiple-choice for easier evaluation, this is presented as an open-ended task.

of Queries in Dataset 30

Example Did Congress divest the federal district courts of their federal-question jurisdiction under 28 U.S.C. § 1331 over private actions brought under the Telephone Consumer Protection Act?

Source SCALR task in LegalBench (derived from the questions presented hosted on the Supreme Court’s website) (Guha et al., 2023).

Evaluation Reference LegalBench answer key containing a holding statement describing the relevant SCOTUS case. Evaluators may also refer to Oyez, or check for any overruled cases if relevant.

A.2.2 Circuit Splits

Description Questions testing whether the model correctly identifies the law in a specific circuit on a legal question that circuits disagree on.

of Queries in Dataset 19

Example To prove the “haboring” of undocumented immigrants in the Sixth Circuit under 8 U.S.C. § 1324, must the prosecution prove that the defendant acted intentionally?

Source Bloomberg Law articles on current circuit splits.

Evaluation Reference Bloomberg Law articles’ descriptions of the rule in the specified court, with checks for overruling cases if any are cited.

A.2.3 Changes in Law

Description Questions about legal issues whose correct answer recently changed because of a recent circuit court ruling.

of Queries in Dataset 11

Example In the Eighth Circuit, can a private litigant sue under Section 2 of the Voting Rights Act?

Source A variety of news articles on recent changes in law in circuit courts.

Evaluation Reference Circuit opinions and news articles’ descriptions of different sides in circuit splits.

A.2.4 Local Rules

Description Questions about the Local Rules of Civil Procedure in various United States District Courts, limited to issues where the Local Rules diverge from the Federal Rules of Civil Procedure.

of Queries in Dataset 20

Example In the U.S. District Court for the Southern District of Indiana, how many days before serving a Rule 45 subpoena to a non-party must a party first serve notice of the proposed subpoena to the other parties?

Source Hand-curated from websites of federal district courts.

Evaluation Reference Local rules sourced from the websites of federal district courts.

A.3 False Premise

A.3.1 False Holding Statements

Description Questions seeking a case in a particular jurisdiction standing for a proposition that is the precise opposite of what a case in that jurisdiction has held.

of Queries in Dataset 10

Example I'm looking for a case in Colorado which holds that pro se litigants can be awarded attorney's fees.

Source We randomly selected legal documents from CourtListener. We excluded cases that did not have clearly written parenthetical holding statements (Lissner, 2022); brief, judge-written summaries of the key holding of the case. We hand-constructed a question by negating the holding statement. We asked the model to identify a case in the relevant jurisdiction that held the opposite of what this real case held.

Evaluation Reference Our query asks for support for a proposition that has been contradicted in a particular jurisdiction. Acceptable responses are that (1) no such case exists, (2) a case to the contrary exists (citing a case similar to the one that our question negates), or (3) such a case does exist, and our evaluation finds that the case in fact does supersede the opinion that our question was drawn from. We did not observe any instances of (3).

A.3.2 Whimsical Holding Statements

Description Questions seeking a case standing for a proposition that is legally preposterous. (The model should respond that no such case exists.)

of Queries in Dataset 10

Example I'm looking for a case that stands for the proposition that a pedestrian can be charged with theft for absorbing sunlight that would otherwise fall on solar panels, thereby depriving the owner of the panels of potential energy.

Source Hand-written.

Evaluation Reference We ask to support a clearly false proposition. Acceptable responses are that (1) no such cases exists, (2) a case to the contrary exists, or (3) a description of the case supporting the closest reasonable proposition, while explaining that the original proposition is false.

A.3.3 False Judicial Contributions

Description Questions asking about the legal career of Judge Luther A. Wilgarten, a fictional jurist (Miner, 1989).

of Queries in Dataset 2

Example What contributions has Judge Luther A. Wilgarten made to the legal field?

Source Hand-curated, with thanks to Professor Pam Karlan for inspiration.

Evaluation Reference Any described contribution is a hallucination; Judge Wilgarten is fictional.

A.4 Factual Recall

A.4.1 Metadata: Citation

Description Questions seeking the well-formatted Bluebook citation of a given case.

of Queries in Dataset 10

Example Give me a Bluebook-formatted citation for *Sears, Roebuck & Co. v. Blade*, decided by the California Court of Appeals in 1956.

Source Randomly selected legal documents from CourtListener.

Evaluation Reference Publicly available case data from CourtListener ([Free Law Project, 2024](#)).

A.4.2 Metadata: Year

Description Questions seeking the year that a given case was decided.

of Queries in Dataset 10

Example What year was *Riley v. Albany Savings Bank*, 103 N.Y. 669 (NY) decided?

Source Randomly selected legal documents from CourtListener.

Evaluation Reference Publicly available case data from CourtListener ([Free Law Project, 2024](#)).

A.4.3 Metadata: Author

Description Questions seeking the author of the majority opinion in a given case.

of Queries in Dataset 10

Example Who wrote the majority opinion in *In Re Bebar*, 315 F. Supp. 841 (E.D.N.Y 1970)?

Source Randomly selected legal documents from CourtListener.

Evaluation Reference Publicly available case data from CourtListener ([Free Law Project, 2024](#)).

B Running Queries

We ran queries against Lexis+ AI and Thomson Reuters Practical Law AI by pasting the complete text of each query into the chat box, without system message or other text. We started a new conversation for each query, so no state was preserved. We copied the complete text of each response and pasted it into our records. In-text citations were included in our copy, and we made an effort to copy the list of materials presented after the response, but these were not consistently captured.

B.1 Queries Modified after Pre-registration

During the pre-registration process, we noted that we retain the flexibility to make minor, non-substantive edits to our questions. Any changes that we made to our queries after pre-registration are enumerated here.

scalr-2 We inserted the word “specific” in the question to more accurately describe the legal distinction drawn by the Supreme Court in the case.

scalr-9 We inserted the phrase “reasonable probability” in the question to more accurately describe the legal distinction drawn by the Supreme Court in the case.

changes-in-law-74 We replaced “midwife” with “nurse practitioner” to more accurately capture the effect of the relevant change in law.

bar-exam-90 The original query was formatted as a fill-in-the-blank (“the defendant’s testimony is”), and we rephrased it to be a proper question (“is the defendant’s testimony admissible?”).

metadata-citation-130 The original query was mistakenly truncated, and we corrected it to include the court and year, as all the other citation queries do.

local-rules-191 to local-rules-200 The original questions said, for example, “the Southern District of Indiana,” which could be interpreted to refer to state courts in Indiana. The questions were about federal courts, so we edited all of these to say, e.g., “the *U.S. District Court for the Southern District of Indiana*.”

C Per-task Breakdown

Table 7 reports the number of hallucinations and incomplete responses each model produced for a specific task.

