

The Legal Ethics of Generative AI—

Part 3

A robot may not injure a lawyer or, through inaction, allow a lawyer to come to harm.

BY COLIN E. MORIARTY



This is the third and final article in a series discussing the legal implications of generative AI. This installment examines ethical considerations for attorneys using generative AI.

The practice of law has marched in step with improvements in technology. The days of searching through a stack of Pacific Reporters in a library have been replaced with inputting queries into online databases. Instead of combing through paper or microfiche catalogs, lawyers can now rifle through recorded documents online. Undoubtedly, this has made practitioners more efficient. But it also creates a danger of losing track of the analog reality that still dictates how the law is published and argued. The organization behind legislation and opinions was developed on paper, and any lawyer who does not understand that system will miss opportunities or expose themselves to embarrassing and costly mistakes. Unlike fictional artificial intelligence, real-life generative AI is not necessarily focused on helping the lawyers who use it avoid harm.¹

The rise of generative AI is the next major technological milestone in the practice of law, promising great advances in efficiency and training. Large language models (LLMs), a form of generative AI with a peculiarly humanlike capacity to interpret and produce human language, appear poised to have the most transformative impact on the practice of law.² Engineers have been trying to create so-called “legal expert systems” to automate the practice of law since the early computer era and have fallen short largely on the difficult problem of language comprehension. With LLMs, this problem may be solvable for the first time. Some lawyers are already using generative AI tools today to help them summarize or understand large documents or sets of documents (such as discovery), conduct legal research, brainstorm ideas, and assist with any number of other tasks.

Generative AI technology cannot be ignored. A lawyer has an ethical duty to understand and

stay abreast of new technology relevant to the practice of law.³ That means understanding not only how the technology can be used but also its risks. Generative AI presents grave dangers for the uninformed, hasty, or lazy.⁴ This concluding piece of a three-part series explores how AI is being incorporated into legal practice and discusses the range of potential risks attorneys face when determining whether and how to use this rapidly evolving technology. Some of those dangers implicate a lawyer’s ethical duties, such as the duties of competence, candor, supervision of nonlawyer assistants, confidentiality, and avoiding discrimination. Other risks are more subtle, such as the potentially corrosive effect on the development of legal reasoning skills and the training and professional development of new lawyers.

Generative AI Is Being Integrated Into the Legal Profession

Engineers have been trying to make robot lawyers for a long time. In the early computer era, some programmers tried to formalize the rules of law into logical statements and then added interpretative software that would allow a nonexpert to plug in the facts and receive the correct legal opinion. These so-called “legal expert systems” were envisioned as a way to replace lawyers with software.⁵ The idea that the law could be reduced to logical statements has long had its critics.⁶ Witnesses and evidence can be misleading or untruthful, relationships and reputation can make a difference in a lawyer’s success, and judges are human beings—not algorithms or equations to be solved.⁷ Even black letter law requires interpretation.⁸ Coming up with formal rules to systematically encode the ambiguous, sometimes messy natural language of legislation into a formal logical system turned out to be extremely difficult.⁹

These legal expert systems largely did not deliver on their promise. A fundamental reason for their failure has been attributed to the prior software's inability to perform the "very mentally demanding task . . . which allows the lawyer to interpret legislation."¹⁰ Even the most commonsense legal reasoning proved intractable.¹¹

Generative AI may bridge this gap. LLMs have had shocking success in mimicking human understanding and production of language. They have accomplished this not by being taught how to encode language directly, but by being fed enormous amounts of written language and being asked to synthesize a map or algorithm that successfully produces language matching what already existed.¹² Through machine learning, the model eventually developed internal logic—as of yet fully undeciphered by programmers—that was effective at this task. The law is largely dominated by the written word as recorded in statutes, motions, briefs, orders, articles, and other sources. This vast corpus of written words is exactly the kind of thing needed to fine-tune an LLM, and the production of written words is exactly what an LLM does.¹³ As with its use in other professions, generative AI could automate certain intelligent tasks that previously could only be done by a human being. In law, for example, these tasks might include proofreading, searching for applicable authority, drafting a memo summarizing the law or facts, or producing timelines or tables of contents for large documents or groups of documents. Rather than requiring a user to consider precise search terms likely to occur in the material being sought, it could allow contextual searches of authority or documents for a particular subject matter or issue. The software is also very good at helping brainstorm ideas to spur the human creative process, such as generating ideas for voir dire, presenting possible questions for depositions, or suggesting counterarguments to a draft brief.

Several LLMs are available to lawyers today, with more on the horizon. Lawyers can access ChatGPT, OpenAI's LLM that kickstarted the current AI boom.¹⁴ It can be used either directly through ChatGPT's website or app, or by using a different software that queries ChatGPT using API calls.¹⁵ Some practitioners are already using

ChatGPT in their practice and are happy to extoll its benefits and offer tips and techniques for its use.¹⁶ Vendors likewise use the ChatGPT API to promote their own applications that are marketed to lawyers. For example, Casetext, a legal research service, released an AI legal assistant named CoCounsel in March 2023.¹⁷ CoCounsel allows users to explain fact patterns to get applicable law and an explanation of the same in response, summarize large groups of documents, organize questions for depositions, and complete other tasks. Lawgeex sells access to an LLM that promises to help with contract review by analyzing legal language "the same way a human lawyer would."¹⁸ And DISCO offers a chatbot named Cecilia that promises to provide "evidence-based answer[s] with citations to documents" in an eDiscovery database.¹⁹

More applications are on the way. Logikcull, an eDiscovery vendor, is preparing to release a generative AI product that will integrate ChatGPT into its systems.²⁰ Logikcull promises that its software will be able to perform context-based searches, such as "find any potential violations" of a statute or "find where Jane Smith's statements show her public statements were false"—a useful enhancement of the current process of brainstorming keywords for a text-based search.²¹ LexisNexis is also working on its own version of an AI legal assistant,²² which may prove very useful given the company's extensive library of statutes, rules, cases, briefs, orders, and secondary sources.²³ In interviews with the author, LexisNexis had predicted that its AI tools would be available on its web-based research service around September 2023. In addition, LexisNexis is working with Microsoft to make its fine-tuned legal models available through other software, such as Microsoft's CoPilot AI. If all goes well, lawyers will, for example, be able to interact with LexisNexis' model from directly within a Word document.²⁴

Lawyers are already using generative AI in their practices in some capacity and will continue to do so. According to a recent Thomson Reuters Institute survey of 440 lawyers, 82% of lawyers surveyed believe that generative AI will be applied to legal work, though only 51% think it should.²⁵ Another survey from LexisNexis found that half of lawyers surveyed had already used

AI in their practice or are planning to do so.²⁶

It may be natural to expect lawyers to adopt generative AI products. After all, in addition to the legal research tools and online court and public records systems mentioned above, lawyers already routinely use cloud-based file storage and case management systems, communicate by email, coordinate on electronic calendars, and use Google or other search engines to seek data online. In fact, most lawyers are likely already using AI, whether they know it or not, in the form of algorithms fueling their legal research systems.²⁷ In between typing "dog/s bite /p ("warning" or "sign" or "trespass" or "notice")" and the display of results of that inquiry by the legal research service of choice, there is a great deal of calculation going on to decide which cases to present and in what order.²⁸ The black box of generative AI may not be much different from the black box of search algorithms when it comes to the practice of law.

Candor and Supervision

In learning how to better mimic the slippery and ambiguous nature of language, and perhaps because of doing so, LLMs are more unpredictable than the other tools a lawyer may use. The largest risks to attorneys using generative AI may be overestimating the capabilities of the software or being overly credulous as to its output. A lawyer's most fundamental duty is competence. Attorneys must possess the knowledge, skill, thoroughness, and preparation reasonably necessary to represent the client, including the ability to understand and properly supervise the tools they use.²⁹ In the use of generative AI, as with any other tool, this requires diligence to avoid errors. An LLM is trained to predict the next word based on the patterns in its training set and on the feedback it received from reinforcement learning.³⁰ This process often results in accurate information. Truth is not guaranteed, however, and LLMs can "hallucinate"—that is, confidently display false information as true.³¹ Just as it would be dangerous to entrust a random nonlawyer with the authority to sign a lawyer's name to pleadings, lawyers who place blind faith in an LLM can face dire consequences. Lawyers, therefore, must scrupulously review generative

AI results before relying on them.

