

THE DUTY OF EFFICIENCY & GENERATIVE AI PEDAGOGY

Megan E. Boyd* & Brian L. Frye**

ABSTRACT

The emergence of generative artificial intelligence (generative AI) promises to transform the legal profession by improving efficiency, yet its adoption has been hindered by several obstacles. Courts have imposed restrictions on AI use, and law schools have been slow to integrate AI training into their curriculum. The legal profession's reluctance to embrace AI stems from concerns about accuracy, confidentiality, and professional responsibility. This Article contends that law schools must play a proactive role in preparing students for incorporated AI in their practice. It examines the ethical obligations of lawyers regarding AI use, evaluates reactionary regulations, and proposes a pedagogical framework that equips students with the skills necessary to leverage AI effectively. The legal profession must acknowledge AI's growing role and ensure that future lawyers are trained to use it responsibly and strategically.

INTRODUCTION

At least in theory, lawyers have a fiduciary duty to be efficient. The duty of care requires lawyers to exercise reasonable professional care when providing legal services to their clients,¹ while the duty of loyalty requires lawyers to pursue their client's interests to the exclusion of other interests, including their own.² So, in combination, the duties of care and loyalty require lawyers to provide competent legal representation as efficiently as possible.

Of course, in practice, this implied duty of efficiency is hard to apply

* Assistant Clinical Professor of Law, Georgia State University College of Law.

** Spears-Gilbert Professor of Law, University of Kentucky J. David Rosenberg College of Law.

1. See MODEL RULES OF PRO. CONDUCT r. 1.1 (AM. BAR ASS'N 2024); see also *id.* at r. 1.3.

2. See, e.g., *id.* at r. 1.2.

and even harder to enforce. Applying the duties of care and loyalty always requires a judgment call, so applying them in tandem necessarily requires balancing one judgment against another. When does observing the duty of care conflict with the duty of loyalty? Lawyers can always do more work, but only at the risk of diminishing returns. Eventually, the juice isn't worth the squeeze. But different lawyers may have different opinions about how much work is justified under the circumstances, depending on their view of efficiency, and it is hard to imagine disciplining a lawyer over a reasonable judgment call about how much care to take.

While the duty of efficiency is inevitably merely aspirational, it can still inform our practices and expectations. Social norms matter because people observe them, irrespective of the law. If lawyers accept a duty of efficiency, then lawyers will observe a duty of efficiency, at least when their peers are watching. No one wants a reputation for padding hours with unnecessary and unproductive work, especially because it encourages your competitors to poach your clients.

So, how can we encourage lawyers to observe the duty of efficiency? One of the best and easiest ways is encouraging them to adopt new, productivity-enhancing technology. The legal profession often resists change, including technological change. Everyone remembers the partner who refused to use a computer, insisted on printing emails, or resisted e-discovery. But all of those technologies are now universal.

There are still plenty of ways technology can increase the efficiency of the legal profession. One of them is integration of artificial intelligence (AI) tools into legal practice. AI already has many uses, and we are discovering new ones every day. One of the most promising applications of AI for the legal profession is the development of large language models (LLMs) that can automatically respond to a written prompt by generating a plausible written response. Lawyers are in the business of answering questions, often in writing, and LLMs promise to make their work faster and easier. Of course, LLMs also promise to make lawyering cheaper—a prospect most lawyers find less appealing. Even worse, LLMs cannot guarantee the accuracy of their responses, and some lawyers have already filed inaccurate documents that were generated by an LLM.

As always, technology can be used for good or ill, and AI is no exception. Some lawyers have already gotten themselves into trouble by misusing LLMs, causing some courts to adopt rules prohibiting or limiting

the use of AI tools. This is a mistake. The problem was never the AI tool, but its misuse. Courts should instead help lawyers understand how to use LLMs and other AI tools effectively and efficiently. AI can help lawyers provide better legal advice to more people at lower cost—if we let it.

For their part, law schools can and should teach law students how to use LLMs and other AI tools. As it stands, many law schools prohibit the use of AI tools, or at least strongly discourage their use. This is also a mistake. If lawyers do not know how to use AI tools well, they are bound to use AI tools poorly—because there is little question that lawyers will be using them. And the best way to ensure that lawyers know how to use AI tools well, is for law schools to teach law students how to use AI tools well.

In our article *Plagiarism Pedagogy: Why Teaching Plagiarism Should be a Fundamental Part of Legal Education*, we argued that law schools should teach law students how to plagiarize effectively, because plagiarism is an essential lawyering skill.³ After all, in the practice of law, if you are not plagiarizing, you are committing malpractice. Clients do not hire lawyers to produce original works of authorship—they hire lawyers to get results. Often, the cheapest and easiest way to get those results is to reuse an existing document—in other words, to plagiarize. And it is not just permitted, it is an obligation. Refusing to plagiarize when appropriate would violate the lawyer’s fiduciary duty to their client and conceivably even be sanctionable.

In this Essay, we argue that the same is true of AI tools like LLMs. We explain how AI tools work and how lawyers can use them. We describe some of the problems associated with lawyers using AI tools and how courts have responded to these problems. We reflect on the many ways in which lawyers could use AI tools to improve their effectiveness and efficiency. And we conclude by arguing that law schools should be teaching law students how to use AI tools, rather than prohibiting them.

I. WHAT IS GENAI AND WHAT CAN IT DO FOR LAWYERS?

Generative AI (GenAI) products create content in response to prompts

3. See Megan E. Boyd & Brian L. Frye, *Plagiarism Pedagogy: Why Teaching Plagiarism Should be a Fundamental Part of Legal Education*, 99 WASH. U. L. REV. ONLINE 1, 1–2 (2021), <https://wustllawreview.org/wp-content/uploads/2021/11/Plagiarism-Pedagogy-.pdf> [<https://perma.cc/FLQ2-ZTYN>].

from the user. GenAI can produce video, audio, and other types of output, but most importantly for lawyers, GenAI products use LLMs to respond to written prompts by generating texts. The process is somewhat complicated, but in its most basic form, the user asks the GenAI tool a question (“What is the summary judgment standard under the Federal Rules of Civil Procedure?”) or inputs a prompt (“Summarize this hearing transcript.”) and the GenAI tool relies on LLMs to generate a “human-sounding” response.⁴

The most important thing that AI can do for lawyers is increase their efficiency. Lawyers are in the business of selling legal advice to clients. Typically, lawyers bill by the hour. Consequently, lawyers want to bill as many hours as possible at the highest price possible.

But lawyering is also a profession, with the typical information asymmetries between principal and agent. Accordingly, lawyering is not an arms-length business. Lawyers have fiduciary duties to their clients, which are intended to prevent lawyers from taking advantage of information asymmetries and to punish them if they do. The duty of care requires lawyers to exercise reasonable professional judgment, and the duty of loyalty requires lawyers to put their clients’ interests before their own.⁵ But the combination of the two can give lawyers a tempting excuse to maximize their diligence in order to bill as many hours as possible.

The duty of efficiency requires lawyers to be as diligent as prudent under the circumstances, but not more so, and to use technology to provide legal services as efficiently as possible. GenAI tools can already make lawyers more efficient in many ways. They can, for example, write website copy, summarize specific cases, analyze contracts, and assist with due diligence.⁶ And LLMs can respond to queries by producing texts that at least superficially resemble a response that might be provided by a lawyer.

While GenAI tools are getting better and better at these tasks, one thing they still cannot do (and may never be able to do) is perform legal analysis. And, as we will see below, even now, they cannot consistently produce the type of legal writing, supported by citations to authority and complete with

4. See e.g., Adam Zewe, *Explained: Generative AI*, MIT NEWS (Nov. 9, 2023), <https://news.mit.edu/2023/explained-generative-ai-1109> [<https://perma.cc/4VV8-KWQW>] (describing the basics of GenAI as well as the “increase in complexity” of GenAI models).

