BAR COUNSEL

Understanding the Ethics Surrounding Generative AI

Remain Vigilant

By Ankur Doshi

iStock.com/posteriori

Throughout 2023, generative AI and its potential ethical pitfalls were some of the hottest topics for discussion at every legal conference. When ChatGPT opened to the public in 2022, ethics experts raised questions about whether lawyers could use this powerful technology effectively and without creating ethical quandaries.[1] A number of industries, including the legal industry, raced to adopt new technologies based on generative AI.

It was not long until the first case of an attorney misusing a generative AI tool (ChatGPT) arose in court proceedings. In June 2023, attorneys for the plaintiffs in *Mata v. Avianca, Inc.* were sanctioned $5,000 for filing an opposition drafted by ChatGPT that contained fictional citations and authority, and then again using ChatGPT to provide fictional excepts to the court.[2] Shortly thereafter, other courts grappled with ethical issues surrounding generative AI. Many issues within these cases could be avoided through the understanding of the Rules of Professional Conduct (RPC) and their application to these technological tools.

Technology and Ethics

It's no surprise that the legal profession lags in the adoption of new technologies. With new shifts in technology, the law struggles to adopt until technology becomes more mainstream. We've seen this challenge in the past with email and again with cloud computing.[3] As these changes become more frequent, the RPCs are crafted to provide the flexibility necessary to apply to rapidly evolving technology.[4] As the preamble notes, the ABA Model Rules of Professional Conduct and our RPCs are rules of reason.[5]

Even with the unique abilities and challenges of generative AI in the legal field, we find that applying the underlying principles of the RPCs provide guidance as to how to best use this technology while also protecting our clients' interests and the integrity of the legal system.

Competence (RPC 1.1)

A fundamental principle of legal ethics is that lawyers must provide competent representation to their clients. As noted by RPC 1.1, competent representation "requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation."

While RPC 1.1 focuses on substantive knowledge of the areas of law in which a lawyer practices, the rule inherently includes the competence more broadly required in the practice of law, including competent use of technology. The ABA made this inference explicit in the comments through the Commission on Ethics 20/20 in 2012.[6] Comment (8) to ABA Model Rule (MR) 1.1 notes that

> 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.

(emphasis added). Comment [8] recognizes an ongoing duty to update our knowledge of current practices, including understanding the benefits and risks associated with technologies utilized in the practice of law. Within the past 20 years, lawyers have witnessed the practice of law fundamentally change with the introduction of technologies such as online legal research, email, electronic filings in court and e-discovery platforms. Many technologies which were once new and the subject of intense scrutiny are now mandatory to practice law.[7]

Understanding the risks and benefits of technology is an area of competence frequently overlooked by attorneys. It requires learning how the technology functions at a high enough level that we are aware of the risks associated with it. Attorneys are better equipped to evaluate the risks to their clients by understanding the technology.

Many of the ethical lapses involving attorneys and AI start with a misunderstanding of generative AI itself. Within five months after the introduction of ChatGPT to the public, an attorney used the program to generate an opposition to a motion and submitted it to a court in *Mata v. Avianco, Inc.*[8] ChatGPT "hallucinated"[9] in the filing, generating names to cases that sounded correct, but did not exist.[10] Five additional instances of attorneys submitting documents with fictional cases generated by AI to courts followed afterward.[11]

In each instance, the attorney's explanation for submitting documents citing nonexistent case law revealed that the attorney misunderstood the fundamentals of the technology. Attorneys stated that they believed that these generative AI programs were "super search engines" akin to Google.

They did not understand the potential risks of using the new technology, including the potential for fabricating citations, nor did they properly evaluate what had been generated, using their own professional judgment to supervise the work performed by the tool. Attorneys that understand the risks employed with AI know that they must evaluate the content created by AI to ensure accuracy.[12]

These cases highlight the need for attorneys to be competent in technology utilized in their practice. If we want to use these technologies, we must also fundamentally understand how they work.

Confidentiality (RPC 1.6)

One of the unique aspects of AI tools, and especially generative AI tools, is their ability to improve through iteration. However, iteration requires massive amounts of source data.[13] Generative AI incorporates user prompted data to train and improve its responses to user requests for information. The use of data and information within generative AI raises significant concerns for lawyers given our duty to protect confidential client information from unauthorized disclosure.

Our duties to protect client information are found in RPC 1.6. Specifically, RPC 1.6(a) states:

a) 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).

RPC 1.6(c) further specifies:

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.

"Information related to the representation of a client" is a defined term under RPC 1.0(f) denoting "information protected by the attorney-client privilege under applicable law, and other information gained in a current or former professional relationship that the client has requested be held inviolate or the disclosure of which would be embarrassing or would be likely to be detrimental to the client." The Oregon Supreme Court has noted that this definition is broad in scope and may encompass publicly available information if such information could be detrimental or embarrassing to the client and if not generally known.[14]

At first look, generative AI and the attorney's duty of confidentiality may appear incompatible. However, introductions of new technologies in the practice of law and the ethics analysis involved with those new technologies provide attorneys with guidance as to how to ethically utilize generative AI if they wish to do so. Oregon Formal Opinion 2011-188 states that attorneys have the duty to review the capabilities of the technology to securely store client data. That means that attorneys must make a "reasonable effort" to determine whether a third party follows industry standards for confidentiality and security.[15]

RPC 1.6 requires us to understand what happens to our data when it is submitted to an AI program. If an AI program is storing an attorney's data for other purposes, an attorney must evaluate if they can risk utilizing the AI program, even if they're using nonspecific and hypothetical data, without exposing information related to representation of the client.[16] Comment (19) to Model Rule 1.6 provides us with guidance, noting "(f)actors 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."

As an example, ChatGPT generally does not provide protection for confidential information.[17] It provides warnings within its terms and conditions that data provided to ChatGPT will be utilized to "provide, maintain, develop, and improve our Services, comply with applicable law, enforce our terms and policies, and keep our Services safe."[18] An option is also provided to opt out of data collection.[19] Newer AI tools specializing in assisting lawyers are affording more robust data protection and specifically stressing the protection of client data in the same manner as many other legal software solutions.[20] However, an attorney seeking to use such tools must review the terms and conditions of such software to ensure that the necessary protections are available to keep their client's information confidential.

An additional aspect of utilizing an AI tool, and any tool that stores client information online, is the potential for a data breach.[21] Incorporating AI requires an understanding of the security measures in place that secure and protect the data from data breaches and other potential intrusions. Attorneys should undertake reasonable steps, including reviewing the terms and conditions and the privacy policy, of AI tools to determine whether the AI tool complies with industry standards for confidentiality and security.[22]

Duty to Communicate under RPC 1.2 and 1.4

If you decide to use a secure AI tool in your practice, do you need to inform your client that you use such a tool? While AI may feel like a research tool like Westlaw or Lexis, its capabilities may require disclosure of use to the client, along with discussion to allow the client to make an informed decision.

Our duties under RPC 1.2 and 1.4 define how we must communicate with our clients about the representation:

RPC 1.2(a): . . . (a) lawyer shall abide by a client's decisions concerning the objectives of representation and, as required by Rule 1.4, shall consult with the client as to the means by which they are to be pursued. A lawyer may take such action on behalf of the client as is impliedly authorized to carry out the representation.

RPC 1.4(b): A lawyer shall explain a matter to the extent reasonably necessary to permit the client to make informed decisions regarding the representation.

Both of these RPCs create a duty to abide by the client's objectives and provide enough information to the client to allow them to make informed decisions. However, lawyers do not have to inform their client of every single step of the representation.[23] For instance, lawyers do not generally disclose the type of electronic research database or the type of word processing software they use.

Whether there is a need to communicate with the client on the usage of AI can vary substantially depending on how an attorney uses AI. For example, relatively trivial uses of AI, such as word suggestions in word processing[24] can be nonconsequential to the representation and do not require disclosure.