The Colorado Rules of Professional Conduct (the Rules) do not mention generative AI directly, at least not yet. Whether generative AI is viewed more as a neutral research tool or more like a nonlawyer assistant, however, the Rules clearly make the lawyer responsible for ensuring the LLM does not contribute to fraud upon the tribunal. Lawyers have a duty to be truthful in their dealings with the courts and others and a duty to supervise those assisting the lawyer.³² Rule 3.3 states that a lawyer must not knowingly make a false statement of material fact or fail to correct a false statement.³³ This duty further requires the lawyer to disclose authorities known to the lawyer that are directly adverse to the lawyer's position,³⁴ a task that may well be difficult to calibrate using generative AI. Lawyers have a "special obligation to protect a tribunal against . . . fraudulent conduct," including failing to disclose information when required.³⁵ A lawyer also must not knowingly make a false statement of material fact to a third party.³⁶ When an attorney is supervising another lawyer or nonlawyer, the attorney must make reasonable efforts to ensure the other person conforms with the Rules.³⁷ If a nonlawyer assistant commits an act that would violate the Rules, the supervising lawyer is responsible if the lawyer either ordered or ratified the conduct or if the supervising lawyer fails to act to mitigate or take remedial action once the act is known.³⁸

Two recent situations in Colorado and in New York showcase the pitfall of AI hallucinations in court. In the first instance, in Colorado in El Paso County District Court, a young attorney who had been practicing civil litigation for about three months decided to use ChatGPT to conduct legal research in support of a motion to set aside a default judgment.³⁹ He claimed to have seen advertisements from LexisNexis for a new generative AI product, so he used ChatGPT to find cases to support his position to "exponentially augment[]" his fledgling research skills,⁴⁰ and save time. He received some search returns that appeared to be accurate and valid, so he concluded they all were and ended up copying and pasting at least one bogus case from the LLM into his motion.⁴¹ Right before a hearing on the motion, the attorney realized

"my case cites from [C]hatGPT are garbage" and he had "no idea what to do" other than try to find real cases to support his motion.⁴² He failed to correct the error fast enough, however, and the district court judge issued an Order to Show Cause why he should not be sanctioned and reported the attorney to the Office of Regulation Counsel.⁴³

As harsh as it may sound, it could have been

“
The largest risks
to attorneys using
generative AI may
be overestimating
the capabilities
of the software
or being overly
credulous as to
its output.
”

much worse if the attorney had neglected his duties of candor to redress the issue once he discovered it. An even more thorny situation unfolded in New York when two attorneys who had been practicing for at least 25 years placed too much faith in generative AI.⁴⁴ In *Mata v. Avianca, Inc.*, attorney Stephen Schwartz was representing the plaintiff in a personal injury action against an airline.⁴⁵ The defendant removed the case to the Southern District of New York.⁴⁶ Because Schwartz was not admitted

to that court, another attorney, Peter LoDuca, entered as counsel of record while Schwartz continued to do the actual work.⁴⁷ Regrettably, LoDuca did not carefully read or review all of Schwartz's work product.⁴⁸ Also unfortunately, Schwartz, having recently heard of ChaptGPT, decided to fully rely on it to conduct legal research in responding to a motion to dismiss alleging that the complaint was time-barred by the applicable statute of limitations.⁴⁹ Schwartz instructed ChatGPT to "argue that the statute of limitations is tolled" and then followed up by repeatedly asking the LLM to "provide case law," "show me specific holdings," "show me some more cases," and "give me cases."⁵⁰ ChatGPT dutifully provided him with exactly what he asked for, a series of properly formatted citations to cases and statements that their holdings supported Schwartz's position. Schwartz copied them into his brief, even though he could not independently find the cases on his legal research tool.⁵¹ LoDuca then signed the brief, declaring "under penalty of perjury that the foregoing is true and correct."⁵² It was not.

While neither of plaintiff's counsel checked the citations, opposing counsel did, and tactfully highlighted that he "has been unable to locate most of the case law cited" and "the few cases which the undersigned has been able to locate do not stand for the propositions for which they are cited."⁵³ If this warning wasn't clear enough, the district court then issued an order that plaintiff produce copies of the challenged cases.⁵⁴ Despite the clamoring alarm bells, LoDuca merely asked Schwartz to produce the cases, while Schwartz just went back to ChatGPT to ask for assurances that the cases were real.⁵⁵ ChatGPT not only said they were, but actually provided text of the bogus cases, which the plaintiff's attorneys dutifully provided to the court albeit without disclosing how they generated them.⁵⁶

The court was not amused. It set a hearing for plaintiff's counsel to show cause why they should not be sanctioned for the bogus opinions.⁵⁷ This resulted in an affidavit revealing the facts of the situation and their reliance on generative AI, as well as a hearing where, among other things, Schwartz appeared to be confused about what "F.3d" meant,⁵⁸ and LoDuca revealed that he

had misled the court about going on vacation to induce the court to grant him an extension of time, with the effect of concealing Schwartz's role in the case.⁵⁹ Based on these facts and on the garbled text of the bogus cases, the court concluded that none of the plaintiff's lawyers had actually read the cases they were submitting.⁶⁰

Sanctions followed. The court explained that, while there is nothing inherently improper in using generative AI, attorneys have "a gate-keeping role" to supervise junior attorneys and technology.⁶¹ Attorneys Schwartz and LoDuca had "abandoned their responsibilities" by submitting the nonexistent legal citations with fake opinions.⁶² The court found that the plaintiff's attorneys abused the judicial system and violated Federal Rule of Civil Procedure 11 by filing papers "without taking the necessary care in their preparation" because the rule requires counsel to undertake a reasonable inquiry into the viability of a pleading before it is signed.⁶³ The attorneys also violated the New York Rules of Professional Conduct by making a false statement to the tribunal and failing to correct the false statement.⁶⁴ The court stopped short of finding that the lawyers had committed criminal forgery because they had not actually forged a signature or seal.⁶⁵ Determining that the lawyers had acted in bad faith, the court ordered them to pay a \$5,000 sanction and to inform both the plaintiff and the judges who had been credited as authoring the gibberish cases of what had occurred.⁶⁶

Though the lawyers in these two examples made different choices about how to handle the errors once known, in both cases the lawyers claimed to have lacked knowledge of the LLM's initial errors. While some of Colorado's Rules of Professional Conduct only prohibit "knowing" misstatements, lawyers almost certainly have an affirmative duty to prevent misstatements by generative AI in work product under Rule 3.3(a) alone. To the extent that the AI could be seen as a nonlawyer assistant, the Rules require the lawyer not just to correct errors once discovered, but also to "make reasonable efforts to ensure" conformance with the Rules.⁶⁷ For court filings actually signed by an attorney, the attorney has a duty to conduct a reasonable

inquiry into any filing they sign and determine that it is well grounded in fact, warranted by existing law or a good faith extension of the same, and not interposed for any improper purpose.⁶⁸ Interestingly, the rule requiring truthful statements to third parties other than the client or the court does not appear to contain an affirmative duty, although it does generally prohibit making a false statement of material fact or law.⁶⁹ It is probably still wise to take an active role in preventing generative AI from making false statements to third parties, though, since a lawyer may not "engage" in dishonesty, fraud, or misrepresentation.⁷⁰

A lawyer also must not "engage" in conduct that is prejudicial to the administration of justice.⁷¹ As mentioned by the judge in the *Mata* case, polluting the court record with fictitious court opinions is prejudicial:

Many harms flow from the submission of fake opinions. The opposing party wastes time and money in exposing the deception. The Court's time is taken from other important endeavors. The client may be deprived of arguments based on authentic judicial precedents. There is 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. It promotes cynicism about the legal profession and the American judicial system. And a future litigant may be tempted to defy a judicial ruling by disingenuously claiming doubt about its authenticity.⁷²

Famed 18th-century satirist Johnathan Swift noted that "[f]alsehood flies, and truth comes limping after it. . . ."⁷³ Copying and pasting bogus cases from generative AI without checking if they are legitimate is quick and easy. If that is done and falsehood is allowed to fly, correcting the error becomes significantly harder. Often the time it takes to analyze and expose falsehoods far outweighs the effort needed to commit the initial fraud, wasting judicial resources. And drafting affidavits, motions to withdraw, or corrections is time-consuming for the lawyer. The waste is particularly bad since it should be at least as quick and easy for the originating attorney to copy and paste a citation provided

by generative AI into an authoritative legal research service to check it as to copy it into a court filing.

