

Ethics Opinion 388

Attorneys' Use of Generative Artificial Intelligence in Client Matters

Advances in technology have greatly improved the ways in which lawyers provide legal services. What technology has not done is alter lawyers' fundamental ethical obligations, and specifically, the duties lawyers owe to their clients—and to the courts. We anticipate that both of these statements will hold true with respect to lawyers' use of generative artificial intelligence (GAI). Due to the rapid development of technology in this area, we recognize that some of the concerns raised in this opinion may be resolved or mooted for particular products in the future, perhaps even in the near future.¹

The Rules of Professional Conduct require lawyers to be competent. Competence includes understanding enough about any technology the lawyer uses in legal practice to be reasonably confident that the technology will advance the client's interests in the representation. Separately, the lawyer should also be reasonably confident that use of and reliance on the technology will not be inconsistent with any of the lawyer's other obligations under the Rules of Professional Conduct.

Lawyers commonly adopt and use recently developed technology in their practices to achieve competitive advantage and gain efficiencies in providing legal services to their clients. Although technological innovation offers definite advantages, it comes with risks and the potential for adverse consequences. All of this holds true with respect to GAI. It can be a great boon to the practicing lawyer but—as recent events have shown—can sometimes be an untrustworthy and incompetent legal assistant.

Lawyers should understand that GAI products are not search engines that accurately report hits on existing data in a constantly updated database. The information available to a GAI product is confined to the dataset on which the GAI has been trained. That dataset may be incomplete as to the relevant topic, out of date, or biased in some way. More fundamentally, GAI is not programmed to accurately report the content of existing information in its dataset. Instead, GAI is attempting to create *new* content. In the case of a request for something in writing, GAI uses a statistical process to predict what the next word in the sentence should be. That is what the “generative” in GAI means: the GAI generates something new that has the properties its dataset tells it the user is expecting to see.

GAI products sometimes “hallucinate,” meaning they make up things that do not exist. As discussed below and as has been widely reported, a GAI product fabricated the names and citations of several “cases” that did not exist in response to a request for citations to support a particular legal position in

a brief. Believing the GAI to be a “super search engine,” a lawyer included the fake cites in the brief without checking them. That matter ended very badly for the lawyers who signed the brief. Current GAI for a general audience is not a reliable substitute for traditional fact- and cite-checking, and lawyers who blindly rely on outputs produced by GAI do so at considerable peril.²

Lawyers also should understand that many GAI products currently on the market are specifically designed to collect and use information received from users—which may include client confidences and secrets—for the GAI’s own training and for transmission to future users of the technology. That raises potential issues under D.C. Rule 1.6, which requires lawyers not to reveal client confidences or secrets to third parties without the client’s informed consent.

Regarding confidential client information, lawyers should determine whether the product will save information that the lawyers provide to the GAI, and whether the lawyers’ interaction with the GAI product will affect the answers the GAI gives to future users of the product outside of the lawyer’s law firm. Affirmative answers to either of those questions signal a need for caution. Depending on the circumstances, lawyers should either identify a different or more advanced GAI product that can be trusted with Client Confidential Information³ (or negotiate with the product vendor for improved confidentiality terms to make the first product trustworthy), or input only data that is not Client Confidential Information.

For proceedings before a court or other tribunal—especially but not only those that have adopted rules or issued orders regulating the use of GAI—lawyers must be attentive to their duty of candor to the tribunal and their fairness obligations to opposing parties and counsel. Lawyers should also be attentive to their obligations as to the client file with respect to their use of the GAI. Additionally, lawyers whose fee agreements provide for fees based solely on time spent may only bill for the time the lawyers actually spend, even if the GAI reduces the time the lawyers devote to the matter. Absent a prior agreement, a lawyer cannot charge separately for the perceived value to a client of the work done by the GAI, though they may pass through any out-of-pocket expenses for GAI applications where the client has agreed to pay for out-of-pocket expenses. Finally, these issues implicate lawyers’ duties of supervision. Under Rules 5.1 and 5.3, a lawyer should take reasonable measures to ensure that any supervised lawyer’s or nonlawyer’s use of GAI conforms to the Rules of Professional Conduct and the principles discussed in this opinion.

Applicable Rules

- Rule 1.1 (Competence)
- Rule 1.2 (Scope of Representation)
- Rule 1.5 (Fees)
- Rule 1.6 (Confidentiality of Information)
- Rule 1.16 (Client File)
- Rule 3.3 (Candor to Tribunal)
- Rule 3.4 (Fairness to Opposing Party and Counsel)
- Rule 5.1 (Responsibilities of Partners, Managers and Supervisory Lawyers)

- Rule 5.3 (Responsibilities Regarding Nonlawyer Assistants)
- Rule 8.4 (Misconduct)

Discussion

Most lawyers are not computer programmers or engineers and are not expected to have those specialized skills. As technology that can be used in legal practice evolves, however, lawyers who rely on the technology should have a reasonable and current understanding of how to use the technology with due regard for its potential dangers and limitations. So it is with generative AI technology. The widely reported events culminating with *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023), provide an object lesson.

The *Mata* debacle began with a lawyer's fundamental misunderstanding of a technology that was new to him. The principal plaintiff's lawyer in that case regularly practiced in state courts in New York. The legal research database that he normally used was complete as to state court caselaw but had only limited access to federal caselaw, which became a problem when a case he had filed in state court was removed to federal court. He needed a new legal research tool.