5. See *supra* notes 1–2 and accompanying text.

6. See e.g., Michael D. Murray, *Artificial Intelligence for Learning the Law: Generative AI for Academic Support in Law Schools and Universities*, 8 J.L. & TECH. TEX. (forthcoming 2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4564227 [<https://perma.cc/6SMB-LLVZ>].

analysis, expected of lawyers.⁷ Or rather, AI can produce documents that superficially resemble legal reasoning, but cannot actually do the reasoning itself. A text produced by an AI may look like a legal memorandum or brief, complete with what appear to be statements of fact, legal reasoning, and citations to authority. But there is no guarantee the statements of fact are true, the legal reasoning is sound, or the cited authorities actually exist or say what the text claims. Of course, it is possible for an AI to produce an accurate text. But it is also possible for an AI to produce an inaccurate text. And there is no way for the AI to tell the difference—that takes a person.

While lawyers cannot—and must not—rely on AI to produce legal documents that are ready for filing, they can use AI for many other purposes, so long as they ensure the text produced by the AI is factually accurate before using it. For example, they can use AI to generate summaries of facts, cases, or articles. They can use AI to generate draft agreements. And they can even use AI to generate potential legal arguments.

This last use could be especially effective for lawyers, given how AI works. The primary purpose of legal reasoning and argumentation is to be persuasive. While there is some room for creativity in a legal argument, it is cabined by existing laws and legal precedents, as well as conventional wisdom. That's the sweet spot for AI, which cannot express new ideas, only recycle existing ones, and is designed to reproduce conventional wisdom.⁸

As a consequence, AI is good at expressing ideas in ways that people understand and do not find surprising. Indeed, AI is even good at making ideas seem unsurprising, or at least less surprising than they are. AI may also be good at generating arguments that lawyers may not have considered, but that sound convincing and unsurprising once available to them. In other words, AI may help lawyers discover arguments and ideas that they would otherwise have overlooked. Some of those AI-generated arguments may help them advance their client's position, and some of them may help them respond to their opponent's position. In any case, AI may improve the quality of legal argumentation by helping lawyers see the unseen. But to harness AI's power and ensure they do not end up in trouble, lawyers must understand how to use AI ethically and appropriately.

7. *See e.g., id.*

8. *See, e.g.,* Michael Townsen Hicks et al., *ChatGPT is Bullshit*, 26 ETHICS AND INFO. TECH. 37, 38 (2024).

II. AI IN LEGAL PRACTICE

The legal profession resists change. Lawyers' understanding of technology, as a whole, lags far behind members of other professions. Technological competency is required by Comment 8 to Rule 1.1 of the Model Rules of Professional Conduct (MRPC), which many states have adopted.⁹ Some states, like Florida and North Carolina, have gone further, requiring Continuing Legal Education (CLE) hours in technological competency.¹⁰ Still, as a profession, lawyers have only a baseline understanding of basic technology—if you ask ten random lawyers whether they know how to convert a Word document to a .pdf, you might be surprised how many do not know. The federal courts' outdated and clunky Case Management/Electronic Case Files (CM/ECF) and Pacer systems are greatly in need of updating or replacement by more user-friendly and less costly programs. And most law schools still use the Langdellian model of instruction, despite copious evidence that it does not produce the best results for the greatest number of students.¹¹

But change is inevitable, even if it is slow. Consider the discovery process in litigation. Twenty years ago, many Wall Street law firms still conducted discovery on paper. Documents were Bates-stamped by hand, and every page was individually reviewed by a lawyer for relevance and privilege. While computerized document review was possible, we watched many law firms resist it in favor of doing things the old and familiar way. But they could not resist for long. The same was true of AI-assisted document review. At first, we saw many lawyers resist it. But when it became clear that AI was more accurate than people, they had to accept it. Today, it would be obvious malpractice to conduct a substantial document

9. MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 8 (AM. BAR ASS'N 2024); *Jurisdictional Adoption of Model Rule 1.1, Comment [8] on Tech Competence*, AM. BAR ASS'N (Apr. 4, 2024), https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/mrpc1-1-comment-8.pdf [https://perma.cc/S23S-GN7Z].

10. See *General CLE Info & Requirements*, THE FLA. BAR, <https://www.floridabar.org/member/cle/general-cle-info-and-requirements/#:~:text=At%20a%20Glance-,Requirement,mandatory%20Florida%20Legal%20Professionalism%20course> [https://perma.cc/Z2CP-P98G] (requiring three hours of credit from approved technology courses every three years); *CLE Requirements in North Carolina for Lawyers*, N.C. STATE BAR CONTINUING LEGAL EDUC., <https://www.nccle.org/for-lawyers/requirements/renewing-lawyers/> [https://perma.cc/UYT3-QXS7] (requiring one technology credit hour every two years).

11. See Edward Rubin, *What's Wrong with Langdell's Method, and What to Do About It*, 60 VAND. L. REV. 609, 610, 613 (2019).

review on paper without AI assistance.

In the last few years, GenAI has become ubiquitous. Yet lawyers as a group lack any meaningful understanding of how GenAI works and what it can—and cannot—do. Based on the number of advertising emails that the authors have received in the last few months, numerous enterprising CLE providers are happy to help remedy the GenAI confusion for a substantial fee, but we very much doubt that this will lead to any immediate widespread understanding about GenAI.

III. THE FUTURE OF AI IN LEGAL PRACTICE

Within a few years, virtually every legal professional will be using GenAI in some capacity. According to one forecast, the GenAI market in the legal technology industry will reach \$50 billion by 2027, a 69% increase over the next three years.¹² GenAI is here to stay, and it promises to revolutionize the way lawyers practice law. Many large law firms have already bought in. According to an article in the *American Lawyer*, Crowell & Moring uses GenAI to create client alerts on important public hearings.¹³ Dechert has created its own GenAI product to draft disclosures for mutual fund clients, and Dentons' proprietary GenAI is being used to draft witness statements and summarize documents, among other things.¹⁴ They are surely not alone. And of course, many firms are using GenAI in legal research and drafting. Biglaw firms get it: GenAI is an “amazing personal productivity tool,” that can enable firms to enhance the delivery of legal services.¹⁵ But Biglaw firms also understand GenAI's limitations: “[G]en AI's output is only as savvy as the professional using it,” and while GenAI generated content is a “good first step,” it still must be vetted and checked by lawyers.¹⁶

12. Marc Doucette, *GenAI Continues to Spark More Growth in the Legal Technology Market*, KOHO CONSULTING (June 22, 2024), <https://kohoconsulting.com/blog/genai-continues-to-spark-more-growth-in-the-legal-technology-market/> [https://perma.cc/C2ZG-HH6T].

13. Justin Henry, *We Asked Every Am Law 100 Law Firm How They're Using Gen AI. Here's What We Learned*, AM. LAW. (Jan. 29, 2024, 5:00 AM), <https://www.law.com/americanlawyer/2024/01/29/we-asked-every-am-law-100-firm-how-theyre-using-gen-ai-heres-what-we-learned/?slreturn=20241108105729> [https://perma.cc/F8WQ-BLAL].

14. *Id.*

15. *Id.*

16. *Id.*

IV. LAWYERS ABUSING AI

While the legal profession will inevitably adopt AI tools because of their efficiency, it will probably take a while, and there will definitely be problems along the way. We have seen some of those problems already. The most common problem arises when a lawyer uses a text generated by an LLM without verifying its accuracy. As one court put it, “It is not necessarily the use of AI in and of itself that causes such offense and concern, but rather the attorney’s failure to review the sources produced by AI without proper examination and scrutiny.”¹⁷

ChatGPT, a free and popular GenAI product, was launched as an early demo on November 30, 2022, and, within a week, more than a million users had tried it.¹⁸ Some of those were lawyers, who decided that ChatGPT could be a useful tool because of its ability to quickly generate written responses to questions. After all, legal briefs can take many hours to write, while ChatGPT can produce a legal-looking brief in a matter of seconds. Who wouldn’t consider using ChatGPT to do the hard work of lawyering—making legal arguments? The problem for some of those lawyers is that they simply took the responses that ChatGPT generated and filed them without checking that statements of law were accurate, that cited cases actually existed, or that the cited cases stood for the legal propositions advanced. Within a few months of ChatGPT’s release, the courts began dealing with GenAI-generated filings when it became apparent that ChatGPT tended to “hallucinate,” or make up, cases and citations.¹⁹

The first judicial opinion to address GenAI in filings was *Mata v. Avianca* in 2023.²⁰ In *Mata*, plaintiff’s counsel relied heavily on ChatGPT to respond to a motion to dismiss but failed to check the accuracy of the cases cited in the ChatGPT-generated content.²¹ Unfortunately for the plaintiff, the cases cited and corresponding quotations from those cases

17. Will of Samuel, 206 N.Y.S.3d 888, 891 (N.Y. Sur. Ct. 2024).

18. Bernard Marr, *A Short History of ChatGPT: How We Got to Where We Are Today*, FORBES (May 19, 2023, 1:14 AM), <https://www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-how-we-got-to-where-we-are-today/> [<https://perma.cc/BWX3-U7M9>].