On the other end of the spectrum, in a recent criminal trial in D.C., an attorney

utilized a never-before-used AI program to draft the closing argument of the trial without informing the client defendant.[25] The defendant is now seeking a new trial based on ineffective assistance of counsel, noting that the closing argument used by his counsel made frivolous arguments and ignored multiple weaknesses within the prosecutor's case.[26]

With the unique aspects of generative AI, and its relative infancy in legal practice, lawyers are "well-advised to consult with clients before using generative AI to assist with anything other than *de minimis* usage."[27] Guidance from New Jersey reaches a similar conclusion, noting "(The New Jersey) RPCs do not impose an affirmative obligation on lawyers to tell clients every time that they use AI. However, if a client asks if the lawyer is using AI, or if the client cannot make an informed decision about the representation without knowing that the lawyer is using AI, then the lawyer has an obligation to inform the client of the lawyer's use of AI."[28] Considering the breadth of potential uses for AI, transparency will be a critical factor in ensuring that our clients understand how we use generative AI and what information we use with generative AI.[29] It also provide us with an understanding of the concerns the client may have about confidentiality and AI. Many firms experimenting with generative AI have explored placing initial disclosures within their engagement letters to start the discussion with their clients.[30]

Duty to Supervise under RPCs 5.1 and 5.3

Imagine a scenario where you request your paralegal to prepare a letter to your client on a matter. The paralegal sends you a draft letter to review within 30 minutes. Amazed by the efficiency of your paralegal, you review it, don't see any issues with the letter and sign off on it. Two weeks later, you get an angry call from your client. There has been a data breach at FoundationAI, a generative AI tool open to the public. The data was published, and apparently your client's information was in the data, even though your firm does not contract with FoundationAI. After discussing with your staff, you learn that your efficient paralegal has been using FoundationAI for drafting correspondence and other memoranda. In order to do that, she's provided client information through the prompts to tailor the letters.

This scenario has occurred[31] and is a likely scenario that many lawyers may encounter related to generative AI. As these publicly available models become more available and useful, there will be a substantial number of associates and nonlawyer staff tempted to use the programs for certain types of work, especially letters and other correspondence. Attorneys need to be prepared to provide guidance and training to lawyer and nonlawyer staff alike about acceptable technologies and the risks that such technologies pose.

RPC 5.1 and 5.3(a) place responsibility on the supervising attorney for the work of their associate attorneys and nonlawyer staff. RPC 5.1 states:

> A lawyer shall be responsible for another lawyer's violation of these Rules of Professional Conduct if:
>
> (a) the lawyer orders or, with knowledge of the specific conduct, ratifies the conduct involved; or
>
> (b) the lawyer is a partner or has comparable managerial authority in the law firm in which the other lawyer practices, or has direct supervisory authority over the other lawyer, and knows of the conduct at a time when its consequences can be avoided or mitigated but fails to take reasonable remedial action.

And RPC 5.3(a) states:

> With respect to a nonlawyer employed or retained, supervised or directed by a lawyer:
>
> (a) 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

Attorneys in supervisory roles have an obligation to know the technology their associates and nonlawyer staff are using and make sure they conform with the RPCs. This rule extends to third parties that the lawyer contracts with for services to their firm. ABA Comment (3) to Model Rule 5.3 emphasizes:

> "When using such (nonlawyer) 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."

Because of the transformative capabilities of AI, it is paramount that attorneys provide time for discussions and training to staff about how and when to use AI, and when using AI increases risk and exposure to a violation of our ethical duties.

To avoid ethical pitfalls, attorneys must be proactive. Attorneys with staff should consider offering proper training and education on AI tool usage, establishing expectations about what tools employees can use, setting clear policies and procedures for using AI technologies, and implement robust quality control and review to ensure their obligations under RPC 5.1 and 5.3 are met. Absent any type of training, policies, or education for their staff, attorneys face a significant risk.

Candor to the Court

Generative AI does not change our duty of candor to the court encompassed within the RPCs. Under RPC 3.3(a):

> a) A lawyer shall not knowingly:
>
> (1) make a false statement of fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the lawyer;(.)

An attorney's submission of unverified and fictitious cases cited by a generative AI obviously runs afoul of this RPC, as well as several other RPCs related to truthfulness.[32] The publicity of several cases involving fictitious cases cited by generative AI has raised additional concerns with the judiciary. Several judges have standing orders prohibiting the use of generative AI.[33] Additional judges have required disclosure of the use of generative AI.[34] In November 2023, the U.S. Court of Appeals for the 5th Circuit specified certification of either that a document was not drafted by generative AI, or that a document drafted by generative AI was further reviewed for accuracy by a human.[35] Attorneys utilizing generative AI should be aware of any standing orders within the courts they practice in case they do require disclosure of generative AI.

Additional Guidance

Oregon's Legal Ethics Committee is reviewing whether additional guidance on AI would assist Oregon attorneys in clarifying their ethical duties. Currently, California,[36] Florida,[37] New Jersey[38] and North Carolina[39] have issued opinions related to gen-

erative AI. Other states are also considering releasing additional opinions. California issued practical guidance on generative AI late in 2023, which noted several best practices and considerations for using generative AI. It also noted:

> 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.[40]

Florida echoed those sentiments and reiterated an attorney's duty of competence in using generative AI:

> "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."[41]

Each of these states evaluated many of the same questions raised here and provided guidance to their practitioners. If you believe that an opinion related to AI would be helpful, Oregon's Legal Ethics Committee would like to hear from you. You can provide feedback to them at *feedback@osbar.org*. ∎

Ankur Doshi is general counsel for the Oregon State Bar. He wrote about the 2023 Client Security Fund Report in the February/March Bulletin.

ENDNOTES

1. Andrew Perlman, "The Implications of ChatGPT for Legal Services and Society," *The Practice,* March/April 2023, at *https://clp.law.harvard.edu/knowledge-hub/magazine/issues/generative-ai-in-the-legal-profession/the-implications-of-chatgpt-for-legal-services-and-society/* (noting potential challenges in ChatGPT for lawyers related to accurate and reliable results and accounting for the nuances and complexities of law).

2. *Mata v. Avianca, Inc.*, No. 22-cv-1461 (PKC), 2023 U.S. Dist. LEXIS 108263, at *2 (S.D.N.Y. June 22, 2023) ("[Plaintiff's attorneys] abandoned their responsibilities when they submitted non-existent judicial opinions with fake quotes and citations created by the artificial intelligence tool ChatGPT, then continued to stand by the fake opinions after judicial orders called their existence into question.").

3. Andrew Perlman, "The Legal Ethics of Generative AI," at 3, 17, (February 22, 2024), *Suffolk University Law Review*, forthcoming, available at SSRN: *https://ssrn.com/abstract=4735389* (" When [email] first became available, ethics opinions urged considerable caution and even suggested that lawyers might violate their duty of confidentiality by using it. We have now reached the point where lawyers must have an email address in order to remain licensed to practice law.").

4. *Id.* at 18 *citing* 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).

5. Model Rules of Prof'l Conduct: Preamble: Scope, Comment [14] (2024).

6. ABA Commission on Ethics 20/20 Report to the House of Delegates, 105A, August 2012.

7. *See In re Schaefer*, 33 D.B. Rptr. 461 (2019) (attorney expected to know how to communicate with clients and court via email); *In re Fjelstad*, 31 D.B. Rptr. 268 (2017) (attorney expected to know how to utilize electronic filing with the court); *In re Heinzelman*, 31 D.B. Rptr. 11 (2017) (same). *See also* Perlman *supra* note 3 at 17 (attorney requires email address to be member of California State Bar).

8. *Mata v. Avianca, Inc. supra* note 2.

9. *See* Jordan Murphy, "AI in the Legal Sector – an Overview for Information Professionals, *Legal Information Management*, 23 (Sept. 2023), at 150 ("Known limitations of the [ChatGPT] tool include 'hallucinations'; outputs that sound plausible but are factually incorrect. These can arise from misunderstanding or a lack of contextual understanding.").

10. *Mata v. Avianca, Inc. supra* note 2.

11. *Park v. Kim*, 91 F.4th 610, 615 (2d Circ. 2024) (Second Circuit referred an attor-

ney for submitting brief relied on fictional authority created by ChatGPT); *Smith v. Farwell*, 2282CV01197, (Superior Court, Norfolk, Mass., Feb. 12, 2024) (attorney sanctioned $2,000 by court for submitting several legal memoranda relying on wholly fictitious case law generated by an AI system); *U.S. v. Cohen*, 18CR602JMF. 2023 WL 8635521 (S.D.N.Y. Dec. 12, 2023) (client (former attorney) sent attorney fictional cases citations on Google Bard AI that were submitted to the court); *People v. Crabill,* 23PDJ067, 2023 Colo. Discip. LEXIS 64, (Sup. Ct. Colorado, Nov. 22, 2023) (Attorney suspended for one year and one day, with 90 days served, for submitting a motion that containing citations created by AI and not verifying them); *see also* David Wagner, "This Prolific LA Eviction Law Firm was Caught Faking Cases in Court. Did They Misuse AI?," *LAist.com*, Oct. 12, 2023, at *https://laist.com/news/housing-homelessness/dennis-block-chatgpt-artificial-intelligence-ai-eviction-court-los-angeles-lawyer-sanction-housing-tenant-landlord* (attorney was sanctioned for submitting filing with fictional cases in eviction case). *See generally* Pranshu Verma and Will Oremus, "These lawyers used ChatGPT to save time. They got fired and fined." *Washington Post*, Nov. 16, 2023, at *https://www.washingtonpost.com/technology/2023/11/16/chatgpt-lawyer-fired-ai/*.