As the lawyer testified at the sanctions hearing, "I had heard about this new site which I assumed – I falsely assumed was like a super search engine called ChatGPT, and that's what I used." 678 F. Supp. 3d at 456. With that incorrect understanding, he used ChatGPT to look for cases upon which he could rely in opposing a motion to dismiss his client's complaint.

Unbeknownst to the lawyer, the ChatGPT tool that he consulted was not a legal research tool built to search a comprehensive library (a database) of previously published judicial opinions. The "free" version of ChatGPT that he chose had no such library. Indeed, ChatGPT explicitly disclaimed having a database in responding to an inquiry sent in October 2023 (after the events just described):

I don't have a "database" in the traditional sense. I generate responses based on the text data I was trained on up until my last knowledge update in September 2021. My responses are generated by predicting what comes next in a given text prompt, drawing upon the patterns and information present in the text data I was trained on. I don't have access to the internet or real-time databases, and my knowledge is static, meaning I can't provide information on events or developments that have occurred after my last update.

The technical term for what ChatGPT and many other GAI programs use is a "dataset" rather than a *database*.⁴ A dataset may be thought of as a limited pool of materials that illustrate a premise on which GAI learns the vocabulary necessary to generate a reasonable-sounding answer to a user's questions – without any regard for the truth of the answer. For GAI to generate case citations, for example, the dataset merely needs to include information indicating the proper format in which case citations normally appear, *e.g., Party A v. Party B*, 100 F.3d 100 (D.C. Cir. 2020). For the GAI to produce citations for divorce cases, it might rely on information in the dataset establishing that the parties in divorce cases tend to have the same last name, and that such cases are published in a state

reporting service rather than a federal reporter. Similarly, to generate criminal case citations, the GAI might rely on information that a government entity normally appears on the plaintiff side of the caption.

Had the lawyer in *Mata* understood that only a limited dataset was available to ChatGPT, he presumably would not have felt comfortable using the platform to research federal case law. As ChatGPT's disclaimer statement establishes, its initial training might or might not have included subject matter relevant to the *Mata* attorney's research; even if it had, that information would no longer have been current.

A more fundamental problem is that large language GAI tools like ChatGPT simply are not built to supply accurate answers even based on the limited datasets available to them, at least not yet. Instead, they respond to text prompts by creating "new" content that is statistically similar to what they have seen before:

Generative AI refers to deep-learning models that can take raw data – say, all of Wikipedia or the collected works of Rembrandt – and “learn” to generate statistically probable outputs when prompted. At a high level, generative models encode a simplified representation of their training data and draw from it to create a new work that's similar, but not identical, to the original data.

Kim Martineau, [What is Generative AI? | IBM Research Blog⁵](#), (Apr. 20, 2023). For example, a graphical GAI program might be prompted to generate a portrait of a horse that appears to have been painted by Rembrandt. But that output is new, different from what came before, and certainly not a genuine Rembrandt painting.

“Statistically probable outputs” are not what a lawyer searching for existing controlling authorities needs or wants. Generative AI offered to a broad audience is simply no substitute for the familiar legal research tools provided by trusted services like LexisNexis and Westlaw, which have databases of previously published laws, regulations, and cases that are constantly updated as laws change and new cases are decided.⁶

In the *Mata* case, ChatGPT appeared to provide the lawyer with exactly what he requested: six case citations for the propositions the lawyer needed to defeat a motion to dismiss. Unfortunately for the lawyer, however, those cases—reported by ChatGPT with case numbers, court names, citations, and authorship by real judges—did not exist.

Why did ChatGPT make up cases from whole cloth? As the IBM Research Blog explains:

Many generative models, including those powering ChatGPT, can spout information that sounds authoritative but isn't true (sometimes called “hallucinations”) or is objectionable and biased. Generative models can also inadvertently ingest information that's personal

or copyrighted in their training data and output it later, creating unique challenges for privacy and intellectual property laws.

Id. Put differently, “The best way to think about this is you are chatting with an omniscient, eager-to-please intern who sometimes lies to you.” [Emma Bowman, *A New AI Chatbot Might Do Your Homework for You. But It's Still Not an A+ Student*, National Public Radio \(Dec. 19, 2022\)](#) (quoting Ethan Mollick, a professor at the University of Pennsylvania’s Wharton School of Business). Similarly:

“There are still many cases where you ask it a question and it’ll give you a very impressive-sounding answer that’s just dead wrong,” said Oren Etzioni, the founding CEO of the Allen Institute for AI, who ran the research nonprofit until recently. “And, of course, that’s a problem if you don’t carefully verify or corroborate its facts.”

Id. For a visual demonstration of GAI hallucinations in action, see Gerrit De Vynck, Jhaan Elker and Tyler Rimmel, [The Future of AI Video Is Here, Super Weird Flaws and All](#), Washington Post (Feb. 28, 2024).

Again, had the *Mata* lawyer known of the risk of hallucinations at the outset, he presumably would have acted differently by, at a minimum, performing a traditional cite check on the cases ChatGPT claimed to have found.⁷ He would not have done what he did next.

Upon receiving the brief with the fake citations, opposing counsel and the court raised questions about cited cases that they could not find. Still unaware of the “hallucinations” issue, the attorney doubled down with ChatGPT and asked the GAI whether the cases were real. ChatGPT “responded that it had supplied ‘real’ authorities that could be found through Westlaw, LexisNexis and the Federal Reporter.” 678 F. Supp. 3d at 458. Again, that was not true. The cases were fabrications manufactured by the GAI’s innate “desire” to give the lawyer what he was looking for.