19. See *What are AI Hallucinations?*, I.B.M. (Sept. 1, 2023), <https://www.ibm.com/think/topics/ai-hallucinations> [<https://perma.cc/5HNV-D2KR>] (explaining hallucinations of AI).

20. 678 F. Supp. 3d 443 (S.D.N.Y. 2023).

21. *Id.* at 448, 450–51.

were made up by ChatGPT.²² The cases (and the helpful quotations from them) simply did not exist.

Counsel for the defendant first notified the court and the plaintiff that it could not locate many of the cases the plaintiff cited in his brief, and the court, after doing its own research, confirmed that most of the cited cases were non-existent.²³ Yet the plaintiff's lawyers did not withdraw the filing, instead doubling down and defending their work for months before finally admitting that ChatGPT was responsible for the errors.²⁴ Some of their most egregious conduct involved the submission of an affidavit by one of the lawyers that purported to include, as attachments, copies of excerpts of cases identified in the plaintiff's brief but which the defendant's counsel and the court has been unable to locate.²⁵ Ultimately, the court discovered that while those excerpts had "some traits that are superficially consistent with actual judicial decisions," the cases were non-existent, as verified by the courts from which the cases had purportedly been issued.²⁶

The *Mata* opinion begins by acknowledging that "[t]echnological advances are commonplace and there is nothing inherently improper about using a reliable artificial intelligence tool for assistance."²⁷ The court found, however, that Rule 11 "impose[s] a gatekeeping role on attorneys to ensure the accuracy of their filings," and the plaintiff's attorneys "abandoned their responsibilities" by submitting a brief with fake quotes and non-existent citations.²⁸ The court then identified numerous harms that can flow from this type of conduct including: (1) waste of the opposing party's time and money; (2) waste of judicial resources; (3) "potential harm to the reputation of judges and courts whose names are falsely invoked as authors of the bogus opinions and to the reputation of a party attributed with fictional conduct"; (4) the promotion of "cynicism about the legal profession and the American judicial system"; and (5) the possibility that a future litigant may attempt to "defy a judicial ruling by disingenuously claiming doubt about its authenticity."²⁹

22. *Id.*

23. *Id.* at 451–58.

24. *Id.*

25. *Id.* at 451–52.

26. *Id.* at 454–55.

27. *Id.* at 448.

28. *Id.*

29. *Id.* at 448–49.

Still, the *Mata* lawyers received only a light punishment—they voluntarily hired a CLE provider to offer education on technological competence and AI programs to members of the firm.³⁰ Relying on its authority under Rule 11, the court also ordered the lawyers to notify their client of the opinion and send letters and a copy of the show cause hearing transcript to each judge identified as the writer of one of the fake excerpts submitted to the court.³¹ Finally, the court fined the lawyers \$5,000.³²

The United States Court of Appeals for the Second Circuit also relied on Rule 11 to punish an attorney for filing an appellate brief primarily created using ChatGPT. In *Park v. Kim*, a plaintiff's complaint was dismissed by the trial court under Rule 37 as a sanction for failing to comply with discovery requests and the court's orders regarding discovery.³³ On appeal, plaintiff's counsel received multiple extensions to file a reply brief but instead filed a late, non-compliant brief that contained citations to only two cases.³⁴ When the court attempted to locate one of the cases, it could not find the case because it did not exist.³⁵ The attorney ultimately admitted that she had relied on ChatGPT to locate a case to support the legal position for which she was arguing, having successfully used the GenAI tool for other purposes, but "did not cite any specific reasoning or decision" from the non-existent case.³⁶

The court held that Rule 11 requires, at a minimum, "that attorneys read, and thereby confirm the existence and validity of, the legal authorities on which they rely."³⁷ The court found that the attorney had violated Rule 11(b)(2) by failing to "conduct a reasonable inquiry and determine that any papers filed with the court were well grounded in fact, and legally tenable."³⁸ And it held that Rule 11 authorizes courts to "sanction an attorney for, among other things, misrepresenting facts or making frivolous legal arguments."³⁹ The court sanctioned the attorney by referring the matter to the court's Grievance Panel and requiring the attorney to provide her

30. *See id.* at 455–66.

31. *Id.*

32. *Id.*

33. 91 F.4th 610, 613 (2d Cir. 2024).

34. *See id.* at 613–14.

35. *Id.* at 614.

36. *Id.*

37. *Id.* at 615.

38. *Id.* at 614–15 (cleaned up).

39. *Id.* at 615 (internal quotations and citation omitted).

client with a copy of the court's opinion.⁴⁰

State courts have also relied upon state-specific rules of procedure and state specific ethics rules to sanction parties for submitting briefs containing non-existent authority that were generated, at least in part, by GenAI programs. In *Kruse v. Karlen*, the perpetrator was a pro se appellant who submitted an appellant's brief in which "the overwhelming majority of the citations [were] not only inaccurate but entirely fictitious."⁴¹ In fact, twenty-two of the twenty-four cases cited in the appellant's brief were either entirely fictitious or a fictitious citation using a real case name.⁴² The court held that the appellant had violated numerous rules of civil and appellate procedure because, "[c]iting nonexistent case law or misrepresenting the holdings of a case is making a false statement to a court; it does not matter if generative A.I. told you so."⁴³ Relying on its authority under Rule 84.19 of the Missouri Rules of Civil Procedure to award sanctions against a party who files a frivolous appeal, the court ordered the appellant to pay the appellee \$10,000 in appellate attorneys' fees.⁴⁴

Similarly, in a Colorado disciplinary proceeding, the attorney used ChatGPT to draft a motion to set aside a judgment.⁴⁵ The attorney failed to verify that the citations in the motion were accurate and even upon discovering that they were not, the attorney failed to withdraw the motion or otherwise notify the court of the error.⁴⁶ When confronted about the made-up citations in the brief, the attorney blamed the problem on a legal intern, before finally admitting the truth.⁴⁷ The disciplinary council found that the attorney had violated the Colorado Rules of Professional Conduct pertaining to competent representation (Rule 1.1), diligence (Rule 1.3), false statements of material fact or law (Rule 3.3), and dishonesty, fraud, deceit, and misrepresentation (Rule 8.4).⁴⁸ The attorney received a year-long suspension from law practice which was reduced to ninety days upon

40. *Id.* at 616.

41. 692 S.W.3d 43, 45, 48 (Mo. Ct. App. 2024).

42. *Id.* at 48–51.

43. *Id.* at 52–53 (internal quotation marks omitted) (quoting Maura R. Grossman et al., *Is Disclosure and Certification of the Use of Generative AI Really Necessary?*, 107 JUDICATURE 68, 75 (2023)).

44. *Id.* at 53–54.

45. *People v. Crabill*, No. 23PDJ067, 2023 WL 8111898, at *1 (Colo. O.P.D.J. Nov. 22, 2023).

46. *Id.*

47. *Id.*

48. *Id.*

satisfactory completion of a two-year probationary period.⁴⁹

Not all courts have even awarded sanctions for the filing of GenAI pleadings containing false statements of law. Some courts have used problems with GenAI-created pleadings to warn pro se parties (and attorneys alike) that the submission of false citations and cases can land them in real trouble with the courts. For example, in *Anonymous v. New York City Department of Education*, a pro se litigant apparently used a GenAI product to assist in preparing court filings and submissions.⁵⁰ The court addressed the issue by noting that it is “improper and unacceptable for litigants—including *pro se* litigants” to file documents with non-existent or fake citations.⁵¹ While acknowledging that it had authority under Rule 11 to sanction the plaintiff for those errors, it declined to do so, warning that “the [c]ourt w[ould] not look kindly on similar infractions in the future.”⁵² As one court recognized, unlike attorneys who have access to Westlaw, Lexis, CaseText, and similar programs that could easily be used to verify the existence and veracity of the legal arguments presented in GenAI-created legal documents, a pro se party likely is not in a position to “check the veracity of case citations generated by AI programs.”⁵³ While this is true, pro se parties are held to the same standards as lawyers vis-a-vis Rule 11, and thus must understand the limitations of AI programs if they wish to use them in legal proceedings.