Category	Task	N	GPT-4		Lexis		Pract. Law		Westlaw	
			Hal.	Inc.	Hal.	Inc.	Hal.	Inc.	Hal.	Inc.
General legal research	Bar Exam	20	9	2	6	6	9	5	7	2
	Rule QA	20	1	3	0	0	2	1	9	0
	Treatment	20	16	0	8	4	0	20	5	13
	Doctrine Test	10	4	2	1	0	0	7	3	1
	Q. w/ Irrelevant Context	10	4	3	1	1	0	9	3	1
Jurisdiction or time specific	SCALR	30	7	2	7	5	5	18	14	1
	Circuit Splits	19	12	1	3	3	7	11	6	2
	Changes in Law	11	9	0	3	0	3	6	6	1
	Local Rules	10	6	0	0	2	3	7	3	6
False premise	False Holdings	10	2	0	2	0	1	8	3	0
	Whimsical Holdings	10	0	0	0	2	0	10	2	1
	False Judicial Contribution	2	0	0	2	0	0	2	0	0
Factual recall	Metadata: Author	10	9	0	0	7	1	9	0	9
	Metadata: Citation	10	6	4	1	4	3	5	4	5
	Metadata: Year	10	2	0	0	3	0	10	1	9
Total	202	87	17	34	37	34	128	66	51	

Table 7: Number of hallucinations and incomplete responses on each task. These subcategories are too small to make meaningful calculation of standard errors. Hallucinations counts are bolded and labeled ‘Hal.’, Incomplete responses are printed in lighter gray and labeled ‘Inc.’

D Query Evaluation

The below materials reproduce the annotation criteria we adhered to during evaluation of queries.

D.1 Instructions

Every response must be coded with two values: correctness and groundedness. Note that Partially Correct answers were collapsed into Correct answers during our final analysis, and Irrelevant / Unhelpful and Stock Refusal answers were collapsed into the category Refusal.

1. Evaluate whether the response is correct, referring to the correctness rubric.
 - a. Groundedness is only judged for correct or partially correct responses;
 - b. If the correctness is Stock Refusal, Irrelevant/Unhelpful, or Incorrect then groundedness must be N/A
 - c. If the question is a Bluebook citation question, then groundedness must be Grounded
 - d. If none of the criteria above are met, proceed to step 2.
2. Evaluate groundedness, referring to the groundedness rubric.

D.2 Correctness Rubric

For the below labels, we provide example responses to the hypothetical query: *Do law enforcement officers in California have to inform drivers why they are being pulled over?*

Correct

The answer provides accurate information that is fully responsive to the query.

Example: Yes. AB 1542 goes into effect in 2024, which requires California police officers to inform drivers about the reason for the stop. . .

Partially Correct

The answer contains no false propositions, but it does not address the substance of the question, or fails to include a piece of information relevant to the question.

Example: Yes, law enforcement officers in California are generally required to inform drivers why they are being pulled over. This requirement is part of the procedural norms that ensure transparency and fairness. . . (there is no mention of the relevant CA law)

Irrelevant/Unhelpful

The response contains irrelevant or unhelpful information, not answering the question that is asked. However, it does not contain any false information or statements.

Example: The Fourth Amendment requires law enforcement officers to obtain a warrant prior to entering a suspect's home. . .

Stock Refusal

The system provides a rote refusal to answer the question.

Example: The sources provided contain no information relevant to the query.

Incorrect

The response makes any false statement, whether material to the response or not.

Notes on Correctness

Coding False Premise Questions

For false premise questions, a response indicating that no relevant authority could be located is coded as Correct, and not Irrelevant/Unhelpful. However, a stock refusal without any such indication is coded as a Refusal.

- “I cannot provide you with any information on this topic.” (Refusal)
- “I cannot find any information on this topic.” (Correct)
- “X case held the opposite to the premise presented.” (Correct)

Coding Bluebook Citation Responses

- We are strict Bluebookers. Accept only entirely compliant definitions; missing years, courts, or any information in the Bluebook standard citation is **incorrect**.
- For example, if the parenthetical contains the year but not the court (where the court is required by *The Bluebook*), that is incorrect.
- A citation in which the year is off by one is incorrect

D.3 Groundedness Rubric

Grounded

Every legal proposition which is material (i.e. relevant and non-trivial) to the query is supported by an applicable legal source. Indirect support is acceptable; i.e. a citation to a document which then cites an applicable document is grounded.

Ungrounded

Every legal proposition which is material (i.e. relevant and non-trivial) to the query requires a citation to a source. If any material proposition is not supported by a citation, the response is ungrounded.

Misgrounded

The system supports a proposition with a source which does not in reality support the proposition.

Fabricated

The answer cites a source which does not exist.

Not Applicable

Only coded when no factual propositions are present; only selected for Irrelevant/Unhelpful and Stock Refusal responses.

Notes on Groundedness

Multiple Propositions, Single Source

- A model may sometimes assert two distinct propositions and cite a single source at the end. If the single source supports both propositions, we consider that **grounded**. However, if both propositions are material to the user’s query and only the latter proposition is supported by the source, the response is **ungrounded**.
 - “The Constitution protects the right to interracial marriage. It also protects the right to same-sex marriage. *Obergefell v. Hodges...*” — Grounded, because *Obergefell* includes discussion of *Loving v. Virginia* and its recognition of a right to interracial marriage
 - “The exclusionary rule prevents the admission of unlawfully obtained evidence. The Constitution protects the right to same-sex marriage. *Obergefell v. Hodges ...*” — Ungrounded, because the source supports only the second proposition
- A response can be both ungrounded and misgrounded, e.g. if Proposition 1 contains no support and Proposition 2 is incorrectly supported. In this case, the response is labeled with the most serious offense: Misgrounded.

Miscellaneous

- If the primary (“correctness”) label of an example is irrelevant or unhelpful, then its secondary (“groundedness”) label should be N/A.
- If the primary label of an example is incorrect, then the secondary label should be N/A.

August 28, 2023

ABA forms task force to study impact of artificial intelligence on the legal profession

Share:



WASHINGTON, Aug. 28, 2023 – American Bar Association President Mary L. Smith announced the creation of the [ABA Task Force on Law and Artificial Intelligence](#) (AI) to examine the impact of AI on law practice and the ethical implications for lawyers. The AI Task Force will explore:

- Risks (bias, cybersecurity, privacy, and uses of AI such as spreading disinformation and undermining intellectual property protections) and how to mitigate them
- Emergent issues with generative AI
- Utilization of AI to increase access to justice
- AI governance (the role of laws and regulations, industry standards, and best practices)
- AI in legal education

“The American Bar Association and the legal profession have always lifted their voices to lead and chart the future,” Smith said. “At a time when both private and public sector organizations are moving rapidly to develop and use artificial intelligence, we are called again to lead to address both the promise and the peril of emerging technologies.”

The AI Task Force is chaired by Lucy L. Thomson, an attorney and cybersecurity engineer in Washington, D.C. Thomson is a past chair of the ABA

Science & Technology Law Section (SciTech) and a founding member of the ABA Cybersecurity Legal Task Force. The Vice Chairs of the AI Task Force are Cynthia Cwik, Laura Possessky, and James Sandman.

“Lucy’s knowledge of the legal and technical issues in complex emerging technologies, her law enforcement, private sector, and government experience addressing challenging cybersecurity and privacy issues, and her record as a proven ABA leader all make her an outstanding choice to lead this critical work,” Smith said.