At the conclusion of the sanctions hearing, the involved lawyers (the one who used ChatGPT and wrote the brief, and his local counsel supporting him in the federal court to which he was only admitted pro hac vice) were required to (1) pay a \$5,000 fine; (2) notify their client in writing of the court’s decision and of the background leading up to the sanction; and (3) “mail a letter individually addressed to each judge falsely identified as the author of the fake ‘*Varghese*,’ ‘*Shaboon*,’ ‘*Petersen*,’ ‘*Martinez*,’ ‘*Durden*,’ and ‘*Miller*’ opinions,” providing a copy of the sanctions opinion, the transcript of the hearing, and the fake opinion wrongly attributed to each judge. *Id.* at 466.

In the wake of the wide publicity surrounding the *Mata* case, a number of judges around the country issued orders addressing potential use of GAI by attorneys practicing before them. Judge Brantley Starr of the U.S. District Court for the Northern District of Texas was one of the first:

All attorneys appearing before the Court must file on the docket a certificate attesting either that no portion of the filing was drafted by generative artificial intelligence (such as ChatGPT, Harvey.AI, or Google Bard) or that any language drafted by generative artificial

intelligence was checked for accuracy, using print reporters or traditional legal databases, by a human being. These platforms are incredibly powerful and have many uses in the law: form divorces, discovery requests, suggested errors in documents, anticipated questions at oral argument. But legal briefing is not one of them. Here's why. These platforms in their current states are prone to hallucinations and bias. On hallucinations, they make stuff up—even quotes and citations. Another issue is reliability or bias. While attorneys swear an oath to set aside their personal prejudices, biases, and beliefs to faithfully uphold the law and represent their clients, generative artificial intelligence is the product of programming devised by humans who did not have to swear such an oath. As such, these systems hold no allegiance to any client, the rule of law, or the laws and Constitution of the United States (or, as addressed above, the truth). Unbound by any sense of duty, honor, or justice, such programs act according to computer code rather than conviction, based on programming rather than principle. Any party believing a platform has the requisite accuracy and reliability for legal briefing may move for leave and explain why. Accordingly, the Court will strike any filing from an attorney who fails to file a certificate on the docket attesting that the attorney has read the Court's judge-specific requirements and understands that he or she will be held responsible under Rule 11 [of the Federal Rules of Civil Procedure] for the contents of any filing that he or she signs and submits to the Court, regardless of whether generative artificial intelligence drafted any portion of that filing.

Chief Justice John Roberts specifically referenced GAI issues in his [2023 Year-End Report on the Federal Judiciary](#). As he described it, "AI combines algorithms and enormous data sets to solve problems." *Id.* at 5. While "AI apparently can earn Bs on law school assignments and even pass the bar exam... any use of AI requires caution and humility." *Id.*

One of AI's prominent applications made headlines this year for a shortcoming known as "hallucination," which caused the lawyers using the application to submit briefs with citations to non-existent cases. (Always a bad idea.) Some legal scholars have raised concerns about whether entering confidential information into an AI tool might compromise later attempts to invoke legal privileges.

Id. at 5-6.

Closing out an eventful year for GAI, disbarred former lawyer Michael Cohen admitted to (unknowingly) supplying his lawyer in his criminal case with fake citations generated by GAI. Ella Lee, [Michael Cohen Gave Lawyer Fraudulent Case Citations Generated by AI](#), The Hill (Dec. 29, 2023). This time the GAI in question was Google Bard, which Cohen believed to be "a supercharged search engine." *Id.* Cohen wrote that he had "not kept up with 'emerging trends (and related risks)' in legal technology and was not aware" that generative text services "could create citations and descriptions that 'looked real but actually were not.'" *Id.* Falling back on his current non-lawyer status, Cohen wrote that he trusted his lawyer to vet his suggestions before "drop[ping] the cases into his submission wholesale without even confirming that they existed," which is what his counsel did. *Id.*

Unmentioned thus far is what ChatGPT likely did with the fake cases it created in response to the *Mata* lawyer's inquiry. As mentioned above, ChatGPT was initially trained on a limited dataset. As time goes on, however, many GAI programs supplement their datasets based on interactions with users. To ChatGPT, the interaction with the *Mata* lawyer was successful and the fake case names may have been added to the dataset to be reported to future users with similar questions. Until the hallucination issue is resolved, systems prone to this problem are therefore self-corrupting, which is yet another reason their outputs need to be checked carefully.

In early 2024, researchers at Stanford University announced the preliminary results of a study finding that “[l]arge language models hallucinate at least 75% of the time when answering questions about a court’s core ruling.” Isabel Gottlieb & Isaiah Poritz, *Legal Errors by Top AI Models "Alarmingly Prevalent," Study Says*, Bloomberg Law (Jan. 12, 2024). Their testing involved “more than 200,000 legal questions on OpenAI’s Chat GPT 3.5, Google’s PaLM 2, and Meta’s Llama 2—all general-purpose models not built for specific legal use.” *Id.* One of the researchers said:

We should not take these very general purpose foundation models and naively deploy them and put them into all sorts of deployment settings, as a number of lawyers seem to have done.... Proceed with much more caution—where you really need lawyers, and people with some legal knowledge, to be able to assess the veracity of what an engine like this is giving to you.

Id.