V. RESPONSES TO AI ABUSE

The enormous uptick in lawyer use of GenAI has prompted numerous state bars and courts to implement rules regarding AI use. Some courts have required lawyers to disclose GenAI usage in any filed documents, while others have simply reminded lawyers of their professional duties under the relevant rules of professional conduct and Rule 11, and a few have outright

49. *Id.*

50. No. 1:24-cv-04232 (JLR), 2024 WL 3460049, at *1, *7 (S.D.N.Y. July 18, 2024).

51. *Id.* at *7.

52. *Id.*; see also *Morgan v. Cmty. Against Violence*, No. 23-cv-353-WPJ/JMR, 2023 WL 6976510, at *7 (D.N.M. Oct. 23, 2023) (noting that pro se plaintiff filed briefs with court that contained fake or non-existent cases, likely generated by GenAI, and that while the court would not sanction the plaintiff at that time, any future submissions containing false or non-existent citations could result in dismissal of the plaintiff’s case).

53. *Dukuray v. Experian Info. Sols.*, No. 23 Civ. 9043 (AT) (GS), 2024 WL 3812259, at *11 (S.D.N.Y. July 26, 2024).

banned the use of GenAI by certain types of litigants, such as pro se parties or out-of-state attorneys.

These knee jerk responses are misguided and ineffective. First, they misdiagnose the problem. The issue was never lawyers using AI to generate content, rather the problem was lawyers filing documents they did not read or understand. That problem is as old as time, and certainly cannot be pinned on AI. Bad lawyers have always tried to save time by cutting corners. But they can do it just as easily with a form as with an LLM.

Moreover, some courts have responded with ridiculous orders prohibiting the use of AI or requiring its disclosure. For example, the Eastern District of Missouri briefly prohibited pro se litigants—but not licensed and admitted attorneys—from using GenAI in any filings.⁵⁴ Other judges require certifications from lead counsel regarding the correctness of material contained in filings incorporating GenAI content.⁵⁵ A judge from the United States Court of International Trade requires that the portions of any filing generated by AI be identified and, perhaps more critically, that the filer certify that the use of GenAI “has not resulted in the disclosure of any confidential or business proprietary information to any unauthorized party.”⁵⁶ The problem with this rule is that a lawyer who used a Lexis or Westlaw generative AI product can probably make this certification with a reasonable degree of certainty, but a lawyer using a free or open access program like ChatGPT or Claude almost certainly cannot,⁵⁷ effectively preventing the lawyer from using open source GenAI at all.

In perhaps the oddest GenAI order to date, the judges of the United States District Court for the Western District of North Carolina expressed “concern[] regarding whether factual and legal citations or references in

54. See Stephanie Pacheco, *Missouri AI Rules Could Impede Access to Justice*, BLOOMBERG L. (Mar. 7, 2024), <https://news.bloomberglaw.com/bloomberg-law-analysis/analysis-missouri-ai-rules-could-impede-access-to-justice> [<https://perma.cc/6XDG-RJ4G>]. But see *Self-Represented Litigants (SRL)*, U.S. DIST. CT. E.D. MO., <https://www.moed.uscourts.gov/self-represented-litigants-srl> [<https://perma.cc/5BSB-GW6Y>] (noting rules regarding use of Generative Artificial Intelligence including ChatGPT, Harvey.AI, and Google Bard).

55. See, e.g., Judge Araceli Martínez-Olguín, United States District Court for Northern District of California, *Standing Order for Civil Cases* (Nov. 22, 2023), <https://www.cand.uscourts.gov/wp-content/uploads/2023/03/AMO-Civil-Standing-Order-11.22.2023-FINAL.pdf> [<https://perma.cc/D7DT-T2T6>].

56. Judge Stephen Alexander Vaden, United States Court of International Trade, *Order on Artificial Intelligence* (June 8, 2023), <https://www.cit.uscourts.gov/sites/cit/files/Order%20on%20Artificial%20Intelligence.pdf> [<https://perma.cc/299S-GT2E>].

57. Megan E. Boyd, *AI, Lawyers, and the Courts*, 30 GA. BAR J. 14, 18–19 (Nov. 2024).

court filings have been properly vetted by counsel or the pro se parties filing the documents.”⁵⁸ To remedy this perceived problem, the court prohibited the use of all but four specifically enumerated programs—Westlaw, Lexis (initially spelled Lexus and apparently corrected with a pen in the filed order), FastCase, and Bloomberg—and required that parties or their attorneys certify that “[n]o artificial intelligence was employed *in doing the research*” to prepare any document filed with the court.⁵⁹

There are several problems with this court’s approach to the use of GenAI. First, we find it exceptionally problematic that the court appears to be telling lawyers what specific programs they can and cannot use to do their jobs, which is wholly outside of the court’s authority. But second, and perhaps even more importantly, in the words of Inigo Montoya: “I do not think it means what you think it means.”⁶⁰ As written, the order prohibits the use of all but four specific programs “in doing research,” but says nothing about using GenAI programs to produce content. Non-lawyer GenAI programs like ChatGPT and Claude do not “do research”—they generate content. So, under a plain language reading, the order prohibits lawyers and pro se parties from using GenAI programs that cannot “do research” for the purpose of doing research but does not prohibit them from using GenAI programs to generate content and using it in their filings. If there is a GenAI problem, it is with filing documents using AI-generated content that has not been checked for accuracy, not with using GenAI to do something—research—that it does not and cannot do. This order exemplifies the extreme misunderstanding within the legal profession about what GenAI is, how it works, and what it can and cannot do.

Judges would never consider implementing disclosure requirements or prohibitions regarding the use of form documents. But how is using AI any different? The problem with AI is not the technology, but how it is used. “The way to remedy AI problems is not to demand disclosure of GenAI content because doing so will discourage its responsible use, could disincentivize companies from developing GenAI programs that greatly increase lawyer productivity, and may result in worse rather than better research and writing.”⁶¹

58. *In re Use of Artificial Intelligence*, No. 3:24-mc-104 (W.D.N.C. June 18, 2024).

59. *Id.* (emphasis added).

60. *THE PRINCESS BRIDE* (20th Century Fox 1987).

61. Boyd, *supra* note 57.

While several courts have felt compelled to create GenAI-specific rules, they are not necessary. As GenAI cases have already shown us, the existing rules governing litigation procedure and state-specific bar rules of professional conduct are more than sufficient to deal with any problems created by hallucinated cases or fake citations. Rules 11(b)(2) of the Federal Rules of Civil Procedure provides that by signing and filing documents with the court, the litigant certifies that “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.”⁶² A litigant or attorney who files a brief containing hallucinated citations or non-existent cases has violated this rule and can—and should—be punished under it. Similarly, state procedural and appellate rules that prohibit litigants from asserting frivolous claims,⁶³ unnecessarily expanding the proceedings,⁶⁴ or filing of frivolous appeals⁶⁵ cover the filing of GenAI created briefs that contain fake citations, as *Kruse* demonstrates.⁶⁶ Finally, state bar disciplinary proceedings against attorneys who engage in this conduct can be appropriate.

The filing of GenAI created pleadings that contain hallucinated or fake citations violates numerous Rules of Professional Conduct,⁶⁷ as the American Bar Association (ABA) indicated in July 2024 when it issued Formal Opinion 512, addressing lawyer use of GenAI tools. The opinion clarifies the ABA’s position on lawyers’ ethical duties vis-à-vis GenAI by providing that lawyers must “have a reasonable understanding of the capabilities and limitations” of GenAI that they use, must protect client confidentiality when using GenAI tools, must determine whether relevant state rules of professional conduct require that the lawyer disclose GenAI usage to the client, and must ensure that GenAI content is factually and legally correct.⁶⁸

62. FED. R. CIV. P. 11(b)(2).

63. *See, e.g.*, GA. CODE ANN. § 9-15-14(a) (2023) (imposing attorneys’ fees against a party that “has asserted a claim, defense, or other position with respect to which there existed such a complete absence of any justiciable issue of law or fact”).