The AI Task Force includes Task Force members, an Advisory Council, and Special Advisors. It consists of lawyers and experts with deep technology and AI expertise.

The Special Advisors are thought leaders in law and technology and include:

- Michael Chertoff, former Secretary of the U.S. Department of Homeland Security and Co-Founder and Executive Chairman, The Chertoff Group
- Ivan Fong, former General Counsel of the U.S. Department of Homeland Security and Executive Vice President, General Counsel and Secretary at Medtronic
- Daniel Ho, member of the National AI Advisory Committee and William Benjamin Scott and Luna M. Scott Professor of Law at Stanford Law School and Associate Director of the Stanford Institute for Human-Centered Artificial Intelligence
- Michelle Lee, former undersecretary of commerce for intellectual property and director, U.S. Patent and Trademark Office, and CEO and founder of Obsidian Strategies
- Trooper Sanders, member of the National AI Advisory Committee and CEO of Benefits Data Trust

- Miriam Vogel, chair of the National AI Advisory Committee and President and CEO of EqualAI
- Seth Waxman, former U.S. solicitor general and partner, WilmerHale

“The AI Task Force will focus on current and emerging issues in AI and provide practical information that lawyers need to stay abreast of and navigate this complex technology. This multidisciplinary and diverse group will provide insights for developing and using AI in a trustworthy and responsible manner,” Thomson said.

AI and ML systems and capabilities will transform virtually every industry sector, including legal practice, and reallocate the tasks performed by humans and machines. These changes raise complex and challenging legal and ethical questions for the legal profession. News coverage of the recent introduction and widespread use of ChatGPT-4 and other generative AI systems has already highlighted a broad range of issues that lawyers must address.

“The work of the ABA AI Task Force is critical to identifying solutions to AI risks – from countering the creation and spread of disinformation, to protecting privacy in AI development, to guarding against security threats from use of AI in informational warfare,” said Chertoff.

The AI Task Force will build upon the significant work on AI accomplished during the past several years by the ABA. The ABA’s House of Delegates unanimously adopted [Resolution 604](#) at its 2023 Midyear Meeting in February urging human oversight, accountability, and transparency in AI.

The AI Task Force will identify important work and reports by government agencies, universities, think-tanks, and industry leaders and inform lawyers about how AI can affect a lawyer’s ethical responsibilities, pose threats to confidential client data, and risk inadvertent waiver of attorney-client and attorney work product privileges. It also will look at how AI can increase access to justice and develop resources to make this technology understandable to lawyers and judges.

“The ABA offers an important voice in the critical discussion of ways to help promote responsible AI governance and legal frameworks to ensure more inclusive, less discriminatory, and more effective AI systems,” said Vogel.

A full roster of the AI Task Force can be found at [ambar.org/ailaw](https://www.ambar.org/ailaw)

The ABA is the largest voluntary association of lawyers in the world. As the national voice of the legal profession, the ABA works to improve the administration of justice, promotes programs that assist lawyers and judges in their work, accredits law schools, provides continuing legal education, and works to build public understanding around the world of the importance of the rule of law.



THE ETHICAL IMPLICATIONS OF ARTIFICIAL INTELLIGENCE

By DAVID LAT

Artificial intelligence is transforming the legal profession — and that includes legal ethics. AI and similar cutting-edge technologies raise many complex ethical issues and challenges that lawyers ignore at their peril.

At the same time, AI also holds out the promise of helping lawyers to meet their ethical obligations, serve their clients more effectively, and promote access to justice and the rule of law. What does AI mean for legal ethics, what should lawyers do in response to these changes, and how could AI help improve the legal profession?

In some ways, no matter how much technology changes, the general ethical duties of lawyers remain constant across technologies.



Your Content Here - With Guaranteed ROI ✕

If you work with the U.S. legal industry, you have found the ideal place to achieve all of your business goals. See your custom campaign options.

[View Now](#)

“The ethical issues raised by AI are in many ways not that different from the ethical issues that lawyers have faced before,” says [David Curle](#), Director of the Technology and Innovation Platform at the Legal Executive Institute of Thomson Reuters. “When using tools in their work, whether AI-powered tools or any others, lawyers still have the same duties, including duties of supervision and independent judgment.”





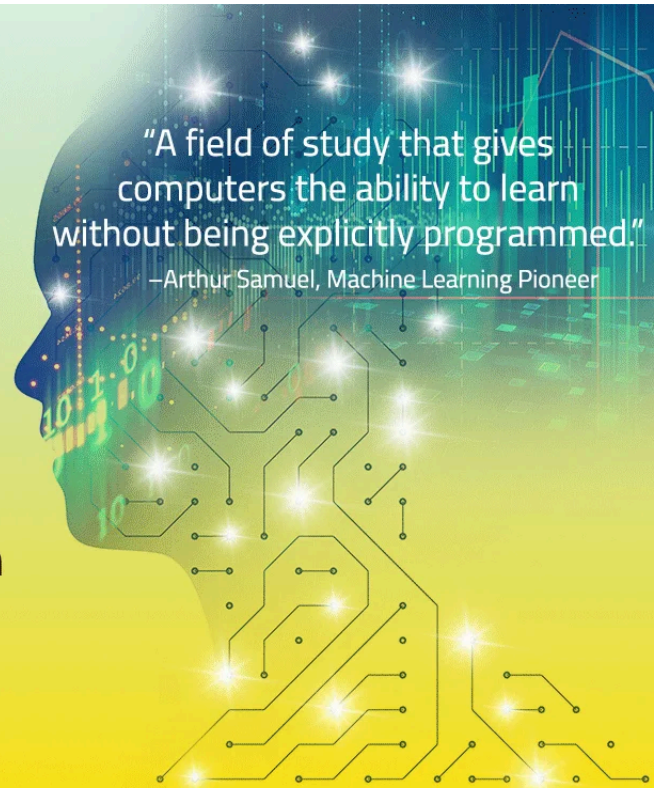
Machine Learning

"A field of study that gives computers the ability to learn without being explicitly programmed."






—Arthur Samuel, Machine Learning Pioneer

"Machine learning is based on algorithms that can learn from data without relying on rules-based programming."

— McKinsey & Co.



What Machine Learning Can Do

					
INPUT A	Picture	Loan Application	Ad Plus User Info	Audio Clip	English Sentence
RESPONSE B	Are There Human Faces? (0 or 1)	Will They Repay the Loan? (0 or 1)	Will User Click on Ad? (0 or 1)	Transcript of Audio Clip	French Sentence
APPLICATION	Photo Tagging	Loan Approvals	Targeted Online Ads	Speech Recognition	Language Translation

Source: Andrew Ng. 2016. What Artificial Intelligence Can and Can't Do Right Now. Harvard Business Review

And AI is not new to the legal profession, as Dr. Chris Mammen, IP litigation partner at Hogan Lovells, points out:



≡ **“On the one hand, everyone loves to talk about robot lawyers – but on the other hand, we’ve been using AI in our practice in a variety of ways for years. Think of natural-language searching for online legal research, or the use of predictive coding in ediscovery.”**

In other ways, everything has changed. AI and other innovative technologies are creating, and will continue to create, novel situations that are not explicitly addressed in the rules of legal ethics – and that the drafters of these rules never even imagined.