More recently, a Washington Post reporter evaluated two GAI products for use in connection with some of her job functions. Danielle Abril, *I Used AI Work Tools to Do My Job. Here's How it Went* - The Washington Post (Feb. 26, 2024). She found that “[t]he AI seemed to do better when it was fed documents or data. But it still sometimes made things up, returned error messages or didn’t understand context.” *Id.* The article ended as follows:

[A]ll results and content need careful inspection for accuracy, some tweaking or deep edits – and both tech companies advise users to verify everything generated by the AI. “I don’t want people to abdicate responsibility,” said Kristina Behr, vice president of product management for collaboration apps at Google Workspace. “This helps you do your job. It doesn’t do your job.”

And as is the case with AI, the more details and direction in the prompt, the better the output. So as you do each task, you may want to consider whether AI will save you time or actually create more work.

“The work it takes to generate outcomes like text and videos has decreased,” Rahman [a professor at Northwestern University’s Kellogg School of Management] said. “But the work to verify has significantly increased.”

Id.

Also quite recently, a law firm invoked ChatGPT in a court filing for new and different purpose. *J.G. v. New York City Dept. of Education*, 2024 WL 728626 (S.D.N.Y. Feb. 22, 2024). Given the now well documented issues with ChatGPT, the court was not receptive. The case involved a motion for attorneys' fees from the losing party by counsel for a prevailing plaintiff in a situation in which the law permits such fee shifting. Movants must establish the "reasonableness" of their fee request. After making what the court deemed to be "aggressive" arguments under case law with respect to the reasonableness issue, the plaintiff's attorneys submitted a report from ChatGPT "as a 'cross-check' supporting" what the court deemed to be "problematic sources." *Id.* at *7. The court found the law firm's "invocation of ChatGPT as support for its aggressive fee bid" to be "utterly and unusually unpersuasive." *Id.* "As the firm should have appreciated, treating ChatGPT's conclusions as a useful gauge of the reasonable billing rate for the work of a lawyer with a particular background carrying out a bespoke assignment for a client in a niche practice area was misbegotten at the jump." *Id.* The court discussed *Mata* and another case where lawyers submitted fake citations manufactured by ChatGPT. And then the court focused on the absence of any basis for believing that ChatGPT could be a reliable source of information in this context:

In claiming here that ChatGPT supports the fee award it urges, the [law firm] does not identify the inputs on which ChatGPT relied. It does not reveal whether any of these were similarly imaginary. It does not reveal whether ChatGPT anywhere considered a very real and relevant data point: the uniform bloc of precedent, canvassed below, in which courts in this District and Circuit have rejected as excessive the billing rates the [law firm] urges for its timekeepers. The Court therefore rejects out of hand ChatGPT's conclusions as to the appropriate billing rates here. Barring a paradigm shift in the reliability of this tool, the [law firm] is well advised to excise references to ChatGPT from future fee applications.

Id. at *7.⁸

Against that background, we address specific Rules of Professional Conduct implicated by attorneys' use of GAI in legal practice. We recognize that there are different kinds of GAI in existence now, and increasingly powerful ones in the pipeline. Use of certain AI or AI-like products for document review and discovery has become an accepted part of legal practice because experience has shown that, when used properly, some of these products yield accurate results and reduce the cost of document review and production. That is not yet the case for the current versions of certain GAI products when used for legal research. This opinion, which is confined to *generative* AI products, is intended to provide guidance to attorneys as the technology evolves to determine when and how those products can be ethically incorporated into their practice. Due to rapid development of technology in this area, we anticipate that many of the concerns raised in this opinion may be resolved or mooted for particular products in the future, perhaps even in the near future.

As discussed below, use of GAI in legal matters implicates lawyers' duties of competence, confidentiality, communication, candor to the court, and fairness to opposing parties and counsel; their obligation to supervise nonlawyer assistants and hold them to lawyer standards of conduct; their obligations as to the fees and expenses of a representation; and their obligation to maintain and make available a complete client file at the conclusion of a representation.

A. Competence

Under Rule 1.1(a), "[a] lawyer shall provide competent representation to a client." "To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, and engage in such continuing study and education as may be necessary to maintain competence." Rule 1.1 cmt. [6]. As we noted in [D.C. Legal Ethics Opinion 371](#) regarding the use of social media in the practice of law:

We agree with ABA Comment [8] to Model Rule 1.1 that to be competent "a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology." Although the District's Comments to Rule 1.1 do not specifically reference technology, competent representation always requires the legal knowledge, skill, thoroughness, and preparation reasonably necessary to carry out the representation. Because of society's embrace of technology, a lawyer's ignorance or disregard of it, including social media, presents a risk of ethical misconduct.

We hold the same view as to the use of GAI in legal practice.⁹ Indeed, there may come a time when lawyers' use of GAI is standard practice.

Before using any particular form of GAI, attorneys should have a reasonable and current understanding of how it works and what it does, with due regard for (a) its potential dangers, including the risk of "hallucinations" or misuse or exposure of Client Confidential Information, (b) its limitations, including whether it uses a narrow dataset that could generate incomplete, out-of-date, or inaccurate results, and (c) its cost. Attorneys also should have a reasonable basis for trusting the GAI outputs, or must review and validate GAI outputs, before incorporating these outputs in their work product for clients or relying on them in support of a legal proceeding.