More insidiously, the injection of false citations into the legal system risks causing perpetual confusion into the future. Once a bogus citation has wormed its way into the court record, it risks being described in motions and orders that may, in turn, be used to fine-tune generative AI on legal cases. Such citations might confuse or mislead LLMs or even less diligent human researchers. Any pollution of the legal corpus in these early, wild, and woolly days of generative AI may thus compound future problems.⁷⁴

The judiciary has taken notice of the threat. Judge Baylson of the US District Court for the Eastern District of Pennsylvania requires that if any litigant uses AI in the preparation of a paper filed with the court, they "**MUST**, in a clear and plain factual statement, disclose that AI has been used" and "**CERTIFY**, that each and every citation to the law . . . has been verified as accurate."⁷⁵ Judge Fuentes of the US District Court for the Northern District of Illinois likewise requires litigants to "disclose in the filing that AI was used and the specific AI tool that was used to conduct legal research and/or to draft the document."⁷⁶ The US Court of International Trade requires a similar disclosure.⁷⁷ Judge Starr of the US District Court for the District of Texas has entered a standing order, with a template certificate, requiring all litigants to file an attestation along with their notice of appearance stating that "no portion of any filing will be drafted by generative artificial intelligence" or that anything so drafted "will be checked for accuracy, using print reporters or traditional legal databases, by a human being."⁷⁸ Judge Starr explained:

While attorneys swear an oath to set aside their personal prejudices, biases, and beliefs to faithfully uphold the law and represent their clients, generative artificial intelligence is the product of programming devised by humans who did not have to swear such an oath. As such, these systems hold no allegiance to any client, the rule of law, or the laws and Constitution of the United States (or, as addressed above, the truth).

Unbound by any sense of duty, honor, or justice, such programs act according to computer code rather than conviction, based on programming rather than principle.⁷⁹

Lawyers have been using algorithms to assist with legal research and writing for decades, so why is this warning necessary? The reason may be the way human lawyers interact with the software (i.e., the subjective user experience when interacting with an LLM versus a traditional search engine). When formulating a search query employing “AND,” “OR,” “/s,” and nested parentheses, most users understand they’re interacting with a digital database and will be getting results based solely on the terms provided. Even when using a natural language search, the modern user probably knows that the algorithms’ results are merely the mechanical production of an algorithm and should be treated accordingly. In other words, users know they have to vet the results. Users of an LLM, by contrast, can be lulled into thinking they’re dealing with something with human-like intelligence instead of an algorithm.⁸⁰ Instead of simply producing a list of search results, most LLM-fueled applications also offer up summaries, discussions, or other human-like explanations, inviting the user to anthropomorphize in a way that would not happen in a traditional search.

That, however, is not the whole answer. Just because an LLM can sound like a human being does not mean that a lawyer will trust it. Lawyers are accustomed to humans providing bad information or incorrect results. Generative AI may appear more trustworthy than it actually is, not just because it sounds human, but also because it borrows (undeserved) indicia of reliability that lawyers would use to help guide their confidence in the answers of another human. In the case of a human assistant, there would be other clues that could corroborate reliable and truthful citation of legal authority. An assistant’s experience, reputation, writing ability, confidence level, nonverbal communications, and other considerations would offer a great deal of data that may correlate with a trustworthy answer. An LLM can display some of these same indicia of reliability, such as properly written prose, convincing arguments, or confident statements, and so suggest that a

lawyer should drop their guard. In the case of the unfortunate Colorado attorney discussed above, he received some correct answers from the software in well-written prose and interpreted them as indicia that the software was trustworthy in all responses—in other words, a type of confirmation bias in favor of the LLM. Ultimately, any lawyer using generative AI must remember that no matter how eloquent or convincing the software seems, it may be less

“
Generative AI can
help locate the
primary sources
that matter, but it
should never be
used as a primary
source itself.
”

trustworthy than a conventional legal search algorithm, and far less accountable. Simply put, the lawyer must review every bit of an LLM’s work product.


A lawyer’s best defense against AI-assisted violations of the Colorado Rules of Professional Conduct or CRCP 11 is keeping firmly in mind that *generative AI is a secondary source, at best*. In the law, primary sources include court decisions, statutes, regulations, and certain other official documents produced by the government. Generative AI can help locate the primary sources that matter, but it should never be used as a primary source itself. Despite being

dressed up in confident and well-written prose, generative AI’s citations should be treated no differently than the results of any other legal research system. Lawyers should check citations and independently consider the reasoning and conclusions based on the lawyer’s own review of the primary sources. Understanding how the law is recorded and organized on paper is critical to this endeavor.

Over time, AI tools offered to lawyers probably will become more reliable and less likely to hallucinate. Providers are using controls at various levels to minimize incorrect information. One approach, used by Casetext and by LexisNexis, involves fine-tuning the instantiations of the LLM legal data from their legal databases.⁸¹ Fine-tuning is the process of retraining the model on a new dataset. The original LLM may have been trained on a great deal of bad, incorrect, or simply irrelevant information. Theoretically, this means the new model is better at predicting text like the legal documents it reads.⁸² Another way to control misbehaving LLMs is to put software layers between the LLM and the user that force the generative AI to produce specific kinds of output. For example, when a lawyer asks CoCounsel a legal research question, software first prompts the LLM to respond with a restatement of the question to ensure it was properly understood. Then, for CoCounsel and for the system being developed by LexisNexis, other software code appears to prompt the LLM to run searches using the legal database provider’s existing search systems, analyze the results, and iterate until it has located appropriate cases. Finally, the LLM returns an answer that lists citations to specific authority, perhaps with hyperlinks, for easier verification. All of these controls likely make generative AI more useful and probably more accurate. However, none of them eliminates the need for human lawyers to exercise their own judgment.⁸³ After all, it is the lawyer who will be sanctioned for an error, not the LLM.

Confidentiality

Another major ethical risk for attorneys using LLMs is the danger of violating client confidentiality or privilege. Unless a law firm has paid to develop and host its own LLM in-house,



prompting an LLM involves sending that prompt and any client-related information in it out into the Internet to a third-party provider. Several aspects of the technology make prompts more likely to include sensitive client information than a garden-variety Boolean search. First, LLMs are prompted using natural language in a manner similar to how a lawyer would ask another lawyer for a research result. So, instead of the lawyer internally translating a client's dog bite situation into terms such as

"dog" /s "bite" /s "statute of limitations," the lawyer might write something descriptive such as "What is the statute of limitations for a negligence claim against the owner of a dog who bites a visitor in Colorado?" Second, part of the competitive advantage of using an LLM in the first place may be to augment the attorney's own reasoning skills, but to do that requires providing the LLM with more factual information so that it can accomplish this task. For example, a prompt might read, "John Smith

was bitten by a dog while visiting his friend Jane Doe in Denver, Colorado. Please identify the possible causes of action available to Mr. Smith and the statute of limitations for each and citations to authority." By its very nature, the LLM invites the attorney to provide richer and more detailed information about a client to leverage the technology.

The information in a prompt could potentially be used by the LLM vendor and by the machinery of the LLM itself, which may be

constantly updating its model by ingesting the text it receives from users.⁸⁴ In other words, the model may change itself to learn to associate the pattern of words making up the prompt and thus be more likely to produce that same pattern to other users. If that pattern of words reveals client information, this has implications for privileges and confidentiality.