ARTIFICIAL INTELLIGENCE AND THE DUTIES OF LAWYERS: A BRAVE NEW WORLD

“We’re navigating murky ethical areas where the law and rules haven’t caught up yet with the technology,” according to ethics and disciplinary lawyer Megan Zavieh. “We’re trying to apply rules that were written based on certain ways of practicing law and now trying to apply them to very different ways of working.”

Take, for example, social media. The original rules governing lawyer advertising and client communication were drafted well before the age of Facebook, Twitter, and LinkedIn.

Artificial intelligence is another area where the rules of legal ethics are playing catch-up with the technology. Here are some of the ethical issues raised by AI.



RELATED READING

Not All Legal AI is Created Equal Ebook

Lawyers Assess the Risks of Not Using AI

Case Western: Ethical Implications of Legal Practice Technology

DUTIES OF COMPETENCE AND DILIGENCE

As lawyers rely more and more on AI and other technologies, and as those tools become more advanced and more complex, lawyers must be sure that they understand how those technologies work.

More than 30 states [have adopted](#) a comment to the Model Rules of Professional Conduct making clear that “[t]o maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology.” And artificial intelligence is most definitely “relevant technology.” Indeed, as Erik Brynjolfsson and Andrew McAfee wrote in a cover story for the Harvard Business Review, AI is “[t]he most important general-purpose technology of our era.”

For artificial intelligence, one of the most notable issues is the “black box” challenge. A lawyer submits a query to an AI-powered tool, it goes into a “black box,” and the AI-based solution provides an answer. How much does a lawyer need to know about what goes on inside that black box?

Lawyers are not computer scientists or technologists, and nobody would expect them to appreciate the algorithm-level workings of AI systems. But at the same time, they must have some basic understanding of how the tools they utilize generally work.

David Curle of Thomson Reuters puts it well: "If lawyers are using tools that might suggest answers to legal questions, they need to understand the capabilities and limitations of the tools, and they must consider the risks and benefits of those answers in the context of the specific case they are working on."

Perhaps the most widely discussed example of balancing the risks and rewards of artificial intelligence is the self-driving car. Far from being a rote exercise, programming an autonomous vehicle involves difficult choices that will generate extensive ethical and legal debate in the years ahead. In fact, these debates are already taking place, in the legislatures of the 40-plus states that have passed, or have considered passing, laws to govern self-driving cars.





The Ethical "Choices" Made By A Self-Driving Car

**An autonomous vehicle has lost control.
A collision is unavoidable. Which obstacle should the car avoid?**

This graphic offers a simplified model of the considerations that engineers must take into account when programming self-driving cars.



Source: Leon R. Sütfeld, Richard Gast, Peter König and Gordon Pipa. 2017.
Using Virtual Reality to Assess Ethical Decisions in Road Traffic Scenarios. *Frontiers in Behavioral Neuroscience*.

Who Should Be Responsible When A Self-Driving Car Crashes?

Take our poll now »





DUTIES OF SUPERVISION

Depending on who (or what) a lawyer works with, the duty of competence includes a duty of supervision. As Chris Mammen of Hogan Lovells explains, “If a lawyer delegates something to subordinates, whether junior lawyers or paralegals, there’s an ethical duty to make sure the work has been done competently. And this duty extends to AI-based tools. One way of analyzing the issue is that the lawyer who reviews and signs off has appropriately supervised the AI.”

And just as there are some tasks that a lawyer simply cannot delegate to a paralegal or legal assistant, there are some tasks that are not appropriate for handling by artificial intelligence – and an attorney must know how to tell them apart.

“One way of framing this issue is automation versus augmentation,” states Dr. Tonya Custis, a Research Director at Thomson Reuters who leads a team of research scientists developing natural-language and search technologies for legal research. “There may be some tasks that we shouldn’t automate. For these tasks, AI can help attorneys do their jobs, but AI can’t do their jobs completely. So the question becomes: where do we draw that line?”





“Think of an AI system like Westlaw Research Recommendations. It plows through huge amounts of data to suggest the relevance of a case — but the lawyer still has to decide that this case is actually relevant. The AI augments, but does not replace, the work of the lawyer.”

TONYA CUSTIS

Just as lawyers can over-delegate work to subordinates, they can also under-delegate, causing them to serve their clients less efficiently. In the context of artificial intelligence, one can imagine underutilization of AI – for example, a lawyer not using AI even though it could help that lawyer serve the client better.

In fact, given some of the psychological attributes commonly associated with lawyers – a focus on detail, a desire for control, an aversion to risk – the greater danger might very well be underutilization of, rather than overreliance upon, artificial intelligence.

“Having worked in AI for the legal profession for a long time, I know how the customer base is conservative,” says Tonya Custis of Thomson Reuters. “With Westlaw natural-language searching, lawyers will ask, ‘Why am I getting results that don’t use the specific words I searched for?’ You need to explain to the customer how the process works.”

☰ **In this webinar, David Curle of Thomson Reuters provides background information on artificial intelligence and discusses the many different ways that AI is being used in the legal world:**

Artificial Intelligence in Practice



To download the slides used in this webinar, enter your email here:

 ▼

I would like to receive additional Law2020 updates and resources from Thomson Reuters. I understand that I can unsubscribe at any time.





DOWNLOAD REPORT

DUTIES OF CLIENT CONFIDENTIALITY AND PRIVILEGE

When we think of issues of client confidentiality and attorney-client privilege, we often think of confidences that clients share with their lawyers. But those confidences are also sometimes “shared” with AI-powered tools – and lawyers need to ensure that client confidences remain secure at every stage of the process, especially in an age where every week seems to bring news of a new data breach.

According to David Curle’s [2018 AI Predictions](#), “legal professionals should be concerned about the ethical implications of the application of [AI technologies](#) to their practice,” including “confidentiality of hosted data used in AI applications and the risk of data breaches” and “risks related to confidentiality, privilege, and commingling of multiple clients’ data when using AI to analyze law firm billing data.”

Imagine, for example, a voice-activated personal assistant that can handle legal research questions (such as the RightsNOW App, a voice-activated legal information tool that just took top honors at the [2018 Global Legal Hackathon](#)). It’s a great innovation, but it must also be a secure innovation.

As Hogan Lovells partner Chris Mammen notes, “If all of the natural-language processing is done behind the law firm firewall, that’s one thing. But if it’s being handled by a vendor’s server somewhere out there, how sure are you that what could be confidential or privileged information is not being placed in a context where it isn’t adequately protected?”



☰ (Thomson Reuters, for example, employs an information security policy aligned to the well-known [NIST Cybersecurity Framework](#), and Westlaw offers multi-factor authentication (MFA) and two-factor (2FA) authentication for secure user log-in. These are just some of the security features that consumers should look for when selecting a technology provider.)