12. 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," Nov. 16, 2023 at *https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf* ("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").

13. *See* Henry H. Perritt, Jr., "Robots as Pirates," 73 *Cath. U.L. Rev.* 57, 75 ("Large language models learn from enormous databases of text . . . [these] databases are constructed from crawling the Web, for specialized sources such as FreeLaw and PubMed, from online discussion groups like Reddit, from the Gutenberg books collection, and from Wikipedia").

14. *In re Conry*, 368 OR 349, 362 (2021).

15. Or. Formal Op. 2011-188; *see also* Or. Formal Op. 2005-141.

16. California's Practical Advice on Generative AI notes "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." Cal. Practical Advice, *supra* note 12 at 2.

17. Foster J. Sayers, "ChatGPT and Ethics: Can Generative AI Break Privilege and Waive Confidentiality?" 20 *N.Y. Law Journal* 269 3(2) (January 31, 2023); *see also* David Canellos, "What to Know About Sharing Company Data with Generative AI," *Forbes*, Aug. 10, 2023) at *https://www.forbes.com/sites/forbestechcouncil/2023/08/10/what-to-know-about-sharing-company-data-with-generative-ai/?sh=7135e06c6022*.

18. OpenAI, *Terms of Use* (November 14, 2023) at *https://openai.com/policies/terms-of-use*.

19. *Id.*

20. *See* John Tredennick and William Webber, "Attorneys Using AI Shouldn't Worry About Waiving Privilege," *Law360* (Aug. 22, 2023), at *https://www.law360.com/articles/1706972/attorneys-using-ai-shouldnt-worry-about-waiving-privilege*.

21. Romaine C. Marshall and Gregory Cohen, Artificial Intelligence Applications and The Rules of Professional Conduct, 36 *Utah Bar J.* 18 (Sept./Oct. 2023) ("Generative AI applications are at risk of data breaches to the same extent as any other software").

22. See Or. Formal Op. 2011-188 at 1-2.

23. See Pearlman, *supra* note 3 at 7.

24. *Id.*

25. *United States v. Prakazrel Michel*, Case No. 19-CR-148-1 (CKK), Def's Mot. For New Trial, (Oct. 16, 2023) 22.

26. *Id.* at 23-24*. see also* Jon Brodkin, "Rapper Pras' Lawyer Used AI To Defend Him in Criminal Case—It Did Not Go Well," *Ars Technica*, October 18, 2023, at *https://arstechnica.com/tech-policy/2023/10/fugees-rapper-blames-conviction-on-his-lawyers-ai-fueled-closing-argument/*.

27. See Perlman, *supra* note 3 at 7.

28. N.J. Supreme Court Committee on Artificial Intelligence and the Courts, "Notice-Preliminary Guidelines on The Use of Artificial Intelligence By New Jersey Lawyers," (Jan. 24, 2024) at 4-5.

29. Kassi Burns, "Demystifying Generative AI," 63 *Judges' Journal* 13 (Winter 2024).

30. See Isabel Gottlieb, "Law Firms Wrestle with How Much to Tell Clients About AI Use," *Bloomberg Law*, (Nov. 29, 2023) at *https://news.bloomberglaw.com/business-and-practice/law-firms-wrestle-with-how-much-to-tell-clients-about-ai-use*.

31. *See* Wagner, *supra* note 11 (noting that lawyer sanctioned for filing indicated that a first-year attorney had drafted the opposition citing fictional cases).

32. *See* RPCs 3.1, 4.1, 8.4(a)(3).

33. 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). *See generally* Perlman, *supra* note 3 at 11-12 (discussing standing orders banning the use of generative AI).

34. Magis. J. Gabriel A. Fuentes, Standing Order for Civil Cases Before Magistrate Judge Fuentes, (N.D. Ill. May 5, 2023); J. Brantley Starr, Mandatory Certification Regarding Generative Artificial Intelligence, (N.D. Tex. May 30, 2023). *See generally* Perlman*, supra* note 3 at 13-15 (discussing orders requiring disclosure of the use of generative AI).

35. 5th Cir. R. 32.3 (Jan. 4, 2024).

36. Cal. Practical Advice, *supra* note 12.

37. Fla. Eth. Op. 24-1 (Jan. 19, 2024), at *https://www.floridabar.org/etopinions/opinion-24-1/*.

38. N.J. Guidance at *supra* note 28.

39. N.C. Proposed 2024 Formal Eth. Op. 1, "Use of Artificial Intelligence in a Law Practice" (Jan. 18, 2024).

40. Cal. Practical Advice, *supra* note 12 at 1.

41. Fla. Eth. Op. 24-1, *supra* note 37 at 7.

Our Editorial Policy

All articles published in the *Bulletin* must be germane to the law, lawyers, the practice of law, the courts and judicial system, legal education or the Oregon State Bar. All opinions, statements and conclusions expressed in submitted articles appearing in the *Bulletin* are those of the author(s) and not of the editor, other editorial staff, employees of the Oregon State Bar, or members of the Board of Governors.

Publication of any article is not to be deemed an endorsement of the opinions, statements and conclusions expressed by the author(s). Publication of an advertisement is not an endorsement of that product or service.

Any content attributed to the Oregon State Bar or the Board of Governors is labeled with an OSB logo at the top of the page or within advertising to indicate its source or attribution.

THE LEGAL WRITER

Lean on Generative AI but Hang Onto Your Skills and Style

Writing in a New Age

By Rebekah Hanley

Generative artificial intelligence has taken legal professionals on a fascinating journey over the past year or so. As lawyers have read about and experimented with large language models, they've experienced a range of reactions: disbelief and denial; amazement and wonder; fear and loathing; enthusiasm and hope. Where opportunistic entrepreneurs see potential profits, risk-averse lawyers see likely liability, for both themselves and their clients. Practical and ethical questions abound. Optimistic lawyers hope that generative AI will, by streamlining tedious tasks, enhance access to justice and improve lawyers' professional satisfaction and longevity. Meanwhile, pessimists worry that this technological advancement will reduce the demand for lawyers — licensed, knowledgeable, skilled problem-solvers.

However you might feel about generative AI, there is no avoiding it. It's becoming more prevalent and more reliable with lightning speed. General-purpose large language models quickly revealed their potential — and their problems. But law-specific models with enhanced features, like data privacy protection and integrated citations linked to primary sources, are addressing many of those concerns.

As legal writers, we have a lot to think about: What can large language models do for us, what writing-related risks do they introduce and how can we minimize those risks?

Cautious Embrace Despite Serious Risks

For anyone who has chosen to look the other way since December 2022, here is a summary of where things stand. Move over, templates, forms and brief banks. You can now generate prose — an objective memo, a client email, a persuasive argument, a contract clause or virtually anything else you can imagine — in well under a minute by entering a prompt into a large language model, one type of generative AI tool. What you get in response may not be as accurate or as thorough as you need it to be, but it's a start; you can then tease more information out of the large language model, and you can adjust and expand the work product that your commands prompt the model to generate.

I'm trying to use my words with care here. Even when you turn to a large language model for drafting help, you, the human, remain the "writer." The quality of the output depends on your instructions — your precision, specificity and judgment. The large language model cannot think, craft or compose — not in the ways we understand those words. Instead, it connects units of language based on patterns. That means it's not really a co-author or a ghost-writer; it can't assume responsibility for work quality or receive credit for ideas. It can merely provide some drafting assistance, or automation, to the legal writer, who must then assess and revise with care.

Still, lawyers are cautiously embracing large language models because those tools are capable of materially improving the lawyers' efficiency. And to the extent that those tools support the creation of thorough, accurate legal prose, they can facilitate the delivery of cost-effective representation. Indeed, so long as lawyers can adequately address confidentiality, accuracy and other ethical concerns, soon lawyers will likely need to use these tools to demonstrate technological competence and avoid overbilling clients.

I need to emphasize the "cautious" nature of lawyers' growing interest in generative AI. We've received ample warnings about the enormous, costly risks of inappropriate reliance on a large language model. In case the cringe-worthy cautionary tales have somehow escaped your attention, here's the lesson a few lawyers have learned in uncomfortable, embarrassing and expensive ways. Large language models "hallucinate" — or they would, if they were human. That is, the models present fiction as fact, introducing errors. To support their answers, they sometimes offer imagined authorities — descriptions of, and citations to, scientific studies and judicial decisions that do not exist.