How might a lawyer do these things? Given what has befallen fellow lawyers in connection with GAI, we suggest the kind of diligence that any reasonable business owner would undertake before making a significant investment in technology for their legal practice. Depending on the context, this might include the following questions and answers:

- What is in the news about the GAI platform with respect to legal practice?
- Has the GAI been tested for your intended purpose by disinterested third parties?
- Have other legal professionals that you trust used the GAI? Have they encountered any issues?
- Ask the GAI to do something you or a respected colleague or adversary has already done and compare its output with the human work product.¹⁰

- Verify the accuracy and completeness (including any date limitations) of the GAI's output in a test run, especially with respect to citations to laws, regulations and judicial decisions.

What is required or prudent in any specific situation depends on the circumstances. Significantly less diligence may be warranted for GAI products that rely on proven, current databases of the relevant information. At the other end of the spectrum, much more care should be taken with respect to GAI products that are offered to a general audience and lack a positive track record in connection with legal services.

B. Confidentiality

Under Rule 1.6(a)(1), “a lawyer shall not knowingly... reveal a confidence or secret of the lawyer’s client.” Under this rule “[c]onfidence’ refers to information protected by the attorney-client privilege under applicable law and ‘secret’ refers to other information gained in the professional relationship that the client has requested be held inviolate, or the disclosure of which would be embarrassing, or would be likely to be detrimental, to the client.” Rule 1.6(b). As we noted in [DC Legal Ethics Opinion 364](#):

This expansive confidentiality obligation “[t]ouch[es] the very soul of lawyering.” *In re Gonzalez*, 773 A.2d 1026, 1030 (D.C. 2001) (quoting *Fred Weber, Inc. v. Shell Oil Co.*, 566 F.2d 602, 607 (8th Cir. 1977)). “Disclosure of client confidences is ‘contrary to the fundamental principle that an attorney owes a fiduciary duty to his client and must serve the client’s interests with the utmost loyalty and devotion.’” *Herbin v. Hoeffel*, 806 A.2d 186, 197 (D.C. 2002) (quoting *In re Gonzalez*, 773 A.2d at 1031).

Comment [40] to Rule 1.6 notes that, “[w]hen transmitting a communication that includes information relating to the representation of a client, the lawyer must take reasonable precautions to prevent the information from coming into the hands of unintended recipients.”

Separately, Rule 1.6(f) requires lawyers to “exercise reasonable care to prevent the lawyer’s employees, associates, and others whose services are utilized by the lawyer from disclosing or using confidences or secrets of a client.” A third party GAI provider who has access to the inputs to a GAI program is an “other” as to which this obligation extends.

To protect client confidences and secrets, lawyers should ask two questions:

1. Will information I provide to the GAI be visible to the GAI provider or other strangers to the attorney-client relationship?
2. Will my interactions with the GAI affect answers that later users of the GAI will get in a way that could reveal information I provided to the GAI?

From the perspective of confidentiality, an affirmative answer to the first is at least a red flag but, perhaps, one that can be resolved after a negotiation with the GAI provider (or an upgrade to a paid

product with better terms) to improve the data security and prevent third party access.¹¹ An affirmative answer to the second might be more challenging to resolve. A lawyer should be reasonably satisfied that her interaction with the GAI will not reveal Client Confidential Information to future users of the GAI. If the lawyer is not so satisfied, she should not reveal Client Confidential Information to the GAI or should not use the GAI.

Many currently available GAI products invite a conversation with their users. Attorney users seek to have GAI generate legal theories to support their client's claims or defenses. They may also share their own legal or factual theories hoping the GAI can provide legal and factual support. The attorney then reacts to the response with follow up inquiries. Through these inputs, the lawyer effectively may be disclosing her mental impressions to the GAI in the hope of getting further support and refining her strategy. Meanwhile, some GAI programs may be saving both the attorney's inputs (or "prompts") and the GAI's own outputs to make that "content" available to future users.

The lawyer's mental impressions about a client matter are, of course, work product. Separate from the attorney/client privilege, work product is among the most protected categories of confidential information in our adversary legal system. Even when a court orders disclosure of some work product information after a required showing of "substantial need," the court "must protect against disclosure of the mental impressions, conclusions, opinions, or legal theories of a party's attorney or other representative concerning the litigation." Fed. R. Civ. P. 26(b)(3)(B); D.C. Super. Ct. R. Civ. P. 26(b)(3)(B).

Lawyers who input Client Confidential Information, or their mental impressions about a client matter, into GAI products –especially "free" ones—risk violating fundamental rules about client confidentiality. Most technology companies that provide these services make no secret of what they will do with any information submitted to them in connection with their publicly usable services: from their perspective, user inputs are theirs to use and share as they see fit.

For example, the [Privacy Policy](#) underlying ChatGPT's free offering makes clear that ChatGPT and its parent OpenAI:

- "[C]ollect Personal Information that is included in the input, file uploads, or feedback that you provide to our Services ("Content"); and
- "[M]ay use Personal Information...[t]o improve our Services and conduct research" and "[t]o develop new programs and services."

Indeed, the policy shows that they view the information as simply another "asset" of theirs to be exploited and sold to others:

Business Transfers: If we are involved in strategic transactions, reorganization, bankruptcy, receivership, or transition of service to another provider (collectively, a "Transaction"), your Personal Information and other information may be disclosed in the

diligence process with counterparties and others assisting with the Transaction and transferred to a successor or affiliate as part of that Transaction along with other assets.