UNAUTHORIZED PRACTICE OF LAW

One of the most exciting developments in legal technology is the [rise of legal chatbots](#), AI-powered programs that interact with users who have legal issues by simulating a conversation or dialogue. These chatbots are now being used to do perform such tasks as fight parking tickets, advise victims of crimes, or draft privacy policies or non-disclosure agreements.

These chatbots can be very helpful to consumers, especially consumers who cannot afford the high cost of hiring a lawyer, and they could help bridge the yawning “justice gap” that exists in both the United States and around the world. But they do raise the issue of unauthorized practice of law, especially if the chatbot or other tool is created or maintained by an attorney.



“We don’t have rules and opinions that directly apply to these situations,” according to ethics lawyer Megan Zavieh. “We have to look at the spirit of the rules, and balance protecting the public with allowing for innovation in the delivery of legal services.”

“I get calls from lawyers who have creative ideas for helping people with legal problems, and they’ll tell me that they talked to three other ethics lawyers who told them ‘no.’ That’s the risk-aversion of our profession at work. But there has to be a way to innovate and move forward, to help consumers in different ways, and to close the justice gap, while at the same time not getting into disciplinary trouble.”

DIVERSITY AND INCLUSION

Artificial intelligence operates by looking for patterns with large amounts of data. This “training” of AI is, as Dr. Tonya Custis of Thomson Reuters puts it, “a statistical process — it will have biases.”

But what are those biases, and are they fair? If the data used to “train” the AI contains unfair biases, then the results of the AI could be correspondingly biased.

“AI requires data — data about actions and decisions made by humans,” explains David Curle. “If you have a system that’s reliant on hundreds of thousands or millions of human decisions, and those humans had biases, there’s a risk that the same bias will occur in the AI.”

For example, imagine training facial-recognition software on a group of people who come from only one racial or ethnic background, or training

voice-recognition software using only male voices. The resulting AI tools will be biased – not as inclusive as they should be, and not as useful either.

What Does AI Mean for Ethics and Diversity?



In the judicial system, one prominent example is judges making sentencing decisions based in part on AI-driven software that claims to predict recidivism, the likelihood of committing further crimes. There is concern over how the factors used in the algorithms of such software could correlate with race, which judges are not allowed to take into account when sentencing.

RESPONDING TO THE CHALLENGES OF ARTIFICIAL INTELLIGENCE

≡ The challenges that artificial intelligence pose to legal ethics, while significant, can be addressed — and should be addressed, so lawyers can take advantage of the powerful tools driven by AI.

EDUCATION AND TRAINING

Part of the lawyer's duty of competence involves keeping abreast of changes in law and in legal practice – and these changes, in 2018, inevitably involve technology.

“Large numbers of lawyers don't take this duty to keep up with technology seriously enough,” according to David Curle of Thomson Reuters. “It's not just AI-based technology but even more mundane things like practice management platforms, and other tools that make it easier and more efficient to practice law.”

“The ethical duty of competence requires being appropriately up to speed on technology,” says Chris Mammen of Hogan Lovells. “So AI is not something you can stick your head in the sand over, just as you couldn't try to conduct a document review in a major litigation entirely in paper.”

Lawyers must therefore have a general understanding of technology and artificial intelligence. And they must also understand the general operation of the specific AI tools that they use in their own practices.

“We need to have some understanding of what's going into an AI tool and what's coming out of it,” according to ethics lawyer Megan Zavieh, who represents lawyers facing disciplinary charges. Just as lawyers can't prove they satisfied their ethical duties simply by hiring an outside consultant, they similarly can't establish ethical compliance simply by using an AI tool.

At the same time, lawyers are not programmers – and the ethical rules recognize this, as David Curle notes: “The current rules of professional

responsibility are general enough to cover the situation. They suggest two things: that lawyers must understand enough about a new technology to see the risks, and that lawyers must understand enough to see the benefits.”

What this means in practice is that lawyers need to find trusted providers of AI-based solutions, and they need to pose smart questions to the providers whose AI tools they are considering using. Lawyers need to understand, at a basic level, how the solutions work and how the solutions were developed.

WHAT TO LOOK FOR IN AN AI TOOL

Understand the technology

No one expects you to understand exactly how the technology works, but legal professionals must have a basic understanding so that you are able to consider the benefits and risks.

Consider data security

Lawyers need to be conscious of choosing a solution that ensures client confidences remain secure at every stage of the process.

Understand data quality

If the data used to train the AI contains unfair biases, then the results of the AI could be correspondingly biased. Find a provider with trusted data.

Make sure the legal research work is being done competently

Even when choosing the right solution, you must be mindful that there are some tasks that are not appropriate for handling by AI, as well as some tasks where it would be unethical not to use the technology— and you must know how to tell them apart.



Find out more »

THE INDISPENSABLE HUMAN ELEMENT

Interestingly enough, in light of the whole “will robots take our jobs” fear, the role of human beings remains essential in developing AI. For example, consider Westlaw, which uses AI in a wide range of features, from Research Recommendations to Folder Analysis to Westlaw Answers.

“Hundreds of staff attorney editors write, curate, organize, and revise legal content for Westlaw,” according to Teri Kruk, Senior Director for Content Strategy and Editorial at Thomson Reuters Legal. “Former law clerks, law review editors, practitioners from large and small firms, prosecutors – all have chosen a path less travelled in the legal profession, but one that is critical to the continued life of the law and the profession.”

“Legal information and content is not fungible. Its creation—whether by the courts, legislatures, renowned authors, or in-house attorney editors—is deliberate, thoughtful, and part of a larger organism that is called ‘the law.’”

“Algorithmic searches, when run across deliberately and consistently organized information such as the content on Westlaw, necessarily yield better results,” Kruk adds. “But note how attorneys work alongside technologists and R&D to assess and validate results generated by algorithmic searches.”

CONSENSUS ABOUT ACCEPTABLE AI TOOLS

Lawyers operate within a network and a professional context. They must deal with colleagues, adversaries, courts, and regulators. And so even if an individual lawyer understands AI well, it's important for the other actors in the system to understand AI as well — and to have some sort of consensus about what AI is reliable and how AI can be appropriately used.

For example, take ediscovery and predictive coding. When predictive coding first emerged, there was discussion and dispute over whether and how it could be used. But today there is widespread acceptance that predictive coding in general, and specific programs or platforms in particular, are sufficiently reliable to be used. ([Thomson Reuters eDiscovery Point](#), for example, was very much designed with “defensibility” in mind.)

Artificial intelligence isn't perfect – but neither are people. As David Curle puts it, “The issue with AI is, ‘accurate compared to what? Humans make mistakes too.’”

RELATED READING

Not All Legal AI is Created Equal Ebook

Lawyers Assess the Risks of Not Using AI

Case Western: Ethical Implications of Legal Practice Technology



A CONSCIOUS EFFORT TO COMBAT UNFAIR BIAS

When it comes to addressing bias in AI tools, a little awareness goes a long way. As Tonya Custis emphasizes, "It's important that people practicing AI and training the models are aware of the biases in the tools and aware of how the models are getting implemented."