As a result, lawyers must check the veracity of everything a large language model spits out. Updating authorities using Shepards or KeyCite is a well-established professional responsibility, but generative AI requires more. A large language model is like a nonlawyer assistant with near-perfect grammar and punctuation, along with a vivid imagination and supreme overconfidence. On the plus side, that assistant never gets distracted or grows fatigued. Still, "supervising" that assistant requires vigilance, even with the introduction of new, guard-rail-enhanced, law-specific tools.

As responsible agents, fiduciaries and professionals, lawyers must use caution

when leaning on generative AI assistants. Lawyers will need to honor clients' preferences and read the fine print in all terms of service before charging ahead with new generative AI tools. But lawyers also need to accept that writing with the help of a large language model is where our profession is headed.

Less Obvious Costs

The ability to produce effective written communication is among lawyers' most critical professional skills. Indeed, some lawyers' earned reputation as effective writers is *the thing* that reliably brings them new business; the careful writing they produce can help persuade decision-makers, satisfy clients and spur referrals.

All lawyers write (at least occasionally), but they do so with varying degrees of success. Even setting aside incomplete, inaccurate or unpolished work product, not all legal writing is equally effective. Based on 20 years of experience teaching legal writing, I can confidently report that individuals relying on the same facts, legal authorities, and rules can produce wildly different written work products.

That variation occurs because a lawyer's writing is shaped by that lawyer's choices — countless decisions about what to say and exactly how to say it. Some of those choices are deliberate, flowing from careful thought, creative experimentation, and focused revision; others are the semi-automatic result of the writer's deep-seated preferences and experience-based judgment. The cumulative effect of those choices is the writer's style and voice — a signature woven into the fabric of the text.

So what happens when we delegate the creation of a first draft to a generative AI assistant? Perhaps leaning on a large language model is just like delegating a first draft to a law clerk or junior attorney: Either way, the senior lawyer gets something better than a blank page as her personal starting point. Then again, perhaps material differences separate those two strategies. After all, the large language model does not — cannot — "think," as a less-experienced human writer can. Instead, the large language model predicts the next most likely word or character. A law clerk or junior attorney might provide you with a brilliant first draft; in contrast, at least for now, a large language model offers something predictably average — something likely to be serviceable but with little to no chance of being extraordinary.[1]

That's cost number one of turning to a large language model for a first draft: Law school education plus capable supervision can help a junior legal professional craft fabulous work product that you are unlikely to coax out of a large language model (again, at least for now).

The mention of law school education raises a second concern, one that legal educators are wrestling with. We have extensive experience supporting students as they develop the skills and judgment necessary to thrive as written legal communicators. Trial and error, and the accompanying struggle and frustration, have long played an important role in cultivating students' understanding of and ability to perform core lawyering tasks; trial and error has helped law students develop problem-solving skills and grit as well. Now, to the extent students begin turning to generative AI before they develop key knowledge and skills, they risk entering the profession with a shaky foundation, less practice-ready than peers who embraced the "desirable difficulty"[2] of the traditional law school curriculum.

Third, the financial cost of this new technology presents additional challenges. The law-specific tools with safeguards appropriate for our profession are expensive. Not every lawyer can count on having access: Some firms will not invest in the technology, and some clients or courts will disapprove of its use. Therefore, all lawyers, including those just joining the discipline, must be able to generate clear prose without the assistance of a large language model.

Finally, consider the risk of skill atrophy. You've no doubt heard the adage "use it or lose it," which can apply to our muscles and minds alike. You want to retain your outstanding legal writing skills — your unique flare for expressing ideas vividly and making complex, dry legal arguments seem both obvious and important? The more you defer to technology instead of crafting your own prose, the harder it may become to call upon those once-sharp skills when you need them.

Skill-and-Style-Saving Strategies

We are still in the early days of writing with the assistance of large language models; our discipline will build consensus around best practices in time. Until then, I tentatively offer the following strategies, which aim to support skill development by

emerging professionals — and guard against skill atrophy among experienced practitioners—in the generative AI era. They also aim to help professional writers appreciate and preserve their own voice and style.

1. Consider your order of operations. Sometimes large language models offer relevant insights, but often they provide inaccurate facts, outdated law or otherwise flawed analysis. To help you catch those problems, think through your analysis before you prompt a large language model to generate a draft. That preparatory work will help make your prompts more specific and the resulting draft more useful. It will also allow you to maintain the healthy skepticism that is necessary to evaluate (and adjust) the automatically generated output.

2. Be selective.

- *What document?* A large language model that helps you complete some projects may slow your progress on others. For example, you might automate the creation of a billing letter, a common contract or a simple motion, while still drafting from scratch when you tackle an appellate brief raising an issue of first impression.

- *Which tool?* Each generative AI product has its strengths and weaknesses. That is why Lexis+AI is built on the back of a variety of general-use large language models.[3] Pick one that is reliable for the type of work you seek to automate.

- *What task?* Writing occurs in a series of stages; use generative AI to help with just one. For example, a first draft provided by a large language model can be quite helpful to those who struggle with a blank screen. Good writing requires rewriting, but you cannot revise and polish until you have prose on the page. Large language models can also be helpful at the revising and polishing stage. But to ensure that you are producing the best possible written work, resist the temptation to have one large language model generate a document and another revise it. Use generative AI to create a first draft, or to revise aspects of your own first draft, but not both.

3. Add depth and detail. Large language models, so far, tend to be superficial.

They produce generic and sometimes repetitive prose, skimming the surface when thorough analysis requires a deep dive. Many arguments get more precise and persuasive with detailed analogical reasoning; I've not yet observed the analytical depth that is the hallmark of outstanding written legal analysis in AI-generated prose.

4. Be judgmental, not deferential. Apply the same healthy skepticism to style that you already apply to content. Large language models generate prose that is organized, fluid and mechanically perfect, enticing users to believe that the work product is finished despite extensive opportunities for improvement. After you ensure that the material is accurate and thorough, ask more questions: Is there a more compelling way to tell this story? Could I make this language tighter or brighter? Would any word, punctuation, or sentence structure adjustments help the reader more readily understand this point or agree with this argument? Is the tone right for my audience? Does this sound like something I wrote?

Choose Hope

Legal writers have powerful generative AI tools to help them efficiently serve their clients' needs. Time will tell how dramatically these tools affect lawyers' writing processes and shape their written products. Based on my experimentation with large language models, I have been amazed and concerned. Now I choose to be hopeful — optimistic that our profession will harness the power of large language models to create sharper work with increased efficiency. ∎

Professor Rebekah Hanley teaches legal writing at the University of Oregon School of Law. As Oregon Law's current Galen Scholar in Legal Writing, Professor Hanley is studying generative AI and its implications for law school teaching and the practice of law. She thanks her colleague Professor Suzanne Rowe and second-year Oregon Law students Aazaad Burn and Thomas Grossman for their suggestions on this column.

ENDNOTES

1. One study showed that when students rely on generative AI, grades at the bottom of the class get a boost, while the grades of top performers slide in the opposite direction — falling down, toward the average. Jonathan H. Choi & Daniel Schwarcz, *AI Assistance in Legal Analysis: An Empirical Study* (Aug. 13, 2023), 73 J. Legal Educ. (forthcoming 2024), available at *https://ssrn.com/abstract=4539836* (access to "GPT-4 substantially improved the scores of students at the bottom of the class and negatively impacted the scores of students at the top of the class").

2. This term, coined by cognitive psychologist Robert Bjork in 1994, refers to the effort that will improve long-term performance on certain tasks.

3. "LexisNexis deploys ethical, powerful generative AI solutions with a flexible, multi-model approach that prioritizes using the best model for each individual legal use case. This approach includes working with large language models like Anthropic's Claude 2, hosted on Amazon Bedrock from Amazon Web Services (AWS), and OpenAI's GPT-4 and ChatGPT, hosted on Microsoft Azure." LexisNexis, *LexisNexis Launches Lexis+ AI, a Generative AI Solution with Linked Hallucination-Free Legal Citations*, (Oct. 25, 2023), *https://www.lexisnexis.com/community/pressroom/b/news/posts/lexisnexis-launches-lexis-ai-a-generative-ai-solution-with-hallucination-free-linked-legal-citations*.

Be an Author

The *Bulletin* welcomes quality manuscripts from Oregon State Bar members for publication.