Attorney-client privilege is an evidentiary rule that protects all communications for the purpose of obtaining legal advice.⁸⁵ Doing so ensures that clients can frankly discuss their issue and all pertinent facts with their lawyer no matter how personal or inculcating those facts may be, so that the lawyer can provide legal counsel based on the totality of the facts.⁸⁶ Normally, if an attorney-client communication is made in the presence of a third party, privilege is waived.⁸⁷ There are exceptions to this waiver. The attorney-client privilege expressly covers communications to a lawyer's "secretary, paralegal, legal assistant, stenographer, or clerk."⁸⁸ The Colorado Supreme Court has held that the privilege also extends to consulting experts acting as agents of the attorney⁸⁹ and to agents of a client who communicate with the attorney.⁹⁰

Client information can also be protected by the work product privilege. The work product privilege protects against disclosure of materials prepared in anticipation of litigation or trial by a party or the party's representatives, "including the party's attorney, consultant, surety, indemnitor, insurer, or agent."⁹¹ It also protects "against disclosure of the mental impressions, conclusions, opinions, or legal theories of an attorney or other representative of a party concerning the litigation."⁹²

Colorado's Rules of Professional Conduct likewise require that lawyers protect client confidentiality.⁹³ Under the Rules, a lawyer must not "reveal information related to the representation of a client" except as authorized by the client, as authorized to carry out the representation, or in other specific cases.⁹⁴ This duty of confidentiality likely includes protecting a client's privileges but is broader than that.⁹⁵ It applies "not only to matters communicated in confidence by the client but also to all information relating to the representation, whatever its source."⁹⁶ Of course, lawyers are "impliedly authorized to make

disclosures about a client when appropriate in carrying out the representation," such as by furthering the client's interests in the matter of the representation.⁹⁷ The lawyer must, however, "make reasonable efforts to prevent the inadvertent or unauthorized disclosure" of—or unauthorized access to—such information.⁹⁸

Any prompt provided to generative AI for the purpose of working on a client matter may, by definition, be "related" to the representation. Prompts given to generative AI to conduct research or produce work product related to a client matter could involve information protected by attorney-client or work product privilege, such as facts or questions provided by the client in confidence for the purpose of obtaining legal advice, or the lawyer's own thought process, ideas, and theories. When this kind of information is used to prompt an LLM, is the privilege waived by disclosure to a third party? Is client confidentiality violated?

Whether or not this disclosure to an LLM vendor waives attorney-client or work product privileges is unclear. Colorado courts have not yet weighed in, but it probably should not waive the privilege.⁹⁹ After all, legal research generally falls within the reaches of the privilege.¹⁰⁰ Commentators note that the complexity of modern existence prevents attorneys from handling client affairs without the help of nonlawyer digital tools.¹⁰¹ Thus, any tool used to conduct research necessary to providing legal advice is privileged, or so the argument goes.¹⁰² A contrary result would cause massive disruption to how law is practiced in the modern age, where every attorney uses third-party research services such as LexisNexis, Westlaw, and Fastcase to conduct day-to-day research. Still, the privilege is normally lost if the holder, by words or conduct, expressly or impliedly forsakes confidentiality of that information.¹⁰³ Whenever a lawyer entrusts information to an online vendor, there is a danger of waiving protection.¹⁰⁴ If the vendor is not an "agent" of the attorney as the term is used under cases discussing the work product and attorney-client privilege, it may not be clear whether the privilege is preserved. It may be that future decisions in this area will be informed by the better-developed considerations related to client confidentiality.

It is at least clear that electronic communication does not, by itself, violate the duty of confidentiality. The Colorado Bar Association (CBA) has opined that a lawyer's use of electronic communications do not per se violate confidentiality.¹⁰⁵ The American Bar Association (ABA) and many other states have similarly held.¹⁰⁶ When an attorney communicates electronically, however, the lawyer "must take reasonable precautions to prevent the information from coming into the hands of unintended recipients."¹⁰⁷ The level of precautions depends on "if the method of communication affords a reasonable expectation of privacy."¹⁰⁸ For well-established forms of electronic communication already in widespread use, like email and smartphone messaging, the CBA opines that email and smartphones both afford such a reasonable expectation of privacy and so normally do not require special precautions.¹⁰⁹ That said, there are certain "common procedures and safeguards" that all lawyers should consider.¹¹⁰ These include a documented cybersecurity plan, periodic inspection for signs of cyberattack or data theft, and basic cybersecurity measures such as installing firewall and antivirus software, keeping software updated, using strong passwords, and training staff on cybersecurity best practices.¹¹¹

The ABA and many states have also concluded that cloud storage is acceptable so long as reasonable care is exercised and the attorney is knowledgeable about the risks when selecting a third-party vendor.¹¹² Lawyers who provide client information to third parties for cloud storage, copying, or other online services "must make reasonable efforts to ensure that the services are provided in a manner that is compatible with the lawyer's professional obligations."¹¹³ The scope of the lawyer's need to evaluate third-party document storage is fact-specific and depends on "the education, experience and reputation of the nonlawyer; the nature of the services involved; the terms of any arrangements concerning the protection of client information; and the legal and ethical environments of the jurisdictions in which the services will be performed, particularly with regard to confidentiality."¹¹⁴

If these principles are applied to LLMs, then confidentiality in prompts or responses will depend on whether the lawyer took reasonable steps to preserve confidentiality. This should include, at a minimum, analysis of the overt statements by the vendor about how they treat information and some consideration of how reliably the vendor will be able to uphold confidentiality. Lawyers should not engage with generative AI or any other online service without first reading the clickwrap terms and conditions very carefully.¹¹⁵ Some online terms and conditions allow the vendor to scan all information provided and use it for advertising or other purposes, with uncertain implications for waiving confidentiality or privilege.¹¹⁶

In the case of ChatGPT specifically, OpenAI's terms and conditions state that it may use content provided by its users "to help develop and improve" the model unless the content is accessed via API or the user fills out a specific form.¹¹⁷ If a lawyer accesses ChatGPT using its default web interface, therefore, the information is going to be ingested into the model and used for training purposes. API access, by contrast, refers to using third-party software to send HTTP requests to ChatGPT.¹¹⁸ So, other vendors like Casetext that may wish to take advantage of OpenAI's LLM can use the API to send prompts and this, in theory, requires OpenAI to maintain confidentiality.¹¹⁹ Thus, any lawyer interested in using ChatGPT would be well advised to limit their use to API calls only.¹²⁰ Similarly, before using any other service claiming to be powered by an LLM, the lawyer should carefully scrutinize the terms of use to ensure that the vendor is required to maintain confidentiality.

Legal protection of data might also contribute to a reasonable expectation of confidentiality. The Stored Communications Act prohibits unauthorized access or exceeding authorized access to a facility through which electronic communication is provided.¹²¹ The provider of electronic communication is prohibited from knowingly divulging the content or subscriber information of any stored communication except under certain conditions.¹²² A subpoena for discovery does not normally fall into these exceptions.¹²³ Electronic communication, under the Act, means "any transfer of signs, signals,

“
It may take some
time before the
effectiveness of
OpenAI's security
is known. This
is particularly
important given
that other third-
party vendors,
including those
marketing services
to attorneys, may
be using fine-tuned
models ultimately
operated by
OpenAI.
”

writing, images, sounds, data, or intelligence of any nature transmitted in whole or in part by a wire, radio, electromagnetic, photoelectronic or photooptical system that affects interstate or foreign commerce," which may be broad enough to encompass prompts sent to an LLM vendor for the purpose of communicating a query to the model.¹²⁴ This, and similar anti-wiretapping statutes, may further support an expectation of confidentiality.

Even if digital legal research does not automatically waive privilege or confidentiality, however, the recipient's mistakes might. Well-established vendors with a good track record of maintaining confidentiality might require less scrutiny than new vendors at the bleeding edge of technology, like those offering LLMs. Even sophisticated users of the Internet may be unable to effectively keep online activity private due to constantly changing technology used by online researchers and trackers.¹²⁵ Whether sending prompts to generative AI breaches the privilege thus may be dependent on whether the third-party vendor receiving the prompts can be reasonably relied on to keep them confidential. Data breaches are a constant threat online, and OpenAI has suffered a few since it began offering its LLM for use.¹²⁶ These breaches, or general uncertainty about OpenAI's ability to keep user content from being disclosed or incorporated into the model, appear to have prompted the Federal Trade Commission to open an investigation.¹²⁷ It may take some time before the effectiveness of OpenAI's security is known. This is particularly important given that other third-party vendors, including those marketing services to attorneys, may be using fine-tuned models ultimately operated by OpenAI.

Discrimination

Generative AI may also implicate Colorado's rules against discrimination in the legal profession. Lawyers in Colorado may not, in the representation of a client, engage in conduct that exhibits or is intended to appeal to or engender bias against a person based on race, gender, religion, national origin, disability, age, sexual orientation, or socioeconomic status.¹²⁸ So far, attorneys have been disciplined under Rule

8.4(g) for making an anti-gay slur in communication with a client,¹²⁹ making a misogynist slur in communication with a district attorney,¹³⁰ and making derogatory slurs exhibiting bias on the basis of a judge's gender.¹³¹ It is not necessary that the attorney actually harbor any bias to violate this rule because it "only addresses the attorney's outward behavior . . ."¹³²

Insofar as Rule 8.4(g) primarily seems to be applied to the overt use of slurs attacking protected classes, large commercial LLMs might not pose much of a risk, at least not an obvious one. The providers of this service appear to be actively working to prevent and block overt foulmouthed or toxic behavior through training, human feedback, or filters.¹³³ Still, these techniques probably will not be perfect. Researchers have shown that toxic behavior can still be generated using certain kinds of

prompt injection, such as convincing the LLM to adopt a persona¹³⁴ or more sophisticated techniques.¹³⁵ And, of course, local models could be created that lack the safeguards of the large commercial versions. Still, it seems that a lawyer is unlikely to encounter—and wrongly use—an overt racial slur through proper use of a large commercial LLM. If AI produced content including an overt racial slur, presumably it would be easy enough for a lawyer who was properly supervising and reviewing the work product to remove it.