Absent client consent, lawyers who share Client Confidential Information with third party providers who have privacy policies like this risk violating their confidentiality obligations under Rule 1.6. And clients are unlikely to give informed consent—and typically should not be asked to consent—to wide ranging disclosures that could waive attorney/client privilege or otherwise make their most confidential and secret information available for third parties to see and use.¹² This includes potential litigation adversaries and their counsel who also have access to the same GAI. Attorneys who would provide client confidences and secrets to a GAI product should ensure that product has implemented adequate security safeguards and controls to ensure confidentiality and protect against unauthorized access and use of client information.

Comment [5] to Rule 1.6 does say that “[a] lawyer’s use of a hypothetical to discuss issues relating to the representation is permissible so long as there is no reasonable likelihood that the listener will be able to ascertain the identity of the client or the situation involved.” This may tempt GAI users to try to protect client confidentiality by anonymizing information that they submit to the GAI. Again, we urge caution. The more information a lawyer provides to a growing GAI dataset, the greater the likelihood that the GAI or one of its other users will be able to connect the dots and link the information the lawyer provided to the client in question. GAI is, after all, artificial “intelligence.”

This is especially true if the lawyer’s representation of a client is available in publicly searchable information, such as docket sheets for litigation, news reports about litigation, or—to the extent the lawyer’s website identifies specific client matters handled by the lawyer—the lawyer’s own website. Similarly, if a lawyer using such a service inputs her client’s name for billing purposes, anyone with access to the service’s records may be able to connect a given research request with an identifiable client.

Even if the lawyer and the client cannot be linked by the GAI as the source of information provided to the GAI, there is still a potential for harm if the information itself is valuable to the client because of the secrecy surrounding it. Imagine, for example, the harm that would occur to the client if a lawyer shared the client’s trade secret manufacturing process with a GAI, and the GAI later revealed that information to others looking for faster and cheaper ways of making the product in question.¹³

One option that some GAI products provide to resolve confidentiality concerns is a zero data retention policy in which the provider of the GAI retains neither the inputs nor the outputs of the GAI’s interaction with a particular user. As noted above, business users who pay to use a GAI product may be able to negotiate better terms than are available to users of a “free” service that the provider makes available for the provider’s own marketing and product development purposes.

C. Responsibilities Regarding Lawyers in a Firm and Their Nonlawyer Assistants



Under Rules 5.1 and 5.3, a lawyer should take reasonable measures to ensure that any supervised lawyer's or nonlawyer's use of GAI conforms to the Rules of Professional Conduct and the principles discussed in this opinion.

Rule 5.1(a) requires managers of law firms and other lawyers with comparable managerial authority in a law firm or government agency to "make reasonable efforts to ensure that the firm has in effect measures giving reasonable assurance that all lawyers in the firm or agency conform to the Rules of Professional Conduct." Rule 5.1(b) requires lawyers with direct supervisory authority over another lawyer to "make reasonable efforts to ensure that the other lawyer conforms to the Rules of Professional Conduct."

Rule 5.3 extends these obligations to nonlawyers working in or for the firm, lawyer, or government agency. Under Rule 5.3(b), "[a] lawyer having direct supervisory authority over the nonlawyer shall make reasonable efforts to ensure that the person's conduct is compatible with the professional obligations of the lawyer." Similarly, under Rule 5.3(a), partners or other lawyers in a firm or government agency "who individually or together with other lawyers possess[] comparable managerial authority in a law firm ... shall make reasonable efforts to ensure that the firm or agency ... has in effect measures giving reasonable assurance that the person's conduct is compatible with the professional obligations of the lawyer." Thus, both the lawyers who retain nonlawyers and the managers of a firm or lawyers with comparable managerial authority in a law firm or government agency that retains nonlawyers must take steps to assure that the nonlawyers abide by the professional conduct rules for lawyers and confirm compliance with attorneys' ethical responsibilities.

Where it is foreseeable that lawyers or nonlawyers within or retained by a firm or government agency will be using GAI in connection with a client representation, the firm and the retaining lawyers should take appropriate steps to ensure that any use of GAI is consistent with the Rules of Professional Conduct.

One step law firms, lawyers, and government agencies could consider is to require employees—lawyers and nonlawyers alike—to satisfy themselves that client confidentiality under Rule 1.6 will be protected before using a GAI product. For example, if a review of the GAI's privacy policy determines that the GAI will not keep client confidential or secret information from the GAI's owner and other third parties outside the law firm, law office, or government agency, Rules 5.1 and 5.3 likely prohibit the supervising lawyers from permitting supervised personnel to use the GAI in the client matter. In the absence of a clear privacy policy, supervising lawyers should consider directing the employees to make other inquiries to satisfy themselves that the client's confidentiality will be protected.

Another step could be to require lawyers and nonlawyers within or retained by a law firm, lawyer, or government agency to take steps to verify the accuracy of the output of any GAI they use. At some point, however, the time and costs associated with the verification may outweigh whatever perceived benefits led the lawyer to consider using the GAI in the first place. Once a firm has vetted a particular

GAI for particular purposes, it could require its lawyers to use that GAI for those purposes rather than some competing GAI not yet approved within the firm.

D. Candor to Tribunal and Fairness to Opposing Party and Counsel

For matters in litigation or arbitration, the use of GAI outputs that contain misrepresentations of facts or law, or that provide fake citations, also implicates the lawyer's duties to the tribunal and to the opposing party and counsel. This is especially true if the tribunal has adopted rules or procedures or issued orders requiring disclosure of the use of GAI and verification or other safeguards with respect to GAI outputs.