We publish articles on a wide variety of subjects and favor such topics as access to justice, legal funding, judicial independence, diversity in the profession, professionalism and future trends. We also publish columns on ethics, practice tips (in specific areas of law), law practice management and legal history, as well as essays on law and life.

The editorial staff welcomes inquiries and is happy to discuss requirements for publication. If you have a manuscript, suggestion or idea, contact editor Michael Austin at (503) 431-6340. He can also be reached by email at *editor@osbar.org*.

Understanding
Generative
AI

A Primer for the Next Step in AI and Its Legal Applications

By Justice Brooks

O ver the past last 25 years, the legal profession has seen a significant evolution in the way technology is utilized to streamline processes and enhance efficiency. Fax machines, snail mail and desk phones are nearly extinct. They have been replaced by email, SMS communications and other digital alternatives. Traditional artificial intelligence (e.g. technology-assisted document reviews, search engine recommendations) is now commonplace. The next step in that evolution is generative artificial intelligence. This article aims to demystify generative AI and explore its applications within the legal sphere.

What Is Generative AI?

Generative AI is a branch of artificial intelligence that involves machines creating content — whether it be text, images, music or other forms of data — that is meant to be indistinguishable from content created by humans.

Unlike traditional AI, which operates based on predefined rules and patterns, generative AI employs techniques such as deep learning and neural networks to generate new, original content autonomously.

How Does Generative AI Work?

At the heart of generative AI are neural networks, computational models inspired by the structure and function of the human brain. These networks are trained on vast amounts of data, learning to recognize patterns and relationships within the data. Generative AI utilizes two main types of neural networks:

Autoencoders: These networks learn to compress input data into a lower-dimensional representation and then reconstruct the original data from this representation. Autoencoders are commonly used in tasks such as image and text generation.

Generative Adversarial Networks (GANs): GANs consist of two neural networks — a generator and a discriminator — that are trained simultaneously. The generator generates synthetic data, while the discriminator evaluates whether the data is real or generated. Through adversarial training, both networks improve over time, leading to the generation of increasingly realistic content.

What Is Machine Learning?

Machine learning plays a pivotal role in the development and implementation of generative AI within the legal field. As generative AI systems rely on large datasets to learn patterns and generate new content, machine learning techniques such as supervised, unsupervised and reinforcement learning are for training and optimizing these systems.

Supervised learning enables generative AI models to understand the structure and semantics of legal texts by learning from labeled datasets of case law, statutes and legal documents.

Unsupervised learning techniques help identify latent patterns and themes within legal data, facilitating the generation of diverse and contextually relevant content.

Reinforcement learning allows generative AI systems to adapt and improve over time based on feedback received from users and real-world interactions, enhancing the quality and accuracy of generated legal content.

How Is Generative AI Different From a Google Search?

It's important to distinguish generative AI from search algorithms like Google, as they serve different purposes and utilize distinct methodologies. While Google focuses on retrieving and ranking existing information based on relevance to user queries, generative AI goes a step further by creating new content autonomously.

Search algorithms operate within the confines of indexed data available on the web, relying on keyword matching, page authority and user engagement metrics to deliver relevant results. In contrast, generative AI generates original content by learning patterns and structures from large datasets.

While both technologies aim to assist users in accessing information, generative AI offers the capability to generate content tailored to specific needs or preferences, making it particularly useful in tasks such as content creation, natural language generation and creative expression. Conversely, search algorithms excel in retrieving existing information efficiently from vast repositories of data, providing users with access to a wealth of knowledge available on the internet.

Applications of Generative AI in the Legal Profession

The following are several ways generative AI can be used in the legal profession.

Document Automation: Generative AI can automate the drafting of legal documents such as contracts, agreements and pleadings. Generative AI algorithms can generate customized documents tailored to specific legal needs.

Legal Research: Generative AI can assist in this process by analyzing vast amounts of legal texts and generating summaries, case briefs or even predictive analyses of potential legal outcomes based on precedent.

Contract Review and Analysis: Generative AI can review contracts for key clauses, obligations and potential risks. It can also analyze contracts and identify relevant clauses, anomalies or discrepancies, thereby expediting the contract review process and reducing the likelihood of oversights.

Predictive Analytics: Generative AI can be utilized for predictive analytics in various legal contexts, such as predicting case outcomes, assessing the likelihood of settlement or identifying potential legal risks for businesses.

Legal Writing: By analyzing vast repositories of legal documents, including case law, statutes and precedent, generative AI can identify relevant arguments, legal principles and citations to support a particular case or position. Generative AI can then generate legal briefs tailored to the specific needs and preferences of legal practitioners. Generative AI can also assist in ensuring consistency and coherence across multiple drafts of legal briefs, reducing the likelihood of errors or omissions. *See The Legal Writer column on page 17 for more information on this topic.*

Ethical and Legal Considerations

While generative AI offers numerous benefits for the legal profession, its use raises important ethical and legal considerations. (*See the Bar Counsel column on page 9 for detailed analysis on this topic*)

Bias and Fairness: Generative AI algorithms are susceptible to biases present in the training data, which can perpetuate and exacerbate existing disparities in the legal system. It is essential for attorneys to be mindful of these biases and take steps to mitigate them when utilizing generative AI tools.

Confidentiality: Documents may contain confidential, proprietary or sensitive information that may only be disclosed in limited circumstances. Attorneys must ensure they are implementing and utilizing appropriate policies, protocols and security measures to protect the confidential information.

Professional Responsibility: Attorneys have a duty to provide competent representation to their clients. While generative AI can enhance efficiency, attorneys must exercise professional judgment and oversee the output generated by AI tools to ensure accuracy and quality.

Hallucinations by Generative AI

While generative AI offers tremendous potential in assisting with legal work, it's essential to recognize the potential for "hallucinations" or the generation of inaccurate or misleading content. Like people, generative AI models are not immune to errors, and there is a risk of producing outputs that may contain inaccuracies, logical inconsistencies or misinterpretations of legal principles. These hallucinations can arise due to biases in the training data, limitations in the AI model's understanding of legal concepts or unforeseen complexities in legal language and context.

Legal professionals must exercise caution and critical judgment when utilizing generative AI. It's crucial to verify the accuracy and relevance of AI-generated content through careful review and analysis by experienced practitioners. Additionally, legal practitioners should continuously monitor and refine AI models to minimize the occurrence of hallucinations and ensure that the generated content aligns with their ethical and professional standards. By approaching generative AI tools with vigilance and skepticism, legal professionals can harness their capabilities effectively while mitigating the risks associated with inaccuracies or misinterpretations in legal work.

Put simply — verify, verify, verify.

Unanticipated Consequences

While generative AI holds promise for enhancing efficiency and productivity in the legal profession, its widespread adoption may have potential implications for the training of new lawyers. As generative AI tools automate routine tasks such as legal research, document drafting and contract review, there is a concern that reliance on AI technology could diminish opportunities for new lawyers to develop essential skills and expertise. Tasks that were traditionally part of a junior lawyer's training, such as conducting extensive legal research or drafting pleadings from scratch, will become increasingly automated, limiting hands-on experience and exposure to the intricacies of legal practice. Moreover, as Generative AI systems evolve to perform more sophisticated tasks, there is a risk that junior lawyers may become overly reliant on AI-generated content, potentially compromising their ability to critically analyze legal issues, exercise judgment and communicate effectively with clients and colleagues.

To address these concerns, legal educators and practitioners must adapt training programs to incorporate instruction on how to effectively utilize AI tools while emphasizing the importance of foundational legal skills, critical thinking and professional judgment. By striking a balance between leveraging AI technology and preserving the essential elements of legal training, the legal profession can ensure that new lawyers are equipped with the competencies needed to navigate the complexities of legal practice going forward.

Conclusion

Generative AI represents a transformative technology with the potential to revolutionize various aspects of legal practice. By harnessing the power of machine learning and neural networks, attorneys can streamline workflows, improve decision-making and deliver enhanced legal services to clients. However, the adoption of generative AI also necessitates careful consideration of ethical, legal and professional implications. As the legal profession continues to evolve in the digital age, embracing innovative technologies like generative AI can empower attorneys to meet the demands of a rapidly changing legal landscape while upholding ethical standards and ensuring access to justice for all. ∎

Justice Brooks is a principal at Foster Garvey in Portland. He recently presented on this topic at the American Bar Association Midyear Meeting, as well as the 2024 Litigation Section Institute and Retreat, cosponsored by the Oregon State Bar Litigation Section, which was held in Gleneden Beach.