64. *See, e.g., id.* § 9-15-14(b) (allowing imposition of attorneys’ fees against a party that “unnecessarily expanded the proceeding by other improper conduct”).

65. *See, e.g.*, CAL. CIV. PROC. CODE § 907 (West 2024) (allowing for the imposition of sanctions where an appeal “was frivolous or taken solely for delay”).

66. *See, e.g., supra* notes 41–44 and accompanying text.

67. *See, e.g.*, MODEL RULES OF PRO. CONDUCT r. 1.1, 1.3, 3.3, 8.4 (AM. BAR ASS’N 2024).

68. *See* ABA Comm. on Ethics & Pro. Resp., Formal Op. 512 (2024).

In short, existing rules are more than capable of addressing any problems associated with GenAI filings that contain hallucinated or false citations. While pro se parties cannot be disciplined for violating attorney ethical rules, the Federal Rules of Civil Procedure, state rules of civil procedure, and state-specific trial and appellate statutes have proven to be powerful tools to discourage the misuse of GenAI-produced content in legal proceedings for both attorneys and non-attorney litigants alike.

Further, though no United States courts or bars appear to have addressed the issue, GenAI raises questions about billing. If an attorney uses GenAI to create a document or help prepare for a deposition or argument, the lawyer can only bill the client for the time actually spent doing the client's work.⁶⁹ As with plagiarizing documents, there is nothing wrong with lawyers using AI in appropriate ways, and lawyers should be encouraged to do so when it will benefit the client. At least one Canadian court has recognized this. In *Cass v. 1410088 Ontario, Inc.* (2018), Justice Whitten of the Superior Court of Justice decided the question of the appropriate amount of fees to award to the defendant, the party that prevailed at summary judgment.⁷⁰ Judge Whitten notes that defendant's counsel billed exceptional amounts for legal research on relatively banal issues, and "[i]f artificial intelligence sources were employed, no doubt counsel's preparation time would have been significantly reduced."⁷¹ Thus, lawyers can and should use GenAI to save their clients' money. But they cannot bill for time they did not actually work. Fraudulent and excessive billing—not GenAI use—is wrong and punishable under current rules applicable in every state.

VI. AI PRAXIS

So, how can lawyers, judges, legal scholars, law students, and others use GenAI in legal practice, pedagogy, and learning? People have already identified many ways, and are sure to discover many more, as they develop

69. See Megan E. Boyd & Brian L. Frye, *Plagiarism Pedagogy: Why Teaching Plagiarism Should be a Fundamental Part of Legal Education*, 99 WASH. U. L. REV. ONLINE 1, 1–2 (2021), <https://wustllawreview.org/wp-content/uploads/2021/11/Plagiarism-Pedagogy-.pdf> [<https://perma.cc/FLQ2-ZTYN>] (observing that the Rules of Professional Responsibility compel lawyers to plagiarize, but prohibit them from billing clients as if they had written plagiarized material).

70. See *Cass v. 1410088 Ontario, Inc.*, 2018 ONSC 6959 (Can.).

71. *Id.* para. 34.

proficiency using GenAI tools and as those tools improve. As previously observed, GenAI is an excellent tool for producing useful texts, with the caveat that any text produced by a GenAI tool must be carefully checked for truth and accuracy before use. As one of us has shown, you can even use AI to produce legal scholarship, so long as you use it thoughtfully.⁷² AI is the angle grinder of legal scholarship tools: a huge time saver that can quickly become incredibly dangerous if you do not pay attention and use it safely.

Accordingly, we must teach our students about GenAI and do so in ways that help prepare them for the future of law practice. That first entails understanding GenAI ourselves and, in particular, the type of GenAI used by common legal research platforms. To give an idea of the types of output that GenAI products can produce, we asked three different GenAI programs the following question: “What are the elements of a slip and fall tort claim in California?” ChatGPT 4.0 responded as follows:

In California, a slip and fall tort claim typically involves proving the following elements:

1. Duty of Care: The property owner or possessor had a legal duty to maintain the property in a safe condition. This duty varies depending on whether the injured person was an invitee, licensee, or trespasser.
2. Breach of Duty: The property owner or possessor failed to fulfill this duty. This can be through negligence, such as not fixing a known hazard or not conducting regular inspections.
3. Causation: The breach of duty directly caused the slip and fall accident. This means showing that the hazardous condition was a substantial factor in causing the fall.
4. Damages: The injured person suffered actual harm or damages as a result of the fall. This can include

72. See generally Brian L. Frye & ChatGPT, *Should Using an AI Text Generator to Produce Academic Writing Be Plagiarism?*, 33 FORDHAM INTELL. PROP. MEDIA & ENT. L.J. 946 (2023); Brian L. Frye, *Apologia Pro Plagio Suo*, J. COPYRIGHT SOC'Y (forthcoming 2025).

physical injuries, medical expenses, lost wages, pain and suffering, and other related damages.

To succeed in a slip and fall claim, it's crucial to provide evidence supporting each of these elements, including proof of the hazardous condition, the property owner's knowledge of it (or lack of reasonable care in discovering it), and the resulting damages.⁷³

You will immediately notice that this summary has no citations to relevant California authority and thus provides no easy way for the user to check its accuracy. And while the summary might be broadly helpful in giving the user some idea about the general state of the law, it is effectively useless to lawyers, who will still have to locate relevant authority to support the summary. Notably, ChatGPT's privacy policy indicates that user information will be used for numerous purposes, including to train its LLMs.⁷⁴ While ChatGPT provides that it takes "commercially reasonable measures" to prevent disclosure of information, ChatGPT cannot ensure full security and warns users to "take special care in deciding what information [to] provide."⁷⁵ This is important because users are explicitly told that ChatGPT does not guarantee confidentiality. Accordingly, lawyers should not input privileged or confidential information into ChatGPT.

Lexis+ AI Assistant produced the following response to the California slip-and-fall question:

In California, the elements of a slip and fall case, which fall under the broader category of premises liability, include duty, breach, causation, and damages. These elements are essential to establish a valid claim for negligence or premises liability arising from a slip and fall incident.

Firstly, there must be a duty owed by the defendant to the plaintiff. This duty is generally to ensure that the premises are reasonably safe for visitors. Secondly, there must be a breach of this duty, meaning that the defendant failed to maintain the premises in a safe condition or failed to warn

73. CHATGPT 4.0, <https://chatgpt.com/> [<https://perma.cc/A2HV-ZFUM>] (submitting the question, "What are the elements of a slip and fall tort claim in California?") (results may vary).

74. *OpenAI Privacy Policy*, <https://openai.com/policies/privacy-policy/> [<https://perma.cc/SN3B-5ABB>].

75. *Id.*

of dangers. Thirdly, this breach must cause an accident or injury, which is the causation element. Finally, the plaintiff must have suffered actual damages as a result of the fall, such as medical expenses, lost wages, or pain and suffering.

Several cases highlight these elements, demonstrating the application of these principles in various circumstances. For instance, in cases where the defendant had control over the premises and was found negligent in its maintenance, leading to the plaintiff's injuries, the courts have upheld the principles of premises liability *Nava v. Shalhoub Mgmt. Co.*, 2022 Cal. Super. LEXIS 95614, *Guzman v. Home Depot U.S.A.*, 2021 Cal. Super. LEXIS 138780, *Ramos v. B & V Enters.*, 2019 Cal. Super. LEXIS 13707. Additionally, the courts have consistently emphasized that the plaintiff must prove all these elements to succeed in a slip and fall claim *McCarthy v. Century City*, 2023 Cal. Super. LEXIS 19026, *Momeni v. Walmart, Inc.*, 2021 Cal. Super. LEXIS 42580.