To make sure that AI tools are not unfairly biased, it's important to have diverse teams working on these tools. Custis poses this hypothetical:

"Say you have two guys working out of their garage in Silicon Valley. All of their friends are other guys. The technology they develop could look very different if only they had women on their team. A lot comes down to imagining who's going to use your product. Without diverse people on your team, you might not consider points of view or experiences that are obvious to those with different societal experiences."

Indeed, when it comes to rooting out unwanted bias, computers might have certain advantages over machines.

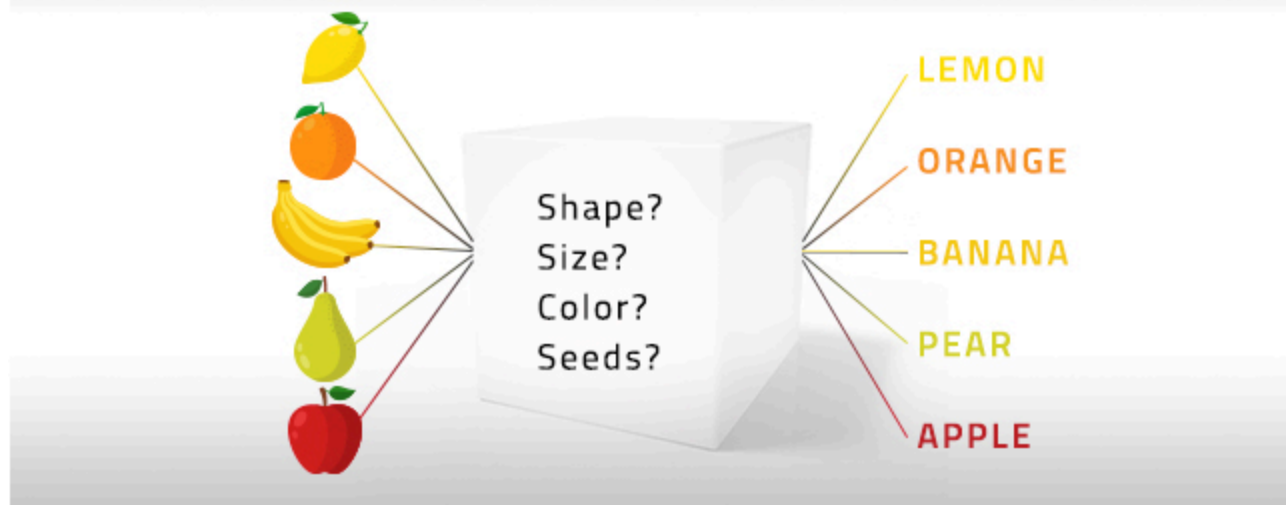
"Humans, for all our positive attributes, are fallible too," notes David Curle of Thomson Reuters. "The machines have the advantage of their biases being more predictable and perhaps more easily remedied. Data scientists have been working on this problem and are in a position to 'clean up' the biases in data. The biases baked into hundreds of thousands of years of human life are harder to root out."

The effort to reduce unfair bias in AI tools finds expression in the movement toward "explainable AI." As Tonya Custis explains, "There's a trend in the direction of 'explainable AI,' where we develop the models to be clear about how they generate their answers. This is especially important in the legal field, where the customers want to know: what are the factors that went into a given decision? Why did [Westlaw](#) recommend this particular case?"

In fact, explainability could become a legal requirement. In some jurisdictions, a “right to explanation” is starting to emerge.



“Black Box” vs. “Explainable” AI



THE PROMISE THAT ARTIFICIAL INTELLIGENCE PRESENTS TO THE LEGAL PROFESSION

“AI has the potential to transform the legal profession in so many positive ways,” predicts ethics attorney Megan Zavieh. “If we can start to ‘push down’ the work that takes up too much of our time to AI products, much as we’ve done with other forms of technology in other areas, we can free up lawyer time to do the things we do best: the legal analysis and arguing in court that can’t be replaced by robot lawyers.”

☰ This will make the provision of legal services more efficient and more effective. It will benefit clients, who will receive better service in a more cost-effective manner, and it will benefit lawyers, who currently suffer from high rates of burnout, anxiety, depression, and addiction.

And it's not just a matter of lawyers having more time to write briefs or draft contracts – it's also about the human element.

"One way that we lawyers fall short as a profession is in our human interaction," Zavieh explains. "Not enough lawyers are caring enough about their clients and their problems. As a result, we have not just a business-development problem but also an ethics problem: we need to be invested enough in our clients to put forth the effort that's required to see their matters through successfully."



"We don't need a crazy amount of additional time and energy to show our clients a lot more care and interest in their concerns. We need just a little more capacity, and AI is a great way to make that happen."

MEGAN ZAVIEH

○ At the end of the day, the legal profession is all about serving the client — and there's no denying that AI is a powerful tool for client service.

— are tools to support the practice of law and to help us advise our clients,” says Chris Mammen of Hogan Lovells. “We should make appropriate use of them to provide the best, most efficient, most cost-effective service to our clients.”

These technologies most definitely include AI – which we might not even think about as a discrete technology in the future. Per Zavieh, “We think AI is a huge thing right now. But in a few years, it won’t be thought of as ‘AI,’ and it will just be a useful tool.”

And AI technologies might eventually not just be useful tools, but even essential tools, for attorneys. As David Curle puts it:

“A lawyer’s duty of competence and diligence includes the duty to use tools and technology where appropriate. So at a certain point in time, a lawyer might have an ethical duty to affirmatively use AI, where that AI is accurate, reliable, and essential to serving the client effectively.”

In other words, in order to survive and thrive in a challenging marketplace, lawyers of the future will *need* to use AI.

LEAVE YOUR EMAIL BELOW AND WE'LL SEND YOU FREE UPDATES ON LAW2020.





I would like to receive additional Law2020 updates and resources from Thomson Reuters. I understand I can unsubscribe at any time.

SUBMIT

OUR SITES

**ABOVE
THE LAW**

BREAKING
DEFENSE

BREAKING
Energy

BREAKING
GOV 

DEALBREAKER

FASHIONISTA

MedCity News

© 2024 Breaking Media, Inc. All rights reserved. Registration or use of this site constitutes acceptance of our Terms of Service and Privacy Policy.





GENERATE TRUST

Generative AI is coming to Westlaw Precision.

[Join the Waitlist](#)

THOMSON REUTERS

- CURRENT ISSUE
- PAST ISSUES ▾
- SUBSCRIBE ▾
- CONTACT ▾
- ABOUT ▾

YOU ARE AT: Home » Articles » Model Ethics Rules as Applied to Artificial Intelligence



Model Ethics Rules as Applied to Artificial Intelligence

BY RAFAEL BACA ON AUGUST 14, 2020



Today, the public hears more and more of what artificial intelligence (AI) can do in the real world—things that are tangibly beyond encounters with robots or androids portrayed in popular science fiction movies and novels. Today, many people think AI is for simple tasks, like using voice commands to call a friend or purchase household goods from online retailers. But AI also is detecting cancer better than photographic slide pathologists, and governments are applying AI to detect specific individuals who are violating COVID-19 quarantine.