Reality *is More* Complex

AI's Efficiency Isn't a Direct Line to Increased Access to Justice

— By Shannon Gormley —

For over a decade, John Grant has worked to make lawyers more efficient. A tech worker-turned-lawyer, Grant decided to combine his knowledge of the two fields when he opened his consulting firm, Agile Attorney Consulting, in 2014. Since then, he's used tech-industry software and methodology to help lawyers streamline their practices.

"I have this almost Pollyanna-ish belief that if I can teach more lawyers and law firms to make their practices more efficient, then they will be empowered to serve more people in an affordable way," says Grant. "If we can get more law firms serving more people, then that makes legal help more accessible to the public."

Now, with recent, rapid advances in artificial intelligence, Grant's mission to increase access to justice by improving lawyer efficiency seems prophetic. Publications from *Forbes* to *The New York Times* and Reuters have all proclaimed that AI will be a boon for lawyer productivity, which in turn will make legal services more accessible. One American Bar Association article, which predicts that the "obvious" time-saving benefits of AI will translate into cost-savings for clients, claims that AI is "here to save us from ourselves."[1]

The advent of advanced large language models like OpenAI and ChatGPT comes at a time when the need for affordable, widely available legal services is increasingly acute. Ninety-two percent of low-income Americans do not receive any or enough legal help.[2] Here in Oregon, 85% of attorneys are located in the upper Willamette Valley, and four counties have no private practitioners at all.[3]

Could AI help? The logic used in many articles about the new technology seems simple enough: AI will make lawyers more efficient, and more efficient lawyers will be able to help more people. But the reality is much more complex.

"I haven't seen any data to firmly support that conclusion," says Rebekah Hanley, Oregon Law's Galen Scholar in Legal Writing who is spending the year researching AI. "I think there's some conjecture."

Though Grant began working at the intersection of lawyer efficiency and access to justice long before the current AI boom, he describes himself as a "cautious optimist" about the new technology. But he doesn't see the potential benefits as inevitabilities.

"You can't just unplug one tool, plug in another tool and have the whole system magically get better," he says. "You have to reapproach the design of the entire system"

Access and Equity

Growing up, Shiwanni Johnson had what she describes as a "distant" relationship to technology. She didn't have a computer in her family home, and got her first cell phone later than many of her peers. Then, she went to law school. It didn't take long for Johnson to see technology and law as deeply intertwined. During her second year at University of Oregon School of Law, she founded the Technology Law Club, aimed at exploring how tech — from e-discovery to spreadsheets — can make lawyers better at their jobs and improve access to justice.

Now a third-year associate, Johnson chairs OSB's Technology Law Section. She regards the hype around large language models with skepticism.

"I think AI is undoubtedly a tool that will lead to more lawyer efficiency," says Johnson. "But just because a lawyer is more efficient doesn't mean that they're providing access to justice."

The assertion that AI will make lawyers more efficient and thus increase access to justice is sometimes treated as if it's a single byte of information. But really, it's two, and both rely on assumptions. And as Johnson points out, access doesn't always mean equity.

"For me, access to justice means using the law to get people what they are entitled to under the law. Not only that, but also using the law in a fair and equitable manner," says Johnson. "Artificial intelligence is so dynamic and so new that I don't know if it necessarily serves either of those purposes."

AI has already been used to perpetuate existing racial bias in the legal field.[4] Generative models rehash the biases of the information they are fed, and the history of the legal field is replete with biases and inequitable outcomes.

The Oregon Judicial Department has been closely monitoring AI's recent developments, its potential to make legal services more accessible and its possible pitfalls. "Artificial intelligence is not something to be ignored," says Chief Justice Meagan A. Flynn. "I've heard multiple presentations about AI over the past year, including one at the Conference of Chief Justices. This conversation will also be happening in the coming months and years with our judges and staff at OJD. AI has exciting potential to help the courts, but it's also something we must approach cautiously and with solid governance in place."

In addition to the issue of algorithm bias, Flynn points to "hallucinations" — fake information that a large language model presents as factual — as a concern. In order for AI to successfully facilitate access to justice, Flynn believes that human oversight will be key.

"If bias or misinformation is built into AI products from the beginning, then the AI will reproduce that misinformation," says Flynn. "We are responsible for protecting the integrity of the court process, checking that the information we rely on is accurate and ensuring that the decisions we make are based on case law and evidence that really exists, not something that is only computer generated."

Additionally, many of the attorneys the *Bulletin* spoke to were concerned about who will and who won't have access to top-of-the-line AI technology. For the moment, it's mostly big law firms with big budgets. At a recent presentation, Hanley heard from several public interest attorneys who were anxious that AI would only put them at a greater disadvantage against better-resourced private parties.

It's a trend that Grant has noticed, too. "If you're using the public version of GPT, that pales in comparison with the bespoke systems that AMLaw 100 firms are all already building," he says. "It has the risk of making the gap between the haves and the have nots even wider."

Angela Laidlaw, a family law attorney based in Oregon City and a member of the State Family Law Advisory Committee's Futures Subcommittee, recently helped draft AI policy guidelines for the Oregon Judicial Department (at the time of the *Bulletin*'s interview, the guidelines were in committee review, and Laidlaw declined to discuss the specifics or speak on behalf of the committee). She points out that AI could also put pro se litigants at an additional disadvantage if they are up against a party with access to representation and advanced AI.

On a scale of 1 to 10, with 10 being the most optimistic about AI, Laidlaw would place herself at a 3 or 4. She is particularly concerned about the confidence with which generative AI presents false

information. As a test, Laidlaw has asked ChatGPT to write essays based on untrue information. Each time, the algorithm has taken the prompt as if it was completely factual and authoritatively replied with an essay of entirely false information. She's noticed that, for some reason, ChatGPT seems to believe that interrogatories are allowed in Oregon when they are not. Laidlaw describes it as an "echo chamber of false information."

"AI is only as good as the information we put into it," she says. "I have a hard time imagining in my lifetime that we could use AI without having a human interpret some of it or check it."

Laidlaw does use ChatGPT in her own practice, mostly for summarizing information and what she refers to as a "tone check" – asking the large language model to make sure what she's written comes across the way she intended it to. It has made her practice more efficient and has allowed her to help more clients.

"It's still on a very small scale," she says. Even so, "I really do think we can figure out ways to do more stuff in less time, which is better all-around."

Large Language Models, Human Scale-Efficiency

Last February, Grant sat in a meeting room at the Multnomah County Courthouse and met with Oregonians facing eviction. He was volunteering with the Commons Law Center's eviction defense clinic, which sometimes serves up to a third of the court's morning docket. The access to justice gap is particularly prevalent in landlord-tenant law. Last year, only 8% of Oregon tenants and 42% of landlords had legal representation in eviction proceedings.[5]

The eviction defense clinic uses practice management software to help with basic procedures. But on the day that Grant volunteered, he was bemused to find a much more analog technology in use — a large paper chart used to track cases, with columns of Post-It notes divided by blue painters tape. The visual organization tool is known as a Kanban board, a tool that Grant often uses with his consulting clients. It helps him and his clients identify where work gets stuck due to system bottlenecks, something that AI, if used poorly, could make worse.

Some AI enthusiasts have suggested that software could help with client intake or new matter filing. But Grant worries that if the influx of filings overwhelms the human workers involved, work will just become more backed up.

"What most people will self-identify as a 'I need to be more productive in my work' problem is actually a 'the system needs to have a more productive in the flow of work' problem," he says. The way Grant sees it, expecting lawyers to work faster through too great of a workload is like expecting a car to drive full-speed down a backed-up highway: "The car is capable of going 100 mph. It's the traffic on the freeway that is keeping it from going that fast."

The distinction between "efficiency" and "productivity" is important, but often ends up muddled. At face value, increasing efficiency is good for lawyers. It means that tasks take less time and are less tedious. Productivity, however, is a measure of how much work you can churn out — and most lawyers are already overworked. A recent Bloomberg Law study found that while work hours have gone down, the average rate of attorney burnout has not.[6]

To increase efficiency without increasing burnout, Hanley proposes "selective deployment" of AI. Given all of the review and oversight necessary, trying to use a large language model to spit out a final product for a client or judge may require more work, not less, and may also lead a lawyer to generate lower quality work. Instead, Hanley suggests using AI to help with small pieces of a project, especially those that you find particularly challenging or tedious. It can even be used to suggest new phrasings if you're feeling stuck or write segments of a first draft to bypass the pressure of a blank page.

While working with Commons Law Center's eviction defense clinic, Grant did find a use for ChatGPT when he wanted to check the compliance of some proposed changes to Multnomah County court procedures. Instead of manually searching through a thicket of Uniform Trial Court Rules, supplementary local rules, and Chief Justice and Presiding Judge Orders, he created a custom GPT with the relevant sets of documents, put guard rails around its ability to hallucinate and then asked the model specific questions about the proposed changes.