In summary, understanding the elements of duty, breach, causation, and damages is crucial for navigating slip and fall cases in California. These elements are consistently applied across various cases to determine the liability of the defendant in maintaining safe premises for lawful visitor. *Nava v. Shalhoub Mgmt. Co.*, 2022 Cal. Super. LEXIS 95614, *McCarthy v. Century City*, 2023 Cal. Super. LEXIS 19026, *Ramos v. B & V Enters.*, 2019 Cal. Super. LEXIS 13707.⁷⁶

There are several problems with this text. You might immediately notice that many of the cases that the Lexis+ AI Assistant cites are from the California Superior Courts and are of little or no precedential value. And while this GenAI tool might be a good first step for someone unfamiliar with California slip-and-fall law, it does little more than provide a basic,

76. LEXIS+ AI, <https://plus.lexis.com/zhome/?crd=f7c57218-8ffe-484b-b84c-648afe36ee0d> [https://perma.cc/44DD-8LZT] (on Lexis+ AI homepage, click top right corner with AI bubble, then click "ask a legal question"; authors submitted the question, "What are the elements of a slip and fall tort claim in California?") (results may vary).

high-level outline of this nuanced law. To produce content worthy of a motion to dismiss or a motion for summary judgment (or even a legal memo), the user will have to do a lot more work.

Finally, we also asked Westlaw's Practical Law AI the same question and it produced the following response:

The elements of a slip and fall claim, which falls under the umbrella of trespass in California, involve proving that:

- 1) The plaintiff owned or exclusively possessed the real property at issue,
- 2) The defendant entered (or caused a substance or object to enter) the plaintiff's property,
- 3) The entry was intentional or a result of recklessness, negligence, or an ultrahazardous activity,
- 4) The plaintiff did not authorize the entry or the entry exceeded the permission the plaintiff granted, and
- 5) The defendant's conduct caused the plaintiff's harm.

An entry is considered unauthorized if the defendant entered the plaintiff's property without the plaintiff's permission or acted beyond the scope of the permission granted by the plaintiff. The entry can be either a tangible (physical) trespass, such as a person or object entering the property, or an intangible intrusion, such as noise, odor, or light, provided the intrusion causes actual physical damage.⁷⁷

The Practical Law AI response cites Westlaw Practice Notes for original sources and contains the disclaimer: "The above response is AI-

77. WESTLAW PRECISION, [https://1.next.westlaw.com/Conversation/LandingPage/conversation/395890ee-6f0d-4d9c-a223-64e12c8c51b6?contextData=\(sc.Default\)&transitionType=Default](https://1.next.westlaw.com/Conversation/LandingPage/conversation/395890ee-6f0d-4d9c-a223-64e12c8c51b6?contextData=(sc.Default)&transitionType=Default) [https://perma.cc/3ETM-NHS3] (on Westlaw home page, click AI-Assisted Research tab; authors submitted the question "What are the elements of a slip and fall tort claim in California?") (results may vary).

generated and may contain errors. It should be verified for accuracy.”⁷⁸ Practical Law AI makes the user work a little harder than Lexis+ AI Assistant to find the precedent to support its summary of California slip-and-fall law, but the result is much the same.

What do all of the GenAI tools have in common? The texts they provide in response to legal questions provide a start, but nothing more.⁷⁹ Anyone using those GenAI tools in the practice, teaching, or learning of law cannot rely on them alone. They must do more work to verify the accuracy of the AI-generated content and find precedential support for the propositions it advances.

VII. INHERENT ISSUES WITH GENAI

While AI tools like LLMs already have many uses, and we will undoubtedly discover many more, they also have risks, which are inherent to the technology. Some people believe that AI poses an existential risk to the human race.⁸⁰ Essentially, they worry that we will create a super-intelligent AI that will destroy us for its own benefit. Thankfully, at least for the time being, their worries consist entirely of speculation, and so can safely be ignored.⁸¹

Others worry about more concrete risks associated with AI: discrimination, privacy harms, economic harms, and so on.⁸² The promise

78. WESTLAW PRECISION, [https://1.next.westlaw.com/Conversation/LandingPage/conversation/395890ee-6f0d-4d9c-a223-64e12c8c51b6?contextData=\(sc.Default\)&transitionType=Default](https://1.next.westlaw.com/Conversation/LandingPage/conversation/395890ee-6f0d-4d9c-a223-64e12c8c51b6?contextData=(sc.Default)&transitionType=Default) [https://perma.cc/3ETM-NHS3].

79. Our prompt was specific, so we received a decent response from these GenAI products. But GenAI can easily lead even careful users astray. We know of a legal librarian who asked one of these GenAI programs for relevant statutes regarding battery in a specific state, and the program returned statutes and regulations regarding the sale of batteries and battery-containing products. A law professor told us that he asked one of these GenAI products for some non-legal sources for a proposition he was advancing. That product returned two sources that were legitimate and useful, but also misquoted another article, which did not say what the product claimed it did. It also provided a citation to a non-existent journal and provided a citation to an existing journal for a non-existent article. Another law professor told us that she questioned a GenAI program after it pointed her to a hallucinated journal article, and it admitted that it had made up the citation and warned her to carefully check sources.

80. See Kevin Roose, *AI Poses 'Risk of Extinction,' Industry Leaders Warn*, N.Y. TIMES (May 30, 2023), <https://www.nytimes.com/2023/05/30/technology/ai-threat-warning.html>.

81. See *New Study Says AI Poses No Existential Threat to Humanity*, THE ENGINEER (Aug. 13, 2024), <https://www.theengineer.co.uk/content/news/new-study-says-ai-poses-no-existential-threat-to-humanity/> [https://perma.cc/DH8J-8M3W].

82. See, e.g., Yonathan Arbel et al., *Systemic Regulation of Artificial Intelligence*, 56 ARIZ. ST.

and peril of AI are two sides of the same coin. Or rather, the things that make AI technology useful are the same things that make it risky. For example, LLMs work by identifying and reproducing linguistic patterns.⁸³ An LLM is “trained” on a vast corpus of texts to identify the most likely words to appear in response to a prompt.⁸⁴ It then generates those words.⁸⁵ In other words, when you ask an LLM a question or tell it to do something, it does not respond by answering your question or doing what you say. It responds by generating a text that is likely to be responsive to your question or to satisfy your request.

The most important thing to remember about GenAI tools is that they do not really “answer questions” or “respond to prompts.” They produce texts consisting of words in an order that an LLM algorithm predicts is most likely to be responsive to the question or prompt it received. Or rather, to put it another way, a text generated by an LLM is meaningless until it is given meaning by a person who reads it. An LLM cannot generate meaning. It is not designed to generate meaning. It is designed to generate the simulacrum of meaning. We give that simulacrum meaning by recognizing whether it is responsive to our prompt.

A. Hallucinations

GenAI text generators often produce texts that include false statements of fact. In the case of GenAI text generators producing texts in relation to prompts on legal subjects, those false statements of fact typically consist of inaccurate statements of law, incorrect summaries of judicial opinions, or fabricated citations.⁸⁶

These false statements of fact are colloquially referred to as “hallucinations” because they appear to represent the GenAI text generator imagining something that doesn’t exist.⁸⁷ But the use of the term hallucination is at best a hackneyed metaphor and at worst a misleading

L.J. 545, 557–69, 618 (2024) (summarizing the literature on the systemic risks posed by AI and proposing a regulatory approach).

83. *Introducing ChatGPT*, OPENAI (Nov. 30, 2022), <https://openai.com/index/chatgpt/> [<https://perma.cc/PAM6-5M9H>].

84. *Id.*

85. *Id.*

86. *See supra* Section IV.

87. *See What are AI Hallucinations?*, *supra* note 19.

misnomer, because GenAI does not and cannot imagine anything. A GenAI text generator merely produces generic text. Or rather, a GenAI text generator is merely a bullshit generator.

Calling false statements of fact produced by a GenAI text generator “hallucinations” wrongly encourages people to conceptualize what a text generator is doing as thinking, rather than merely producing statistically probable words in sequence. For law students to use GenAI text generators productively, they must internalize the understanding that text generators are not actually answering their questions or responding to their prompts, but producing text that may or may not be accurate or even responsive.