For LLMs, though, the bias problem may be more subtle than the kind of overt insults that have triggered action under Rule 8.4(g). Models will absorb whatever biases exist in their training data and may learn to make associations common in written language even if they are discriminatory.¹³⁶ Similarly,

unbalanced questions suggesting one view can produce biased answers.¹³⁷ The resulting answers may be incorrect or incomplete if they are based on stereotypes or if they pick up on a bias inherent in the way the question was asked.¹³⁸

Professional Concerns

Apart from the practical dangers and ethical questions involved in using generative AI, there are deeper questions about what effect its widespread adoption can have on the legal profession. Adopting LLMs without care also risks corroding the profession from the inside. Using generative AI is the next big step away from the pages, books, and libraries that formed the basis for legal research and reasoning. Like past steps such as moving online for legal research, generative AI probably makes



Licensed Lawyer™

We've added more categories and eased searchability so that your practice is more visible to a wider range of clients.

Be sure to log into LicensedLawyer.org/CO using your FindALawyer password to update your bio. Email membership@cobar.org or call 303-860-1115, ext. 1 for questions and support.

CBA
Est. in 1897
Colorado Bar Association

individual lawyers more productive and more efficient. But generative AI increases the likelihood that practitioners operate without a firm understanding of how to find and use legitimate legal authority. When the unfortunate New York lawyer told the judge he did not know what “E:3d” stood for,¹³⁹ it is entirely possible he was telling the truth. Perhaps his practice simply never required him to know what “E:3d” meant. When the secondary sources and tools used to search the law seemingly become sufficient on their own, oversights and mistakes occur.

This is not unique to generative AI. More familiar online research systems pose dangers as well. For example, a legal research tool may report that a case has “not yet been released for publication,” but that does not mean it was not selected for publication.¹⁴⁰ It may simply be too new. LexisNexis has recently been incorporating Colorado state court filings into its database, which include proposed orders and short form orders that often appear as the first pages of the filing, with the original motion following. Sometimes, a judge will stamp “GRANTED” or “DENIED” on the proposed order submitted with the motion. When these kinds of pleadings are added to an online database, all of the text that may or may not have actually been part of the judges’ order can come along for the ride. When these documents are viewed in their original context, it should be obvious that a proposed order was not signed or that the text included in the order was not meant to be part of the ruling. When reduced to plain text in a search engine, though, it is easy to miss these distinctions.¹⁴¹ The use of generative AI seems likely to introduce yet more shortcuts and thus more pitfalls for those who use it without a full understanding of where the shortcut is taking them.

The temptation to rely on generative AI will probably be greatest among the newest attorneys and those seeking to cut corners. At the moment, LLMs do not appear competitive with seasoned lawyers who have already had their brains fine-tuned in their field of practice. But LLMs may be competitive with the inexperienced. The day is likely not far off when a law firm may be tempted to replace junior associates with generative AI since the work

product will be reviewed by a senior attorney in either event, and an attorney costs more. Similarly, generative AI may be attractive to new attorneys seeking to punch above their weight class. In both cases, in addition to the ethical dangers of relaxing the lawyer’s duty of competence, a more subtle problem exists in the form of reducing the incentive for experienced lawyers to teach and for new lawyers to learn.

This is more than mere objection to a new technology. Whereas online research diminished the importance of knowing how to research an issue in physical libraries and print resources, AI threatens to diminish the importance of legal reasoning and writing itself. Lawyers who do not hone these skills give up their own futures. A firm that neglects training the next generation of lawyers in favor of automating the entry level may find itself with nothing but entry-level skillsets without the capacity to ensure that generative AI is employed in compliance with the attorney’s duty of competence. A new associate who simply prompts an LLM does not fine-tune their own brain to be able to master and build on the law.

Simply put, the better generative AI becomes, the more lawyers and firms will have to make a conscious effort to teach young lawyers the skills and facts related to the profession to ensure proper professional development and avoid

the temptation to seek short-term savings over long-term benefits of developing practice, talent, and reputation. Fortunately, the legal profession has unique safeguards to encourage this effort and hold accountable those who seek to use generative AI as a shortcut to greater short-term productivity. Lawyers exist in an adversarial space, where their work is open to scrutiny by clients, judges, regulators, and opposing counsel. Particularly in litigation, lawyers have a vested interest in exposing any oversights or problems with opposing counsel’s work product. If a firm or new lawyer relies on generative AI to replace, instead of supplement, human reasoning, it will likely suffer for this decision at some point as adverse lawyers take notice.

Conclusion

Generative AI will and maybe should be used by firms to assist in the practice of law. But it should be firmly regarded as a secondary source and a tool, not an end-all-be-all. We must not buy into the idea that AI has entirely eliminated the need for human curiosity, creativity, insight, and oversight in the legal profession. Perhaps there will come a time when generative AI is built into software that does replicate a lawyer’s competence. That day has not yet arrived. Until it does, our unique, human competence is our own privilege and responsibility. CI



Colin E. Moriarty practices with Underhill Law, P.C. in Greenwood Village. Focusing on business and commercial litigation and arbitration, he has litigated business disputes, construction and fabrication defect claims, employment discrimination lawsuits, subcontractor litigation, state RICO fraud lawsuits, civil theft disputes, insurance appraisal and adjustment disputes, and other lawsuits involving complex commercial and construction matters—colin@underhilllaw.com. The author thanks Joseph Michaels, the coordinating editor for the Professional Conduct and Legal Ethics column, for his review of this article.

Coordinating Editors: K Kalan, kmkalan@yahoo.com; William P. Vobach, bill@vobachiplaw.com; Joseph Michaels, joseph.michaels@coag.gov

NOTES

1. Asimov, *I, Robot* (Gnome Press 1950) (Asimov’s first law of robotics states: “A robot may not injure a human being or, through inaction, allow a human being to come to harm.”).
2. Ziffer, “The Robots are Coming: AI Large Language Models and the Legal Profession,” ABA (Feb. 28, 2023), <https://www.americanbar.org/groups/litigation/committees/products-liability/practice/2023/the-robots-are-coming>.
3. See Colo. RPC 1.1, cmt. [8].
4. This article is the third in a three-part series concerning generative AI. The first two parts discussed the basics of how the software operates and the general risks when used in business, and because this article will build on those concepts, we suggest reading the first two articles before reading this one. Moriarty, “The Legal Challenges of Generative AI—Part 1: Skynet and HAL Walk Into a Courtroom,” 52 *Colo. Law.* 40 (July/Aug. 2023), <https://cl.cobar.org/features/>

the-legal-challenges-of-generative-ai-part-1; Moriarty, “The Legality of Generative AI—Part 2: I’m sorry, User. I’m afraid I can’t do that.” 52 *Colo. Law.* 30 (Sept. 2023), <https://cl.cobar.org/features/the-legality-of-generative-ai-part-2>.

5. Greenleaf, “Legal Expert Systems—Robot Lawyers?” at 6 (presented at the Austl. Legal Convention, Sydney, Austl., Aug. 1989), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2263868.

6. Leith, “Fundamental Errors in Legal Logic Programming,” 29 *Computer J.* 545–52 (Jan. 1986).

7. Leith, “The Rise and Fall of the Legal Expert System,” 1 *Eur. J. of L. & Tech.* (2020), <https://ejlt.org/index.php/ejlt/article/view/14/1>.

8. Leith, *supra* note 6 at 545–46.

9. See Schweighofer and Winiwarter, “Intelligent Information Retrieval: KONTERM—Automatic Representation of Context Related Terms Within a Knowledge Base for a Legal Expert System,” Proc. 25th Anniversary Conf. of the Instit. (1994), <https://citeseerx.ist.psu.edu/doc/10.1.1.22.4751> (describing one attempt to overcome this problem).