"It was great," says Grant. "It really helped me focus on a small handful of provisions that were relevant to the questions that I had."

For Grant, such improvements don't change that organizations and individuals need to be realistic about how much work they can actually take on. He points out that even with the help of efficiency tools, the eviction defense clinic caps the number of people it can help each day according to how many lawyers and paralegals are physically on-site to help.

"We're not going to be able to use the same systems and workflows that are built around human capabilities for a single person doing legal work, plug in AI and have it magically work at 1.5 or 2 times the human's productivity level," he says. Grant believes that, in a literal sense, keeping law at a human scale is good for clients, too: "When people hire a lawyer, they want to hire a human. They're buying some element of, 'there's this other person in this with me.' Meeting people's social-emotional needs is an important component of meeting their legal ones."

Beyond the Black Box

Sitting in her office on a sunny, unseasonably warm day at the end of winter, Johnson remains as insatiably curious about law and technology as she was when she started law school, despite her acknowledged hesitancy about embracing generative AI. Johnson believes that if lawyers are going to use predictive algorithms, as many of them already are, they should at least know what they're dealing with.

Last fall, the Technology Law Section hosted a CLE aimed at just that. The day-long event functioned as a crash course in AI in the legal field, from ethics to practical applications. A talk from an Oregon State University robotics professor titled "Everything You Ever Wanted to Know About Large Language Models (But Were Afraid to Ask)" dug into how predictive algorithms work, when they fail and "the real risks associated with unrealistic expectations" of their capabilities.

If legal technology is going to help the public, though, Johnson believes that the public needs to be a part of the conversation. "I think it will require not only people in the tech industry — the people creating these AI tools — working directly with lawyers, but

also working with people directly in the community to whom they'll be providing these services," she says. "I also think it requires transparency to people using these services."

Hanley, too, believes that more transparency is necessary moving forward. The inner workings of deep-learning AI are invisible to the user and in many cases, not fully understood even by its creators. Generative AI can still produce unexpected, unwanted outcomes, even as the technology rapidly improves. "There's just more questions than there are answers," says Hanley. "It seems like things could change dramatically tomorrow."

In the meantime, finding the best ways for AI to increase access to justice might require rearranging the order of operations in a much simpler equation. For Johnson, if the goal is to help the public, then benefits to the public have to be discussed first, not as a downstream benefit of efficiency.

"The advances in technology should improve the system from a non-lawyer's perspective, from people who are trying to access the legal system," she says. "That makes it more fair and equitable." ∎

Shannon Gormley is the associate editor of the Bulletin.

ENDNOTES

1. https://www.americanbar.org/news/abanews/publications/youraba/2017/september-2017/7-ways-artificial-intelligence-can-benefit-your-law-firm/
2. https://justicegap.lsc.gov/#:~:text=Low%2Dincome%20Americans%20did%20not,restraining%20order%20against%20an%20abuser
3. https://www.osbar.org/bulletin/issues/2022/2022November/index.html
4. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing
5. https://www.evictedinoregon.com/disparity-in-legal-representation
6. https://news.bloomberglaw.com/bloomberg-law-analysis/analysis-female-lawyers-report-more-stress-burnout-than-males

The**Next**
Generation

Law Schools Navigate AI Amid Rapidly Changing Landscape

— By Melody Finnemore —

Rebekah Hanley, a professor who teaches in the Legal Research and Writing Program at the University of Oregon School of Law, was prepared to speak about new developments in generative artificial intelligence at a mid-January faculty meeting when the presentation was postponed because of the snow and ice storm that wreaked havoc across the state.

Hanley, Oregon Law's 2023-24 Galen Scholar in Legal Writing, is spending the year studying generative AI and its impact on law school teaching and the practice of law. She frequently speaks about AI in legal education and says she was hesitant to update her materials for the rescheduled presentation until right before it actually happened, because AI is changing and advancing so rapidly.

"It highlights one layer of challenge we're all grappling with," she says. "There is no point at which we can all pause, assess, make a plan and commit to sticking to it because the situation is so dynamic."

Like Hanley, professors at Willamette University's College of Law and Lewis & Clark Law School are navigating the use of AI in legal education in real time. From monitoring how students use it for assignments, projects and exams to teaching them how to incorporate AI responsibly and ethically as they begin practicing, Oregon's trio of law schools joins others across the country in adapting to this light-speed evolution in the legal profession.

The American Bar Association last summer established the ABA Task Force on Law and Artificial Intelligence to examine the impact of AI on law practice and the ethical implications for lawyers. AI in legal education is one of five specific areas the task force is charged with exploring.

The task force seeks to inform the legal community 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 lawyer and judges.

In December, LexisNexis announced that it would make its generative AI platform, Lexis + AI, available to 100,000 second- and third-year law and masters of law students at ABA-accredited law schools during the spring semester. Faculty training webinars were to be scheduled during the rollout.

The rollout followed a test run last fall in which 450 law school librarians, legal research and writing instructors, and legal technology professors were given access to the platform, which it said supports "conversational search, intelligent legal drafting, insightful summarization and document analysis."

"In my classroom I have started experimenting with Lexis + AI. That requires me to be vulnerable in front of my students because I'm learning along with them," Hanley says, adding it is important to educate students about the tools now so they can gain experience and understand the pros and cons of using them.

Photo on Opposite page: Policies at Willamette University College of Law are constantly changing to the evolving nature of AI.

"Some students are very wary, nervous and unenthusiastic about these tools, while others are hungry for the opportunity to learn about them and how to use them well," she says. "It feels tricky to continue to tell students not to use generative artificial intelligence when we see the profession moving in that direction."

Hanley notes that AI can support student learning and "be a partner at the starting point" by helping to generate ideas, overcome writer's block and facilitate research. Used later in the writing process, AI can assist with proofreading to polish writing. However, she adds, students must also understand their ethical responsibility to supervise the tool and verify all information that goes into their legal work product, a responsibility that will never change.

She led a class exercise in which her students prompted Lexis + AI to draft arguments and then critiqued the results to evaluate accuracy and thoroughness. They reviewed the depth and detail of the arguments and found that they were somewhat shallow, repetitive and superficial.

"I'm trying to teach students to think very critically about the information the generative AI spits out, and the students are learning that its ability to create meaningful, analytical legal prose is lacking in many respects," she says. "I think that was an important lesson for the students. These tools will maybe save them some time, but the tools don't compare to the knowledge, skills and judgement they have developed so far in law school and their own ability to explain things to their clients."

Hanley says that faculty at Oregon Law have enjoyed the academic freedom to proceed with AI as they see fit for their individual teaching decisions, but the default position for students is to avoid using AI unless faculty says otherwise. She describes her own teaching strategy as a "task-by-task, day-by-day, assignment-by-assignment, tool-by-tool approach." With the academic integrity section of class syllabi as a guide, in some courses students may be permitted — or even encouraged — to use AI for any purpose other than producing graded written work. Still, AI detection tools are not entirely reliable, which makes it problematic to determine when generative AI has been used for work that is graded, Hanley says.

The Next Iteration of Legal Tools

David Friedman, Willamette Law School professor and associate dean for strategic initiatives, likens talking about AI in law as "a little like an arms race" and that it is difficult to really know what it means for the profession.

"This information could be quaint and archaic a year from now, and a year ago we would have been surprised that we are even talking about it," he says.

One thing Friedman is sure about is that human beings will always play the most important role in the final product. For example, two sides may be trying to resolve a dispute and both may have the same AI available, but their unique human insight and experience is needed to gain the edge in the argument.

Friedman points out that electronic discovery changed the structure of litigation and how staffing is managed, and AI is another version of that evolution in that it can provide economical shortcuts that might help by making legal services more accessible, efficient and cost-effective. Still, clients will need their lawyer's advice.

"Part of this is you have to embrace it, there is no point in being afraid of it. If you're afraid of something, you're not looking at the upside of how it can help," he says. "At some point using AI will be part of what we expect when we think of lawyer competence, just the same way we assume they use the internet or computers. This is kind of the next iteration of that, it's just faster."

Willamette is now offering a seminar on law and AI, which Friedman says will continually change as the instructor strives to stay on top of the latest advances. He says faculty members expect to learn as much as students do through the seminar.

In terms of monitoring how students use AI for their work, Friedman calls it still very much a work in progress, in part because of how rapidly AI has developed and also how fast it changes.