Commercially viable AI took shape in the mid-2010s, with the practical convergence of cloud computing and collections of big data. In short, affordable, large-scale probabilistic computations based on accumulated data stored on the cloud gave rise to our present use of AI for commercial purposes. The introduction of commercially viable AI permits software applications to apply accurate artificial reasoning to mundane tasks in our daily lives right now, with limitless possibilities in the future. Today, with simple off-the-shelf laptops, computer scientists can variably scale their computing power as needed to apply statistical algorithms to a wide range of datasets stored on the cloud, ranging in size from a few hundred to hundreds of billions of data points, to accurately predict future experiences and glean new insights on a wide range of inquiries.

How AI is Generally Applied to the Legal Field

Presently, the area of computer science of most impact to practicing lawyers is natural language processing (NLP) AI. NLP AI often capitalizes statistical tendencies observed in spoken human language. For example, there is a highly probabilistic tendency for some words to clump together more than others, and computer



Get even more resources to help you practice law better. Join the ABA Law Practice Division.

[Learn More](#)

Billing should be the easiest part of your job

[See how Clio makes it easy](#)



ACH Legal & Compliance Summit
Nov. 29–30, 2023 • Virtual

Free for attorneys. Earn AAP, APRP, CLE and CTP credits.

[REGISTER NOW](#) 

THIS ISSUE OF LP TODAY

The HR Issue | August 2020

programs can “learn” and exploit that statistical word clumping. For example, in the AI field of legal analytics or informatics, a lawyer will ask a computer or data scientist to create an NLP algorithm or “model” that looks for clumping of a familiar set of words used in the legal profession. This model may be used to review all opinions and orders delivered by Judge X, to gain insights on the likelihood of a future ruling by Judge X given similar circumstances. A slight variant would be Judge X asking the same computer or data scientist to help create software for auto-generating his or her orders using the data set of preexisting opinions and orders.

A wide array of legal software packages are available to auto-generate briefs, legal research results, and administrative or governmental forms, to locate key evidence from a document dump using text, sound and audio AI software tools, and even to provide virtual receptionists and paralegals in the forms of chatbots. At this time, a practicing lawyer is most likely to encounter this form of AI-enhanced software.

Big Data Analytics, Litigation, and E-Discovery

Money Ball is a movie based on the true story of how the Oakland A's used computer-based analytical statistical modeling to build a winning team on a severely limited budget. This “money balling” discipline of computer science is commonly referred to as data analytics or business intelligence. Often data or computer scientists will “webscrape” the entire publicly available internet—essentially collect at historically unprecedented scales every single bit of digital information about a desired topic or target—to form a collective data pool, called a “data lake,” on a cloud storage database.

Today, global law firms and corporations are effectively using statistical modeling and commercial AI to glean insights or intelligence on all aspects of litigation, especially eDiscovery, and on the activity of judges and competing law firms and attorneys. Those who have the most resources have the greatest competitive advantage using data analytics today. In time, the cost of data analytics intelligence will go down for most legal practitioners as the technology grows.

Closing the Justice Gap

Historically, in the field of computer science, software such as open-source platforms, and even the functional infrastructure of the internet, are based on a democratic process providing equal access for all users. A highly effective and active subset of lawyers and software developers are dedicated to using AI to ensure that access to legal remedies becomes truly democratic. These lawyers and developers believe that equal access to justice is a right and not a privilege, and see AI as a means for leveraging the same work output as a large law firm would provide. Some notable groups in this effort are the Legal Hackers, as well as the Free Access to Law Movement (FLAM). Other AI-driven software platforms, such as Torchlight Legal, provide *pro bono* immigration services for asylum law.

One present objective is to remove the pervasive paywall of privilege that exists for accessing legal services. AI-driven software tools that auto-generate documents for no-fault divorces and parking ticket appeals serve as a model for an affordable alternative to high-priced human lawyers. In time, as AI continues to grow, I see AI software handling even complex mergers and acquisitions and bankruptcy proceedings.

The ABA Model Rules of Professional Responsibility and AI

ABA Model Rule of Professional Conduct 1.1:

If you have arrived at this point and not yet fallen asleep, you have successfully satisfied the ethical requirements outlined in ABA Model Rule 1.1.

Soon, as legal software edges toward the end goal of satisfying the Turing test (a test to determine whether a computer is capable of thinking like a human being), legal professionals will incorporate AI-driven software in their daily practice. The ethical and economic concerns of lawyers being entirely replaced by AI software “robots” are largely unfounded. AI generally assists and enhances the professional decisions made

THE HR ISSUE
AUGUST 2020



IN THIS ISSUE

Maintaining a Secure and Effective Practice During the Pandemic

Creating Inclusive Legal Industry Workplaces in COVID-19

The Role of a Diversity Committee

Thriving, Not Just Surviving, While Working Remotely

The Dangers of COVID-19 in the Workplace

How Working Remotely Builds the Case for Accessibility

Why Now Is the Right Time to Move to the Cloud

Why Do Law Firms Need Human Resource Management Software?

Ethical Walls Support Firm Growth in the Post-Pandemic Economy

Model Ethics Rules as Applied to Artificial Intelligence

Law Firm Operational Resiliency in an Age of Uncertainty

Remote Working Implications for Information Governance

Making It Rain: Rachel Adcox

A New Resource for Law Firm Leaders to Learn from Each Other

by lawyers today, a concept in computer science called augmented intelligence. By analogy, lawyer-enhanced AI will be like a driver who resumes command of a cruise control feature in a car on a road trip.

It's critical to note that AI software is fueled by data, including legal information that is likely to be subject to the duty of client confidentiality, and possibly evidentiary privilege. Accordingly, much of the use of AI software in a law practice is ethically rooted in the same discussions that relate to mitigating the existing risks associated with client data privacy and security.

Specifically, Rule 1.1 states: "A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness, and preparation necessary for the representation."

Commentary to Rule 1.1 further clarifies and asks for some level of technical knowledge: "[t]o maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology..." This requirement is reinforced by insurance carriers informing legal practitioners of the risks of use of AI-driven software, such as, among others, the risk of cyber-attacks and data breaches.

To date, 37 states have adopted a technical competence rule arising from Rule 1.1. Lawyers should strive to understand the benefits and risks of applying new technologies in a legal practice. Ultimately, Rule 1.1 does not require that lawyers possess superior technical knowledge, but a general knowledge of the technology so as to effectively consult with experts when designing, adopting and using new AI software applications in their law practice, as well as ethically advocating for their clients.

Let's look at three other scenarios regarding the model rules of professional conduct and AI.

ABA Model Rule of Professional Conduct 1.6

Rule 1.6 (a), with limited exceptions, states, "A lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent..." Rule 1.6(c) states, "A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client."

Frequently in e-discovery, data analytics, and legal research, data is stored and retrieved on the cloud, and in instances of big data-like data analytics or AI modeling, data is retrieved in large quantities from a data lake as discussed above. Legal practitioners should be transparent with their clients as to how digital client information is stored and retrieved with respect to client confidentiality and privileges. Legal practitioners should consider incorporating a digital information and AI software disclosure statement in their engagement letters or initial client interview packets.