Under Rule 3.3(a), lawyers shall not knowingly:¹⁴

(1) Make a false statement of fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the lawyer, unless correction would require disclosure of information that is prohibited by Rule 1.6;

* * *

(3) Fail to disclose to the tribunal legal authority in the controlling jurisdiction not disclosed by opposing counsel and known to the lawyer to be dispositive of a question at issue and directly adverse to the position of the client; or

(4) Offer evidence that the lawyer knows to be false, except as provided in paragraph (b) [*dealing with representation of the accused in a criminal case*]. A lawyer may refuse to offer evidence, other than the testimony of a defendant in a criminal matter, that the lawyer reasonably believes is false.

These duties "continue to the conclusion of the proceeding." Rule 3.3(c). "A lawyer who receives information clearly establishing that a fraud has been perpetrated upon the tribunal shall promptly take reasonable remedial measures, including disclosure to the tribunal to the extent disclosure is permitted by Rule 1.6(d)." Rule 3.3(d).

The comments to this rule specifically address use of fake citations akin to those generated by the GAI in *Mata*. "Legal argument based on a knowingly false representation of law constitutes dishonesty toward the tribunal. A lawyer is not required to make a disinterested exposition of the law, but must recognize the existence of pertinent legal authorities." Rule 3.3 cmt. [3].

Similarly, Rule 3.4 imposes certain duties of fairness to the opposing party and counsel. Among these duties, the lawyer shall not:

- "[f]alsify evidence" (Rule 3.4(b)); or

- “[k]nowingly disobey an obligation under the rules of a tribunal except for an open refusal based on an assertion that no valid obligation exists” (Rule 3.4(c)).

Rule 3.4(c) comes into play when the tribunal has adopted rules or has issued orders restricting or otherwise governing use of GAI in connection with the proceeding, and the lawyer does not comply with those rules or orders.¹⁵

E. Fees

Under Rule 1.5(a), “[a] lawyer’s fee shall be reasonable.” The rule provides a non-exclusive list of eight factors to be considered in assessing the fee’s reasonableness. Separately, when “the lawyer has not regularly represented the client,” the lawyer must send the client a writing stating “the basis or rate of the fee, the scope of the lawyer’s representation, and the expenses for which the client will be responsible.” Rule 1.5(b). That writing must be sent “before or within a reasonable time after commencing the representation.” *Id.*

If the lawyer intends to bill the client for use of GAI for which there is an out-of-pocket cost to the lawyer, that expected cost is an expense that should be communicated to the client under this rule.

Separately:

[i]t goes without saying that a lawyer who has undertaken to bill on an hourly basis is never justified in charging a client for hours not actually expended. If a lawyer has agreed to charge the client on this basis (i.e., hourly), and it turns out that the lawyer is particularly efficient in accomplishing a given result, it nonetheless will not be permissible to charge the client for more hours than were actually expended on the matter. When that basis for billing the client has been agreed to, the economies associated with the result must inure to the benefit of the client.

[D.C. Legal Ethics Opinion 267](#) (1996) (quoting ABA Formal Ethics Opinion 379 (1993)).

A familiar variation of this issue occurs when a lawyer expends considerable time and effort to prepare a detailed legal research memo for one client and, to that first client, considerable expense based on the lawyer’s hourly rate. Shortly thereafter, a second client happens to ask the same legal question. It will take far less time to adapt the first memorandum for the second client’s use than it took to create the memorandum in the first place. While the lawyer may believe it is not fair or reasonable to charge the second client only a fraction of what the first client paid for such a valuable piece of legal research, that is what is required if the lawyer’s billing arrangement with the second client is based exclusively on an hourly rate.

The same is true when the use of GAI reduces billable time and the lawyer’s fee agreement with the client is based exclusively on the time the lawyer spends working on the matter. No matter how good or valuable the GAI’s output is, absent a different fee arrangement, the lawyer can only bill for the

time the lawyer spent. As discussed above, the reasonable expense of the GAI itself may be billed as an expense item if the lawyer's agreement with the client permits the lawyer to bill for such expenses.

F. Client File

When a representation is terminated, Rule 1.16(d) requires a lawyer to do several things, including "surrendering papers and property to which the client is entitled." As we discussed in [D.C. Legal Ethics Opinion 333](#) (2005), this rule requires production of the "entire file," including "copies of internal notes and memoranda reflecting the views, thoughts and strategies of the lawyer." Although lawyers are not required to retain every piece of paper or electronic datum generated or received during a client representation, a lawyer should consider whether specific interactions with GAI in connection with a client matter should be retained as part of the client file.¹⁶

Conclusion

We anticipate that GAI eventually will be a boon to the practice of law. Moreover, lawyers who use generative artificial intelligence do not need to be computer programmers who can write AI programs or critique AI code written by others. But they do need to understand enough about how GAI works, what it does, and its risks and limitations to become comfortable that the GAI will be helpful and accurate for the task at hand, and that it will not breach client confidentiality. Lawyers should also be mindful of the implications GAI creates for their duties of supervision; their duty of candor to the tribunal and their fairness obligations to opposing parties and counsel; the reasonableness of their fees; and their obligations with respect to the client file.