"We realized very quickly that we had to be clear with ourselves and our students about what we are academically assessing, what we are teaching and in which context," he says. "I think we're kind of in an interim phase across education, in general, in figuring out what this means."

As an example, students have been afforded a more open-source environment for exams, especially in remote courses held during the COVID pandemic. Now that may change drastically so faculty can better measure students' abilities and ensure everyone is on a level playing field without the use of AI.

For a research paper or project, some faculty might permit AI to be used for brainstorming potential topics and research but not for the actual writing. "We had to make it very clear, first of all to instructors, about what these tools could do and how you have to manage that. Be clear with students about rules and what AI can be used for," Friedman says.

The experiential component of law school has become more important as Oregon expands its types of examination for licensure. These include more experiential opportunities such as simulation courses and clinics in which AI could improve productivity, he adds.

Well-publicized dangers do exist, Friedman cautions, referring to attorneys in New York who were sanctioned by a federal judge for submitting legal briefs written by ChatGPT that included fake cases and citations. He describes those incidents as attorneys using AI recklessly as shortcuts.

"Lawyers have an obligation to figure out how to harness AI for clients without hurting them," Friedman says. "For the foreseeable future, you're going to have to use your human eyes to make sure everything is accurate."

Clear Rules About Use Are Crucial

John T. Parry, associate dean of faculty and Edward Brunet Professor of Law at Lewis & Clark Law School, says a faculty member there also has developed a class on AI and the law. "One thing I think is interesting is that it's not just the law and policy of AI, but the students will be using AI in the class and some of the same things they learn in the lawyering class so it will be both policy and hands-on," he says.

A couple of robust faculty discussions led to the question of whether to draft an institutional policy about student use of AI. The difficulty is that such a policy either has to be so general that it

Lewis & Clark Law School has developed a class on AI in which students learn about law and policy, but also will use it in a hands-on way.

wouldn't be very useful or so specific that it would be dated in a few weeks, Parry says.

"Students are free to use AI, for example, as a study tool to outline a class or a case and that can be really useful. There is also a lot of AI that has already been out there. Lexis and Westlaw have added AI components and we don't want to outlaw that because they will need to use that as lawyers," he says.

Lewis & Clark Law School has a general understanding with its faculty that it is their responsibility to inform students that AI is only to be used with a professor's permission to ensure there are no code violations. For exams, the law school can disable internet access on student laptops and professors have to opt in if they want their students to have internet access during an exam.

Photo Above: At the University of Oregon School of Law, the default position for students is to avoid using AI unless faculty says otherwise. The Knight Law Center on campus is pictured.

"We want faculty to be able to work with students if they want to on AI issues or rely solely on their own knowledge," Parry says. "Students also realize it's one thing to use it as a tool and they have seen the horror stories of lawyers turning things in with hallucinations included."

Robert Truman, associate dean and director of the law school's Paul L. Boley Law Library, says generative AI tools and concepts for research are introduced to students during their first year and that education continues into the advanced courses. The law school is exploring the potential for AI to be particularly useful for contract drafting classes.

When Truman polled professors of writing classes, they reported asking students who were creating briefs to run their research, writing and analysis through AI tools such as ChatGPT and perform a comparative analysis about how well these tools do — or don't do.

"Like everyone else, we were thrown into this about a year ago and it's all changing so fast, but the overall understanding is that AI tools of all sorts are being rolled into just about every legal tool and practice," he says. "We know it is useful and our goal is to teach it in context so students use it to best serve their clients when they start practicing." ∎

Melody Finnemore is a Portland-area freelance writer. She recently profiled Kenna West of Polk County in the December 2023 Bulletin.

TECHNOLOGY

Having an AI Workplace Policy Is Smart Business

In the Shadows

By Sharon D. Nelson, John W. Simek and Michael C. Maschke

Remember Shadow IT? Say hello to Shadow AI.

There were plenty of articles written about Shadow IT — defined by Cisco as "The use of IT-related hardware or software by a department or individual without the knowledge of the IT or security group within the organization." Shadow IT included cloud services, software and hardware.

Welcome to the sudden rise of Shadow AI. Its use, like that of Shadow IT, is often unknown to a law firm's IT or security group. As lawyers gravitated with haste to using generative AI in 2023, the conversation at law firms rapidly turned to controlling the use of Shadow AI.

Do you have any idea how many of your firm employees are using AI? The likely answer is no. We've all been so busy exploring what AI can do in our practices that only the largest of law firms are likely to have thought about AI policies, much less tracking the actual use of AI in their firms.

While law firms and companies don't like reporting on shadow IT problems, in 2023 Samsung issued a temporary ban forbidding any unauthorized AI applications after an internal data leak. We are sure similar edicts have been issued elsewhere, but that's the kind of subject that companies and law firms prefer to keep quiet.

Hard to Track

AI is everywhere, but it's not always visible. We forget that AI is embedded in videoconferencing programs, in many legal research programs, in our e-discovery software, in the browsers we use to search for information, in our smartphones — and the list goes on and on.

Sometimes it is more apparent that we are using AI — we understand that we are using it when working with ChatGPT, Harvey (which some major law firms use), Bard, Bing Chat, etc.

Has your law firm authorized you to utilize an AI? There's the rub. In general, employers are often unaware of what generative AI is being used by their employees. And the employees like their AI — you might send a survey asking employees if they use AI and they might well say no even if the correct answer is yes — they don't want to get in trouble, but they have no intention of giving their beloved AIs up. AIs have become addictive.

Do You Need an AI Usage Policy?

Absolutely. At least you need to document what is and is not allowed. You may choose to identify generative AIs, which your lawyers and staff may utilize. However, you will certainly want to underscore certain things. Do they need to tell the client if they are using AI? Most ethicists would say yes. Do they need to get permission for that use? If time is saved, is billing reduced? Do you ensure that no confidential data is given to the AI, either placed in its database or used for training?

In the end, your policy will constitute a set of guidelines and regulations which make sure that the law firm's use of AI is ethical and responsible. The policy should address any cybersecurity issues, data privacy laws, federal/state regulations, ethical considerations, etc.

Finally, it should be made clear that no unauthorized AI may be utilized. This may reduce the amount of Shadow AI at your firm, but never imagine that a mere policy will put an end to the use of shadow AI by rogue employees.

Another Area for Employee Training

AI training is an industry these days — and the dangers of Shadow AI can certainly be addressed in a training session. We also suggest that Shadow AI be addressed in employee cybersecurity awareness training.

Many cyberinsurance companies require such training, so it is an additional "guardrail" to include a segment on the dangers of Shadow AI — indeed, on the cybersecurity dangers that may come with authorized AIs as well.

While most lawyers know they shouldn't give client information to an AI, they may not realize how dangerous it is to give information about the firm itself to an AI. As an example, don't construct a prompt along the lines of "How do you configure a TZ500 SonicWall firewall to allow FTP traffic?" Anyone who has access to that data (whether authorized or unauthorized) may use such information in a cyberattack. You'll certainly want to underscore that danger in the training.

How Can You Monitor the Use of Shadow AI?

That's the hardest question. There are solutions which provide full visibility into what applications are running and who is using them. Basically, you can scan for installed

software and/or devices that are used to access your data and environment. What if you do not allow personal devices to access firm information? It's a simple matter to check the device "partnerships" in your Microsoft 365 account to see if any of the phones appear to be non-firm-issued devices used to synchronize a user's mailbox.

There are also software monitoring tools to capture a user's activity even if a browser is used to access an AI environment. To be totally transparent, make sure you notify employees that you may be monitoring their activity. Some states even require that such notice be given to employees in a very prominent way.

Final Words

Embracing AI unreservedly is tempting. But go slow and be careful. Expect rogue behavior and have a plan to deal with it! ∎

Sharon D. Nelson *is a practicing attorney and the president of Sensei Enterprises, Inc. She is a past president of the Virginia State Bar, the Fairfax Bar Association and the Fairfax Law Foundation. She is a co-author of 18 books published by the ABA.* snelson@senseient.com

John W. Simek *is vice president of Sensei Enterprises, Inc. He is a Certified Information Systems Security Professional (CISSP), Certified Ethical Hacker (CEH) and a nationally known expert in the area of digital forensics. He and Sharon provide legal technology, cybersecurity and digital forensics services from their Fairfax, Virginia firm.* jsimek@senseient.com

Michael C. Maschke *is the CEO/Director of Cybersecurity and Digital Forensics of Sensei Enterprises, Inc. He is an EnCase Certified Examiner, a Certified Computer Examiner (CCE #744) a Certified Ethical Hacker and an AccessData Certified Examiner. He is also a Certified Information Systems Security Professional.* mmaschke@senseient.com