B. Confidentiality

Lawyers have a fiduciary duty of confidentiality to their clients, which prohibits them from disclosing any information that is not publicly available without their client’s consent.⁸⁸ This is a problem because, as discussed, information input into GenAI products is not confidential. In fact, because GenAI products use LLMs to generate output, the product may actually use the information to further “learn” how to write.⁸⁹ Lawyers must understand that the ethical rules, specifically Rule 1.6, prohibit them from using confidential information in GenAI unless that GenAI has sufficient guarantees of confidentiality, which most current GenAI products that are non-proprietary lack.⁹⁰ Thus, lawyers must understand that they cannot put confidential information into GenAI products, and must carefully evaluate whether client-specific information is or is not confidential.

In any case, lawyers should never use an AI tool to evaluate confidential or privileged information without an ironclad guarantee that the AI will maintain absolute confidentiality. Of course, there is always the possibility of a leak. But that is always true. Even manuscripts can be copied and leaked.⁹¹

And yet, concerns about the confidentiality of GenAI tools may be

88. MODEL RULES OF PRO. CONDUCT r. 1.6 (AM. BAR ASS’N 2024).

89. See Owen Wolfe & Eddy Salcedo, *With AI Use, Lawyers Need to Ponder Confidentiality Stipulations*, BLOOMBERG NEWS (July 1, 2024, 3:30 AM), <https://news.bloomberglaw.com/us-law-week/with-ai-use-lawyers-need-to-ponder-confidentiality-stipulations> [<https://perma.cc/28AM-4KKL>].

90. See MODEL RULES OF PRO. CONDUCT r. 1.6 (AM. BAR ASS’N 2024); see e.g., *OpenAI Privacy Policy*, <https://openai.com/policies/privacy-policy/> [<https://perma.cc/SN3B-5ABB>].

91. See, e.g., *Harper & Row v. Nation Enterprises*, 471 U.S. 539, 565 (1985).

misplaced or overblown. Even if LLMs use the texts provided to them as training tools, the LLMs themselves do not contain copies of those texts. They only use those texts to further refine the weights they give the probability of particular words appearing in particular contexts. What's more, even if an LLM is trained on confidential information and then produces what appears to be confidential information in response to a prompt, there is absolutely no guarantee—or even reason to think—that the text produced by the LLM is an accurate representation of the confidential information on which it was trained. It is merely an AI output, with no truth value whatsoever. Further, it would be trivial for parties to introduce false information to an LLM, in order to throw off their opponent. In sum, it should be sanctionable for a litigant to present a text generated by an LLM as fact evidence, because there is no basis for asserting that such a text is factual in nature.

C. Bias

LLMs are trained using human language and thus can reflect human biases. OpenAI, the makers of ChatGPT, have acknowledged that the product can and has generated outputs that reflect bias.⁹² Perhaps surprisingly, LLMs may also reflect political biases.⁹³ And lawyers must understand that. But while bias is a problem with AI, it is also a problem in general. Indeed, if AI is biased, it is only because we are biased, and AI reflects our collective biases.

However, concerns about AI bias may also reflect unreasonable bias against AI.⁹⁴ Yes, AI is biased. It is biased because the data is biased. A biased society necessarily produces biased documents which result in biased AI. But biased in relation to what? Is AI more or less biased than the people whose judgment it supplements or replaces? AI is probably less biased than

92. See *How Should AI Systems Behave, and Who Should Decide?*, OPENAI (Feb. 16, 2023), <https://openai.com/index/how-should-ai-systems-behave/#OpenAI%20> [https://perma.cc/9CH4-W8SW].

93. See Amy Cyphert et al., *AI Cannibalism and the Law*, 22 COLO. TECH. L.J. 301, 304–05 (2024).

94. See, e.g., Andrew Keane Woods, *Robophobia*, 93 U. COLO. L. REV. 51, 114 (2022) (“If they sound absurd, it is because our conversation about the appropriate role for machines in society is inflicted with a fear of and bias against machines.”); Edward Lee & Andrew Moshirnia, *The AI Penalty: Is There a Bias Against AI-Generated Works?*, MICH. ST. L. REV. (forthcoming 2025).

people, because documentary evidence of bias is precisely what people try to avoid when they create the kinds of documents used to train AI models.

And while AI models can be adjusted to reduce bias, and at least in theory, a properly adjusted AI cannot discriminate.⁹⁵ It is easy to prohibit people from discriminating, but it is very hard to actually stop them from discriminating, if they are permitted to exercise any discretionary judgment. While AI bias is a problem to consider and mitigate, AI might actually be a tool for reducing or even eliminating bias, so long as reducing bias is required and expected.

VIII. GENAI PEDAGOGY

In order to ask whether and how AI should be used and taught in law schools, we must identify the purpose of legal pedagogy. Law schools have many purposes. But their most important purpose is training law students to become lawyers. We believe that the purpose of legal pedagogy must be to train law students as efficiently and effectively as possible. We should do things that work, abandon things that do not, and always be looking for ways to make legal pedagogy work better. It does not matter how we teach law students to become good lawyers; it matters that we teach them to be good lawyers by practicing law as efficiently and effectively as possible.

AI is already becoming an essential element of legal practice, and it will surely become an essential element of legal pedagogy.⁹⁶ Law schools are already slowly beginning to include AI in the curriculum, and as a subject of conversation, debate, and speculation.⁹⁷ Some law schools are even beginning to teach their students how and why to use AI tools.⁹⁸ More is sure to come. AI has been an intrinsic but largely invisible part of the

95. See Fawn Fitter & Steven T. Hunt, *How AI Can End Bias*, SAP (July 24, 2024), <https://www.sap.com/resources/how-ai-can-end-bias> [<https://perma.cc/26ZD-GEYU>] (discussing that while AI technology is not always unbiased, rooting out bias is possible).

96. See Jonathan H. Choi et al., *Lawyering in the Age of Artificial Intelligence*, 109 MINN. L. REV. 147, 205–206 (2024). See generally Daniel Schwarcz & Jonathan H. Choi, *AI Tools for Lawyers: A Practical Guide*, 108 MINN. L. REV. 1 (2023).

97. See, e.g., Michael D. Murray, *Artificial Intelligence for Learning the Law: Generative AI for Academic Support in Law Schools and Universities*, 8 J.L. & TECH. TEX. (forthcoming 2025), https://papers.ssm.com/sol3/papers.cfm?abstract_id=4564227 [<https://perma.cc/6SMB-LLVZ>].

98. See, e.g., Michael D. Murray, *Artificial Intelligence and the Practice of Law Part 1: Lawyers Must be Professional and Responsible Supervisors of AI*, KY. BAR ASS'N, BENCH & BAR MAG., Nov.–Dec. 2023, at 36; Michael D. Murray, *Artificial Intelligence and the Practice of Law Part 2: Working With Your New AI Staff Attorney*, KY. BAR ASS'N, BENCH & BAR MAG., Mar.–Apr. 2024, at 40.

technology powering legal practice for decades. Now, it is becoming far more visible, as its uses become more readily apparent.

Of course, AI is not a panacea. It does a lot of things incredibly well but does other things incredibly poorly or not at all.⁹⁹ Part of teaching law students how to use AI is helping them understand how AI works, what it can do, and what it cannot do. Michael Smith has persuasively argued that teaching law students how to use GenAI text generators is less important than teaching them how to plagiarize existing documents.¹⁰⁰ After all, as we have previously observed, “As a practicing lawyer, if you aren’t plagiarizing, you’re committing malpractice.”¹⁰¹ Lawyering consists primarily of copying, and good lawyering consists primarily of knowing what to copy and why to copy it. So, it stands to reason that law schools should teach those skills.

After all, lawyering is the craft of assembling a collection of words intended to convince the court to reach the result the client wants, or at least to get as close as possible. A tried-and-true method of accomplishing that goal is using a collection of words that caused a court to reach a similar result in similar circumstances, ideally with changes reflecting any different circumstances.

GenAI is not particularly useful for this purpose because it produces new texts, rather than copies of texts that have worked in the past. In fact, GenAI may change or even eliminate the very features that made the original effective. So, law students should learn to be careful when using GenAI to produce a document, because it may depart from form documents in undesirable ways.