10. Leith, *supra* note 6 at 546.

11. Franklin, “Discussion Paper: How Much of Commonsense and Legal Reasoning Is Formalizable? A Review of Conceptual Obstacles,” 11 *L., Probability and Risk* 225, 225 (June–Sept. 2012), <https://academic.oup.com/lpr/article/11/2-3/225/916300>.

12. Moriarty, “The Legal Challenges of Generative AI—Part 1,” *supra* note 4 at 42.

13. Elwany et al., “BERT Goes to Law School: Quantifying the Competitive Advantage of Access to Large Legal Corpora in Contract Understanding” (Nov. 1, 2019), <https://arxiv.org/abs/1911.00473>.

14. Many examples in this article discuss ChatGPT because it was the first and most famous of the current generation of LLMs. The questions raised should be the same for other LLMs, but the analysis might not be. Each LLM is a complicated web of relationships between prompts and output data born from the specific initial data and reinforcement learning. It is not yet known exactly what algorithms each LLM develops internally to excel at manipulating language, so there is no particular reason to conclude that each trained LLM is doing the same thing as another. At the moment, it appears possible that each different model used in generative AI may end up with particular quirks or differences.

15. “API” stands for “application programming interface” and generally means a hook that one piece of software uses to interact with another piece of software. See, e.g., “What is an API?,” IBM, <https://www.ibm.com/topics/api>. In this context, a ChatGPT API means giving software written by others the ability to prompt and receive responses from the ChatGPT model.

16. Attorney Enrico Schaefer is enthusiastic about his firm’s use of ChatGPT and has made a series of instructional videos on how he implements the same into his workflow. Schaefer, “Traverse AI,” <https://www.youtube.com/@TraverseAI/videos>. Some cloud-based

case management software companies are also encouraging lawyers to use ChatGPT. Barkved, “6 ChatGPT Prompts for Lawyers,” Clio (blog) (July 11, 2023), <https://www.clio.com/blog/chatgpt-prompts>. Spiegel, “ChatGPT for Lawyers: Everything You Need to Know,” Smokeball (blog) (Apr. 26, 2023), <https://www.smokeball.com/blog/chatgpt-for-lawyers-everything-you-need-to-know>.

17. Ambrogi, “Casetext Launches Co-Counsel, Its Open-AI Based ‘Legal Assistant’ to Help Lawyers Search Data, Review Documents, Draft Memos, Analyze Contracts, and More,” LawSites (Mar. 1, 2023), <https://www.lawnext.com/2023/03/casetext-launches-co-counsel-its-openai-based-legal-assistant-to-help-lawyers-search-data-review-documents-draft-memos-analyze-contracts-and-more.html>. In a likely not-unrelated move, Thompson Reuters, owner of Westlaw, recently purchased Casetext. See Reuters, “Thompson Reuters to Acquire Legal AI Firm Casetext for \$650 million” (June 27, 2023), <https://www.reuters.com/markets/deals/thomson-reuters-acquire-legal-tech-provider-casetext-650-mln-2023-06-27>.

18. <https://www.lawgeex.com/platform/managed-ai>.

19. <https://www.csdisco.com/offerings/ediscovery/cecilia>.

20. <https://www.logikcull.com/blog/ai-powered-document-review>.

21. *Id.*

22. LexisNexis, “LexisNexis Announces Launch of Lexis+ AI Commercial Preview, Most Comprehensive Global Legal Generative AI Platform” (May 4, 2023), <https://www.lexisnexis.com/community/pressroom/b/news/posts/lexisnexis-announces-launch-of-lexis-ai-commercial-preview-most-comprehensive-global-legal-generative-ai-platform>.

23. LexisNexis, “Finding Briefs, Pleadings, and Motions,” https://supportcenter.lexisnexis.com/app/answers/answer_view/a_id/1096234/-/finding-briefs%2C-pleadings%2C-and-motions. As this article discusses, the fact that Lexis compiles motions and proposed orders, not just actual orders, is a potential pitfall for users who may be moving too quickly to notice the authority they are citing was rejected (not accepted) by a prior court.

24. As of the date of this article, the pricing of LexisNexis’s proposed offering is uncertain. However, a user wishing to integrate LexisNexis AI with Microsoft products may need to have subscriptions to (1) the latest version of Microsoft 365 (a challenge for larger businesses and governments using older versions of the software); (2) LexisNexis, including its AI tools; (3) Microsoft Co-Pilot; and (4) an add-on to integrate LexisNexis AI with Microsoft.

25. Thomson Reuters, *ChatGPT and Generative AI Within Law Firms* 7–8 (2023), <https://www.thomsonreuters.com/en-us/posts/wp-content/uploads/sites/20/2023/04/2023-Chat-GPT-Generative-AI-in-Law-Firms.pdf>. The survey included 443 responses from midsize and large law firms, and from the Thompson Reuters Influencer Coalition Panel in the United States, United Kingdom, and Canada. *Id.* at 5.

26. LexisNexis, “Generative AI Captures Imagination of Lawyers, Law Students, Consumers Alike” (Mar. 20, 2023), <https://www.lexisnexis.com/community/pressroom/b/news/posts/generative-ai-captures-imagination-of-lawyers-law-students-consumers-alike>.

27. Yao, “Mysterious Search Algorithms,” *Library Innovation Lab* (May 24, 2023), <https://lil.law.harvard.edu/blog/2023/05/24/mysterious-search-algorithms>.

28. *Id.* For example, Westlaw uses a set of vertical search engines, each tuned to one or more content types so that the criteria for a good case is different from that for a good statute or regulation. Mart et al., “Inside the Black Box of Search Algorithms,” *AALL Spectrum* 6, 11 (Nov.–Dec. 2019), <https://scholar.law.colorado.edu/articles/1238>. Lexis Advance uses “a suite of algorithms to identify the most relevant documents based on that.” *Id.* at 6.

29. Colo. RPC 1.1.

30. Moriarty, “The Legality of Generative AI—Part 2,” *supra* note 4 at 31.

31. *Id.* at 34.

32. See generally Colo. RPC 3.3, 5.1, 5.2, and 5.3.

33. Colo. RPC 3.3(a)(1). See also CBA Ethics Comm., Formal Op. 123, *Candor to the Tribunal and Remedial Measures in Civil Proceedings* (June 18, 2011).

34. Colo. RPC 3.3(a)(2).

35. *Id.* at cmt. [12]. The comments, however, are concerned primarily with fraudulent evidence, not fictitious cases.

36. Colo. RPC 4.1(a).

37. Colo. RPC. 5.1(a) and 5.2(b).

38. Colo. RPC 5.2(c).

39. Affidavit of Z.C. at 1, *Gates v. Chavez*, No. 2022CV31345 (El Paso Cnty. Dist. Ct. filed May 11, 2023) (previously published by Ritzdorf, “Colorado Springs Attorney Says ChatGPT Created Fake Cases He Cited in Court Documents,” KRDO NewsChannel 13 (June 13, 2023)), <https://krdo.com/news/2023/06/13/colorado-springs-attorney-says-chatgpt-created-fake-cases-he-cited-in-court-documents> (affidavit shown in embedded video of television broadcast at 0:54).

40. *Id.* at 2.

41. *Id.*

42. *Id.* at exhibit 4, attachments to Affidavit.

43. Order, *Gates v. Chavez*, No. 2022CV31345 (El Paso Cnty. Dist. Ct. filed May 5, 2023). As of August 29, 2023, no further orders for sanctions appear to have been issued.

44. Opinion and Order on Sanctions, *Mata v. Avianca, Inc.*, ___F.Supp.3d___, No. 22-cv-01461, 2023 U.S. Dist. LEXIS 108263, *6–7 ¶ 6 (S.D.N.Y. filed June 22, 2023).

45. *Id.* at *4–5 ¶¶ 1–2.

46. *Id.* at ¶ 2.

47. *Id.*

48. *Id.* at *8 ¶ 10, *13–14 ¶ 23.

49. *Id.* at *5 ¶ 3, *8–9 ¶ 11, *22 ¶ 41.

50. *Id.* at *21–22 ¶ 39.

51. *Id.* at *22 ¶ 41.