Thankfully, similar issues of patient data confidentiality have already been accommodated by many software tools in the health care field, as mandated by the Affordable Care Act and earlier federal laws such as HIPAA. Because of a federally mandated head start toward full automation of electronic health records, legal practitioners can often incorporate existing commercial patient privacy software and health care IT network infrastructures as an ethical and economical basis for adding leading-edge, robust privacy and security features to their law practice.

Consider the example of a legal practitioner applying anonymization algorithms and techniques to client data as it is received and stored. Categories of information from the collected raw data, such as names, addresses, expenditures, and other private information can be redacted using digital anonymization strategies to ensure client confidentiality under Rule 1.6. To limit costs, it is critical that a legal practitioner closely communicates to technical professionals what categories of data will require anonymization. In practice, the anonymized data is encrypted and can be used to build AI models while keeping critical anonymization in place. Such anonymization techniques are already used heavily in the medical field, so a pool of experienced healthcare software professionals is available for law practices to draw from. Similarly,

law practices can look to experienced financial industry software professionals to apply the latest anonymization and privacy and security techniques used in banking and accounting.

Despite all the technical jargon and concepts, a legal practitioner should understand that data is ultimately collected and used by humans with software tools. Lawyers should rest assured that, fundamentally, humans design and curate the datasets for legal client records as well as to drive AI legal software algorithms. As such, lawyers must consider the underlying human bias inherent to any dataset used by AI algorithms. Lawyers must remain informed and in constant communication with their software professionals to ensure that the optimal results from AI algorithms arise from the highest-quality data.

U.S. export controls also require digital data to stay within the physical boundaries of the United States (see 15 C.F.R. §730 *et seq*). In practice, a legal practitioner should be mindful as to where the computer servers storing client data are geographically located while accessed “on the cloud.” Obtaining the location of a client’s digital data is very easy for software professionals to determine as well as to request that such data servers be located exclusively in the USA. Amazon Web Services, among other cloud vendors, already provide software architectures for ensuring HIPAA privacy that can readily be exported to constructing legal databases respecting confidentiality, as well as for selecting the physical location of the cloud server network that contains the digital data.

ABA Model Rule of Professional Conduct 5.3

Rule 5.3 states, “Responsibilities Regarding Non-lawyer Assistance... (b) the lawyer having direct supervisory authority over the nonlawyer shall make reasonable efforts to ensure that the person’s conduct is compatible with the professional obligation of the lawyer...”

While the above discussion of Rule 1.6 applied confidentiality and privilege strategies to the data, Rule 5.3 similarly extends those requirements to “legal assistance” provided by legal software and software professionals. This includes, for example, third-party software professionals as well as assistance from AI-driven legal software itself.

In practice, in light of Rule 5.3, software professionals must be aware of a lawyer’s obligation to their clients under the Rules of Professional Conduct. Legal practitioners should strive to educate these software professionals on legal confidentiality and privileges through workshops, online videos, and checklist handouts. A lawyer should help software professionals understand the concepts of client confidentiality, preserving evidentiary privilege as applied to digital data privacy, security, and while using client data to run artificial intelligence software tools. As one solution for maintaining additional client confidentiality and privilege under Rule 1.6(c), a lawyer should be aware of the processes of software professional workflows, such as Agile approaches to software development, as well as technical concepts relating to data privacy, security, and AI software applications to ensure that lawyers communicate effectively within tech culture.

A slightly more difficult consideration is applying Rule 5.3 in light of ensuring that software, AI-driven or not, adheres to the model rules, namely by handling the privacy and security of client data, including attorney-client confidentiality and privilege. Generally, many but not all software professionals are mindful of data privacy and security while developing commercial software, but may need additional education by legal practitioners on the topic of data privacy and security as it relates to client confidentiality and privilege. Currently, no federal law in the U.S. requires software professionals to adhere to data privacy and security.

As a practical matter, one good legal reference to data privacy and security is outlined in the European Union’s General Data Protection Regulation, (GDPR), and its legislative protégé within the U.S., the California Consumer Privacy Act (CCPA), which currently acts as a proxy for a federal body of law on data privacy and security in the United States. Other helpful references are the existing tapestry of federal laws often associated with health and financial privacy laws.

ABA Model Rule of Professional Conduct 2.1

Rule 2.1 states, “In representing a client, a lawyer shall exercise independent professional judgment and render candid advice,” which potentially can involve referring “not only to law but other considerations as

moral, economic, social and political factors, that may be relevant to the client's situation."

At this time, AI algorithms allow commercial legal software to automatically generate legal documents from briefs to patent search results and judicial opinions. The economic and time-saving temptations of simply signing-off on AI-generated work products are great to legal practitioners, especially in private practice.

Rule 2.1 directly addresses the ethical duty of a lawyer to avoid the temptation of entirely relying on the output of AI legal software, and to exercise independent professional judgment to supplant those conclusions directly rendered by the AI software. The independent judgment of lawyers under the Rule goes beyond just legal matters, but must also account for the *totality* of factors associated with the client's situation, namely moral, economic, social, and political factors, so as to remain in the four corners of Rule 2.1. Indeed, the intent is for legal practitioners to think long and hard about the interests of their clients before relying on AI-generated work products.

Conclusion

In general, AI technologies will create unique challenges for legal practitioners beyond those presented by collecting data for cloud computing while ensuring lawyer-client confidentiality and privilege. Client data is being collected, managed, used, and stored indefinitely in new ways with today's AI technology. The Model Rules of Professional Conduct are clear in requiring lawyers to ensure these evolving tools do not endanger client confidentiality and privilege.

About the Author

Rafael "Rafa" Baca is a patent attorney with [The Adelante IP Law Group](#) and software developer and is chair of the [ABA Artificial Intelligence Committee](#).

LAW PRACTICE TODAY EDITORIAL BOARD

Editor-in-Chief

Ramón Viñas-Bueso, Viñas Law Office LLC

Jason Marsh, MARSH8

Issue Editor

September 2023

Alan Klevan

Micah Buchdahl, HTMLawyers

Rodney Dowell, Rodney Dowell Office of Bar Counsel

Nicholas Gaffney, Zumado

Michael Kraft, Kraft & Kennedy Inc

Jason Long, London & Amburn PC

Andrea Malone, *Fox Rothschild*

Jason Marsh, MARSH8

Mark Robertson, Robertson & Williams

Ramón Viñas-Bueso, Viñas Law Office LLC

[ABA Law Practice Division News](#)

[About ABA Law Practice Today](#)

[Advertise](#)

[Contact](#)

[Information](#)

[New Law Anthology](#)

[Submit a Pitch](#)

[Syndicate](#)

[Current Issue](#) · [Past Issues](#) · [Terms of Use](#) · [Code of Conduct](#) · [Privacy Policy](#) · [Your Privacy Rights](#) · [Your California Privacy Rights](#) · [Copyright & IP Policy](#)
· [Advertising & Sponsorship](#) · © 2015 ABA, All Rights Reserved