However, GenAI has uses that plagiarism does not, in large part because it generates new texts rather than copying existing ones. The strength of plagiarism is also its weakness. Plagiarism is good at solving existing legal problems, but bad at solving new ones. If your client needs you to address a familiar problem, then plagiarizing a form document is always the way to go. But if your client needs you to address a novel problem, then plagiarizing a form document is much less likely to be helpful and may even

99. See, e.g., Jonathan H. Choi et al., *ChatGPT Goes to Law School*, 71 J. LEGAL EDUC. 387, 391 (2022).

100. See Michael L. Smith, *Language Models, Plagiarism, and Legal Writing*, 22 U.N.H.L. REV. 361, 365–66 (2024).

101. Boyd & Frye, *supra* note 3, at 1.

steer you in the wrong direction. GenAI may not always provide the best answer to a new problem, but at least it is capable of providing a new answer to a new problem. Of course, lawyers only rarely encounter new problems. But when they do, GenAI could be most useful.

What's more, GenAI may be able to provide plausible arguments that lawyers could not or would not develop on their own. After all, think about how GenAI works. Given a prompt and a corpus of relevant texts, it predicts which words a person would use and in which order they would use those words. Unsurprisingly, the text it produces is typically superficially plausible, though shallow. It generates conventional wisdom because it is designed to generate conventional wisdom.

But there is always the possibility of GenAI producing something else, something people would not produce on their own, maybe even something clever lawyers could not see because it conflicts with conventional wisdom. In other words, there is always the possibility of GenAI producing a text that people, even knowledgeable people, consider "creative."¹⁰² Not because GenAI itself is creative, but because GenAI can perform a reasonable simile of creative thought by asking what it is possible to think, unlike people, who ask themselves what they actually think.¹⁰³ For example, one federal circuit court judge has used AI to help him determine the "ordinary meaning" of a phrase.¹⁰⁴

This is important for legal pedagogy because law students, lawyers, judges, and law professors could use GenAI to produce texts that advance novel ideas, and then claim those ideas for themselves. Currently, the Copyright Office refuses to register a work unless it includes at least one original element that was not created by a machine.¹⁰⁵ But why? Or rather, why not deny registration more broadly? Understandably, authors tend to claim as many property rights as they can. But what about the authors who do not or cannot claim anything? Is their property for the taking? Should it be?

102. Cf. Brian L. Frye, *The Lion, the Bat & the Thermostat: Metaphors on Consciousness*, 5 SAVANNAH L. REV. 13, 40–41 (2018) (observing the compulsion to anthropomorphize).

103. *Id.*

104. *United States v. Deleon*, 116 F.4th 1260, 1270–72 (11th Cir. 2024) (Newsom, J., concurring) (internal quotation marks omitted) (using AI tools to help determine the meaning of the phrase "physically restrained").

105. See *Copyright and Artificial Intelligence*, COPYRIGHT GOV., <https://www.copyright.gov/ai/> [<https://perma.cc/K8WT-ETN8>].

But GenAI also has more prosaic uses. Many students use it to summarize articles or even help them develop an abstract for their own research. They should be allowed to use GenAI in that way, because it helps them learn in ways that other tools cannot. Law schools should not only allow law students to use GenAI tools in that way, but also help law students learn how to use them more efficiently and effectively.

IX. GENAI PEDAGOGICAL CAUTIONS

One critical aspect of any GenAI pedagogy is thorough communication about the current inherent problems with GenAI. The first, of course, is (the inaptly named) hallucinations. Students must be instructed on how to carefully review AI-generated content for misstatements.

Further, as in the example above, students must be reminded that GenAIs, even those developed specifically for the legal field, still do not understand precedential hierarchy. Cases out of the California Superior Courts have no precedential value. And this is not a situation in which no precedent exists—the cases offered by AI Assistant themselves cite precedential cases out of the California Courts of Appeals.

Professors should develop exercises that both show students how GenAI works and ask them to assess and critique GenAI-generated content. One assignment might require students to write a GenAI prompt to answer a client's question. Students should then feed that prompt into several different GenAI tools and assess the output of each, determining the strength of the prompt based on the outputs and assessing both whether the outputs answered the question asked and are legally correct. This will demonstrate for students the “garbage in, garbage out” principle.¹⁰⁶ GenAI tools are only as strong as the lawyers who use them, and learning how to ask the right questions or input the right prompts is essential.

Another assignment might start with a relatively simple inquiry, such as the statute of limitations of a tort claim in a particular jurisdiction. But the activity could also be done with progressively difficult questions, such as the elements of or defenses to a particular claim. Many GenAI tools are relatively proficient at answering simple legal questions, even if the writing

106. See Oksana Zdok, *Why “Garbage In, Garbage Out” Should Be the New Mantra for AI Implementation*, SHELF (May 2, 2024), <https://shelf.io/blog/garbage-in-garbage-out-ai-implementation/> [https://perma.cc/C2X5-Q4VN].

itself cannot be copied and pasted into a legal document, but struggle with more difficult ones. Students could be asked to assess the writing style and organization of GenAI outputs. They could also be asked whether certain outputs could be put into a memo or brief word-for-word, or need to be edited and, if so, in what ways.

Another aspect of GenAI pedagogy involves showing students that even where GenAI tools can provide an accurate summary of relevant law, they cannot make the types of judgment calls that lawyers make every day. To demonstrate this, professors might provide a factual scenario and have students ask GenAI tools whether the defendant should file a motion to dismiss or motion for summary judgment, given the relevant facts. During this exercise, students would need to consider what information can be ethically plugged in to the GenAI prompt (any facts included in a complaint, answer, or counterclaim, and any other facts that are public knowledge through court filings or other publicly available information that are not confidential). Once students determine what information could be used to generate a response, students should be instructed to input that information into various GenAI tools and assess the outputs. They could then be asked whether the tool provided helpful guidance on the question and asked to decide based on their own judgment whether the motion might be successful.

Law students can also use AI productively when writing papers. For example, many students know that they want to write about a particular area of law or subject matter, but do not know specifically what they want to write about. An AI can help them identify specific problems and provide potential arguments. Of course, students will still need to do their own research to determine whether a thesis is actually viable, but an AI can help them get started. If a student provides a thesis, an AI can also provide examples of how the student can structure the argument, which many students find helpful, especially if they are new to legal scholarship. AI can also provide references to statutes, cases, and articles that are likely to be relevant to the student's thesis. This too, though, requires that law professors understand GenAI and can instruct and help students learn to use it in productive and helpful ways.

CONCLUSION

AI is not just coming to the legal profession, it is already here. Lawyers are using AI tools to help them generate legal documents and provide other legal services, and they are going to do it more. That is a good thing, because it makes them more efficient, which is not just good for their clients, but a fiduciary duty they owe their clients. Yes, there are problems related to AI use—but they are not problems caused by AI, they are lawyering problems. All of the problems associated with lawyers using AI disappear as soon as lawyers exercise their fiduciary duty of care to ensure that the advice they give and the documents they produce are accurate. When it comes to the legal profession, there is no need to regulate AI, there is only a need to regulate lawyers, which the profession does more or less adequately.

It seems inevitable that AI will become an essential part of legal practice. The very nature of the profession demands it. AI makes it far cheaper, easier, and faster to produce the kinds of documents clients need lawyers to produce. That does not mean that AI will make lawyers unnecessary. It means that lawyers will be able to automate a mechanical part of their job (writing) and focus their attention on a more important part of their job (critical thinking). After all, good lawyers do not add value by drafting documents. They add value by ensuring the documents say the right things. There is a reason junior lawyers do the drafting and senior lawyers do the editing.

And if AI will become an essential part of legal practice, then law schools have a duty to teach law students how to use it. Just as lawyers have fiduciary duties to their clients, law schools have—or, at least, should have—fiduciary duties to their students. After all, the information asymmetry is similar. Law students want legal education but do not know how to evaluate its quality. Law schools have an obligation to provide a high-quality legal education and to place the interests of their students before their own. Once again, these duties cash out in a duty of efficiency. Law schools have a duty to teach students how to be as efficient as possible and to teach them what they need to know in order to be effective lawyers. That means teaching them how to use AI tools.