52. Affirmation in Opposition, *Mata*, 2023 U.S. Dist. LEXIS 108263 (S.D.N.Y. filed Feb. 28, 2023).
53. Reply Memorandum of Law in Further Support of Defendant's Motion to Dismiss Plaintiff's Verified Complaint, *Mata*, 2023 U.S. Dist. LEXIS 108263 (S.D.N.Y. filed Mar. 15, 2023).
54. Opinion and Order, *Mata*, *supra* note 44 at *10 ¶ 14.
55. *Id.* at *13-14 ¶ 23, *25-26 ¶ 45.
56. *Id.*
57. *Id.* at *23 ¶ 42.
58. *Id.* at *12 ¶ 21. The judge in *Mata* expressed doubt that Schwartz's confusion was legitimate, but the author is not so sure. One of the risks of relying on digital tools in the practice of law is that the older, analog basis for how the law is published and organized risks being overshadowed. To be sure, not every lawyer needs to be excited when they learn that the Federal Reporter graduated to its F.4th series in 2021, but a working knowledge of how the law is documented is good preventative medicine for any practitioner.
59. *Id.* at *10-11 ¶ 17.
60. *Id.* at *15-20 ¶¶ 27-36.
61. *Id.* at *1.
62. *Id.*
63. *Id.* at *30 ¶ 5 (citing *Gell v. Hartmax Corp.*, 496 U.S. 384, 398 (1990)).
64. *Id.* at *30-31 ¶ 6 (citing NYSBA R.Prof.Cond. 3.3(a)(1) and 22 NYCRR § 1200.0). Colorado's Rule 3.3(a)(1) is substantially the same as New York's Rule 3.3(a)(1), except that Colorado's requires a false statement of "material" fact, whereas New York's more broadly prohibits a "false statement of fact." Both impose the same duty to correct a "false statement of material fact or law" made to the tribunal.
65. *Id.* at *7-8 ¶ 7 (citing 18 USC § 505).
66. *Id.* at *14-19 ¶¶ 26-33.
67. Colo RPC 5.3(a) and (b). See also Colo. RPC 5.1(a) and (b) (imposing on partners or supervisory lawyers the same responsibility as all lawyers in the firm).
68. CRCP 11(a).
69. Colo. RPC 4.1 (providing that a lawyer, in the course of representing a client, shall not knowingly make a false statement of material fact or law to a third person).
70. Colo. RPC 8.4(c).
71. Colo. RPC 8.4(d).
72. Opinion and Order, *Mata*, *supra* note 44 at *1-2.
73. Swift, "The Art of Political Lying," *Examiner*, No. XIV (Nov. 9, 1710).
74. LexisNexis appears to be addressing this problem by adding a warning to the top of its results page that the opinion includes bogus cases: "Notice: This decision contains references to invalid citations in the original text of the opinion. They are relevant to the decision and therefore have not been editorially corrected. Linking has been removed from those citations." See, e.g., *Mata*, 2023 U.S. Dist. LEXIS 108263.
75. Baylson, Standing Order re: Artificial Intelligence ("AI") In Cases Assigned to Judge Baylson (E.D.Pa. June 6, 2023), <https://www.paed.uscourts.gov/documents/standord/standing%20Order%20Re%20Artificial%20Intelligence%206.6.pdf>.
76. Fuentes, Standing Order for Civil Cases Before Magistrate Judge Fuentes, at 2 (N.D.Ill. May 31, 2023), bit.ly/3PtGr3e (citing *Mata* and 2001: A Spacey Odyssey (Metro-Goldwyn-Mayer 1968)).
77. See, e.g., Vaden, Order on Artificial Intelligence (U.S.Ct. of Int'l Trade June 8, 2023), <https://www.cit.uscourts.gov/sites/cit/files/Order%20on%20Artificial%20Intelligence.pdf>.
78. Starr, Mandatory Certification Regarding Generative Artificial Intelligence (N.D.Tex.), <https://www.txnd.uscourts.gov/judge/judge-brantley-starr>.
79. *Id.*
80. See Weil, "You Are Not a Parrot and a Chatbot Is Not a Human. And a Linguist Named Emily M. Bender Is Very Worried What Will Happen When We Forget This," *Intelligencer* (Mar. 1, 2023), <https://nymag.com/intelligencer/article/ai-artificial-intelligence-chatbots-emily-m-bender.html>. It may not be fair to proclaim, as linguist Emily Bender does in the *Intelligencer* article and elsewhere, that LLMs have no "understanding" and no intelligence of any kind. There is not yet a good understanding of what kind of algorithms LLMs have developed that allow them to function as well as they do. Perhaps some LLMs contain a pattern similar to parts of our own brains. Perhaps not.
81. Casetext, "CoCounsel Harnesses GPT-4's Power to Deliver Results That Legal Professionals Can Rely On" (May 5, 2023), <https://casetext.com/blog/cocounsel-harnesses-gpt-4s-power-to-deliver-results-that-legal-professionals-can-rely-on>.
82. Fine-tuning also risks erasing some useful information or capabilities the model acquired during its prior training, known as "catastrophic forgetting." See Wolczyk et al., "On the Role of Forgetting in Fine-Tuning Reinforcement Learning Models," Workshop on Reincarnating Reinforcement Learning at ICLR 2023, at 1-5 (Apr. 2023), <https://openreview.net/pdf?id=zmXJUKULDzh>.
83. That does not mean the technology will never reach a point where it surpasses the most careful and skilled human practitioners. Despite being force-fed much human language output, the most recent version of ChatGPT is still only about two years old. It remains to be seen how much ChatGPT, or its successors, can grow in capability as they are given more ways to interact with the world and better software giving them additional functionality beyond simply acting as a language processing unit.
84. Moriarty, "The Legality of Generative AI—Part 2," *supra* note 4 at 37.
85. CRS § 13-90-107(1)(b). See also *Wesp v. Everson*, 33 P.3d 191, 196 (Colo. 2001).
86. E.g., *Wesp*, 33 P.3d at 196.
87. *Fox v. Alfini*, 432 P.3d 596, 601 (Colo. 2018).
88. CRS § 13-90-107(1)(b).
89. *Miller v. Dist. Ct.*, 737 P.2d 834, 837-38 (Colo. 1987), *superseded by statute as stated in Gray v. Dist. Ct.*, 884 P.2d 286, 291 (Colo. 1994). See also *Bellman v. Dist. Ct.*, 531 P.2d 632 (Colo. 1975) (insurance investigator communications privileged). At least for some kinds of experts, the decision in *Gray* suggests that consulting experts of an attorney might not fall within the scope of privilege. Nevertheless, Colorado commentators continue to opine that they generally do. Evans et al., "Managing Risks When Working with Experts and Consultants," 46 *Colo. Law.* 61, 62 n.7 (June 2017).
90. *Miller*, 737 P.2d at 837 n.3.
91. CRCP 26(b)(3). This includes consulting experts. *Gall v. Jamison*, 44 P.3d 233, 240 (Colo. 2002).
92. *Id.*
93. E.g., Colo. RPC 1.6.
94. Colo. RPC 1.6(a).
95. *In re Estate of Rabin*, 474 P.3d 1211, 1219 (Colo. 2020). See also Colo. RPC 1.6(b), cmt. [3].
96. Colo. RPC 1.6, cmt. [3].
97. *Rabin*, 474 P.3d at 1221 (citing Colo. RPC 1.6, cmt. [5], and D.C. Bar Ethics Op. 324 at 2 (2004)). See also Colo. RPC 1.6(b).
98. Colo. RPC 1.6(c).
99. Klinefelter, "When to Research is to Reveal: The Growing Threat to Attorney and Client Confidentiality from Online Tracking," 16 *Va. J. of L. & Tech.* 1, 24-25 (Spring 2011), <https://bit.ly/3PqZR8D>.
100. *Nguyen v. Excel Corp.*, 197 F.3d 200, 206 (5th Cir. 1999); *Schmidt v. Rodriguez*, 2013 Bankr. LEXIS 5048 (S.D.Tex. 2013); 6 *Moore's Fed'l Practice* § 26.70(2)(c).
101. "An attorney's consultation of a legal research tool or service should easily meet a test of necessity in the rendering of legal advice. For some types of research, courts have held that consultation of internet-based research tools is a necessary part of due diligence. Certainly, lawyers are using online research tools on a regular basis, with a majority reporting that they regularly begin legal research using online sources." Klinefelter, *supra* note 99 at 25 (citing *United States v. Kovel*, 296 F.2d 918, 921-22 (2d Cir. 1961), *cited with approval in United States v. Clark*, 847 F.2d 1467, 1468 n.1 (10th Cir. 1988)).
102. Klinefelter, *supra* note 99 at 26. If this is true, then there are additional ethical questions raised by the decisions of the lawyers in Colorado and New York to voluntarily and publicly disclose their legal research logs from ChatGPT to protect themselves, such as whether they obtained client consent before doing so.
103. *Mountain States Tel. & Tel. v. DiFede*, 780 P.2d 533, 542-43 (Colo.App. 1992).
104. Klinefelter, *supra* note 99 at 3.
105. CBA Ethics Comm., Formal Op. 90, *Preservation of Client Confidences in View of Modern Communications Technology* (Nov. 14, 1992, rev. July 2018).
106. ABA Comm. on Ethics and Prof'l Resp., Formal Op. 477R (2017); ABA Comm. on Ethics

- and Prof'l Resp., Formal Op. 11-459 (2011); ABA Comm. on Ethics and Prof'l Resp., Formal Op. 99-413 (1999); State Bar of Cal. Standing Comm. on Prof'l Resp. and Conduct, Formal Op. 2010-179 (2010); Prof'l Ethics Comm. of the Me. Bd. of Overseers of the Bar, Op. No. 195 (2008); N.Y. State Bar Ass'n Comm. on Prof'l Ethics, Op. 820 (2008); Alaska Bar Ass'n Ethics Comm., Op. 98-2 (1998); D.C. Bar Legal Ethics Comm., Op. 281 (1998); Ill. State Bar Ass'n Advisory Opinion on Prof'l Conduct, Op. 96-10 (1997); State Bar Ass'n of N.D. Ethics Comm., Op. No. 97-09 (1997); S.C. Bar Ethics Advisory Comm., Ethics Advisory Op. 97-08 (1997); Vt. Bar Ass'n, Advisory Ethics, Op. No. 97-05 (1997).
107. Colo. RPC 1.6, cmt. [19].
108. *Id.*
109. See CBA Ethics Comm., Formal Op. 90, *supra* note 105 at 3. This, however, is a case-by-case consideration, and particularly sensitive information may require more protection.
110. *Id.* at 4.
111. *Id.* at 4-5.
112. ABA Comm. on Ethics and Prof'l Resp. Formal Op. 95-398 (1995); Lenon, "A List of All the Ethics Opinions on Cloud Computing for Lawyers," *Clio* (Aug. 18, 2022), <https://www.clio.com/blog/cloud-computing-lawyers-ethics-opinions> (compiling links to ethics opinions from various states).
113. Colo. RPC 5.3, cmt. [3]. See also CBA Ethics Comm., Formal Op. 90, *supra* note 105 at 5-6; ABA Comm. on Ethics and Prof. Resp., Formal Op. 95-398, *Access of Nonlawyers to a Lawyer's Database* (Oct. 27, 1995).
114. Colo. RPC 5.3, cmt. [3].
115. See Preston, "Lawyers' Abuse of Technology," 103 *Cornell L. Rev.* 879, 928 (May 2018). This review also should be updated regularly as terms and conditions change. That is even more true with new technology like generative AI that is undergoing rapid changes.
116. *Id.* at 924.
117. OpenAI Terms of Use ¶ 3(c), <https://bit.ly/3r7IAJp>.
118. For the tech-savvy, instructions about how to go about doing this can be found in OpenAI's documentation. OpenAI API Reference, <https://platform.openai.com/docs/api-reference/introduction>. Large commercial users can also purchase private instantiations of OpenAPI's model and fine-tune it on their own data and, theoretically, this private model does not leak data upstream to the public model. See <https://platform.openai.com/docs/guides/fine-tuning>.
119. APA data privacy, <https://openai.com/api-data-privacy>.
120. The Colorado and New York lawyers appear, from their court filings, to have been using the web version of ChatGPT, another potential error since the default behavior of the web version is that the model will ingest prompts for training purposes.
121. 18 USC § 2701.
122. 18 USC § 2702.
123. *Theofel v. Farey-Jones*, 359 F.3d 1066, 1072-73 (9th Cir. 2004); *In re Subpoena Duces Tecum to AOL, LLC*, 550 F.Supp.2d 606, 609 (E.D.Va. 2008).
124. 18 USC § 2510(12). See also Goldberg, "The Googling of Online Privacy: Gmail, Search-Engine Histories and the New Frontier of Protecting Private Information on the Web," 9 *Lewis & Clark L. Rev.* 249, 260 (Spring 2005) (discussing the Stored Communications Act in relation to search queries).
125. Klinefelter, *supra* note 99 at 3 (citing Harris, [Prepared] *Testimony before the House Subcommittee on Commerce, Trade, & Consumer Protection*, Ctr. for Democracy & Tech., 2-3 (July 22, 2010), http://www.cdt.org/files/pdfs/CDT_privacy_bill_testimony.pdf, and Wingfield, "Microsoft Quashed Effort to Boost Online Privacy," *Wall St. J.* A1 (Aug. 2, 2010)).
126. OpenAI, "March 20 ChatGPT Outage: Here's What Happened," <https://openai.com/blog/march-20-chatgpt-outage> (detailing exposure of user's personal data). And lawyers already have their own professional responsibilities with respect to maintaining data and duties arising from a data breach. See CBA Ethics Comm., Formal Op. 141, *Ethical Issues Arising from Data Breach* (July 20, 2020).
127. *Federal Trade Commission Civ. Investigative Demand*, FTC File No. 232-3044, https://www.washingtonpost.com/documents/67a7081c-c770-4f05-a39e-9d02117e50e8.pdf?itid=inline_manual_4.
128. Colo. RPC 8.4(g). See generally CBA Ethics Comm., Formal Op. 145, *Discrimination Bias* (May 14, 2022) (discussing requirements of and differences between Rules 8.4(g) and 8.4(i)).
129. *In re Abrams*, 488 P.3d 1043, 1053-54 (Colo. 2021) (upholding Rule 8.4(g) against constitutional challenge).
130. *People v. Gilbert*, 2010 Colo. Discpl. LEXIS 79, at *12,*16 (Colo. OPDJ 2010).
131. *People v. McGarvey*, 2023 Colo. Discpl. LEXIS 27 (Colo. OPDJ 2023).
132. *Abrams*, 488 P.3d at 1052. This is only true while representing a client, however. While a lawyer "is free to speak in whatever manner he chooses" in his private life, a lawyer "must put aside the schoolyard code of conduct and adhere to professional standards" when representing a client. *Id.* at 1055.
133. See Johnson, "The Efforts to Make Text-based AI Less Racist and Terrible," *Wired* (June 17, 2021), <https://www.wired.com/story/efforts-make-text-ai-less-racist-terrible>.
134. Deshpande et al., "Toxicity in ChatGPT: Analyzing Persona-assigned Language Models" (Apr. 11, 2023), <https://arxiv.org/abs/2304.05335>.
135. Zou et al., "Universal and Transferable Adversarial Attacks on Aligned Language Models" (July 27, 2023), <https://arxiv.org/abs/2307.15043>.
136. Moriarty, "The Legality of Generative AI—Part 2," *supra* note 4 at 35. See also Sheng et al., "The Woman Worked as a Babysitter: On Biases in Language Generation," *Proc. of the 2019 Conf. on Empirical Methods in Nat. Language Processing and the 9th Int'l Joint Conf. on Nat'l Language Processing at 3407-12* (Hong Kong, China, November 3-7, 2019), <https://aclanthology.org/D19-1339.pdf>.
137. Solis, "How to Write Good ChatGPT Prompts," *Scribbr* (Caufield trans., June 13, 2023), <https://www.scribbr.com/ai-tools/chatgpt-prompts>.
138. Bias can be broader than simply racial or sexual biases, moreover. The New York attorney in the example may have contributed to his own demise by the biased way in which he asked ChatGPT for help. He did not ask it to give him an impartial answer or to consider both sides of his issue. Instead, he directed the LLM to argue in his favor and to provide cases that supported his position.
139. Opinion and Order, *Mata*, *supra* note 44 at *25 ¶ 21.
140. A court of appeals opinion not designated for official publication must state: "NOT PUBLISHED PURSUANT TO C.A.R. 35(e)." C.A.R. 35(f).
141. A good general practice tip when researching local Colorado trial orders or motions is always to use ICCES or PACER to access the actual docket and view the information in its native form.