Published: April 2024

-
1. This opinion is based on information available to the Committee as of the second quarter of 2024.
 2. Although not the subject of this opinion, lawyers should also be aware that GAI has the potential to facilitate outright fraud by bad actors. For example, if prompted to do so, GAI can produce genuine looking videos or photographs of things that never happened, and "recordings" making it sound like a person said something that they never said.
 3. As discussed below, Rule 1.6 requires lawyers to maintain the confidentiality of certain information acquired during the professional relationship. This obligation extends to both "confidences" (defined by the rule as "information protected by the attorney-client privilege under applicable law") and "secrets" (defined as other information gained in the relationship "that the client has requested be held inviolate, or the disclosure of which would be embarrassing, or would be likely to be detrimental, to the client."). This opinion will refer to information protected by Rule 1.6 as

“Client Confidential Information.” Unless otherwise indicated, any citations to a “Rule” in this opinion will be to the D.C. Rules of Professional Conduct.

4. Such GAI tools are initially trained on a finite dataset that, unlike a traditional database, may not be updated regularly. *See generally* <https://atlan.com/dataset-vs-database/#:~:text=A%20dataset%20is%20like%20a,for%20ongoing%20data%20management%20tasks> and <https://databasetown.com/dataset-vs-database-key-differences/>. The initial training might not include the subject matter relevant to whatever the lawyer wants to ask about and, even if it did initially, that information might not have been kept current.

5. This link and the other links in this opinion were last visited in April of 2024. Over time, the content of information at links may change or the links themselves may stop working. The links in this opinion will not be checked or updated after publication of this opinion.

6. We understand that both Westlaw and Lexis are now offering GAI-assisted cite checking products. *See* [Westlaw Quick Check](#) and [Lexis + AI](#). While we are unable to vouch for either product, we understand that they seek to make their legal research databases as complete and up to date as technology allows.

7. A check of ChatGPT’s policies many months after the widely-reported *Mata* debacle yielded this disclaimer:

A note about accuracy: Services like ChatGPT generate responses by reading a user’s request and, in response, predicting the words most likely to appear next. In some cases, the words most likely to appear next may not be the most factually accurate. For this reason, you should not rely on the factual accuracy of output from our models.

<https://openai.com/policies/privacy-policy> (last visited April 2024 and reflecting a Privacy Policy updated on Nov. 14, 2023 and effective on Jan. 31, 2024). The version of the policy in effect when *Mata* attorney relied on ChatGPT did not have this disclaimer. The motion to dismiss in *Mata* was filed on January 13, 2023, and the brief opposing it was filed on March 1, 2023. 2023 WL 4114965, at *2. [The WayBack Machine](#) has links to the relevant ChatGPT policy as far back as February 27, 2023. The policy then in force (and dated Sept. 19, 2022) did not have the “note about accuracy” disclaimer.

<https://web.archive.org/web/20230227230602/https://openai.com/policies/privacy-policy> (last visited April 2024). The essence of that disclaimer appears to have been added as of April 27, 2023.

<https://web.archive.org/web/20230601012741/https://openai.com/policies/privacy-policy> (last visited April 2024).

8. The United States Patent and Trademark Office recently issued detailed guidance on the use of artificial intelligence tools by attorneys and others practicing before it. *Guidance on Use of Artificial*

Intelligence-Based Tools in Practice Before the United States Patent and Trademark Office, 89 Fed. Reg. 25609 (Apr. 11, 2024) (“PTO AI Guidance”). It discusses some issues that are not addressed in this opinion, including patentability and the potential violation of export control and national security regulations. Attorneys who practice before the PTO or whose practices expose them to data that is subject to export controls or national security restrictions should study that guidance.

9. At the time of this writing, the D.C. Court of Appeals is considering a [proposal](#) that would add a reference to technology to the comments to Rule 1.1. Under that proposal, Comment [5] to Rule 1.1 would be amended to add the underscored language:

Competent handling of a particular matter includes inquiry into and analysis of the factual and legal elements of the problem, and use of methods, procedures, and technology meeting the standards of competent practitioners. It also includes adequate preparation and continuing attention to the needs of the representation to assure that there is no neglect of such needs. The required attention and preparation are determined in part by what is at stake; major litigation and complex transactions ordinarily require more elaborate treatment than matters of lesser consequences.

10. This assumes that the original work product is (1) competent and (2) not already in the GAI’s database.

11. ChatGPT’s [Privacy Policy](#) (discussed below) states that it “does not apply to content that we process on behalf of customers of our business offerings, such as our API. Our use of that data is governed by our customer agreements covering access to and use of those offerings.”

12. See [D.C. Legal Ethics Opinion 309, n.10](#) (2001) (noting high bar for waiving confidentiality).

13. As noted in the PTO AI Guidance referenced above:

Use of AI in practice before the USPTO can result in the inadvertent disclosure of client sensitive or confidential information, including highly-sensitive technical information, to third parties. This can happen, for example, when aspects of an invention are input into AI systems to perform prior art searches or generate drafts of specification, claims, or responses to Office actions. AI systems may retain the information that is entered by users. This information can be used in a variety of ways by the owner of the AI system including using the data to further train its AI models or providing the data to third parties in breach of practitioners’ confidentiality obligations to their clients.... If confidential information is used to train AI, that confidential information or some parts of it may filter into outputs from the AI system provided to others.

89 Fed. Reg. at 25627.

14. As defined in the Rules of Professional Conduct, “[k]nowingly,’ ‘known,’ or ‘knows’ denotes actual knowledge of the fact in question. A person’s knowledge may be inferred from circumstances.” Rule 1.0(f); *In re Soto*, 298 A.3d 762, 767 (D.C. 2023) (quoting Rule 1.0(f)). This is a high threshold that would not normally be crossed by a lawyer’s unintended misrepresentation due to the lawyer’s ignorance of the limitations of the GAI or other technological application. To the extent that the misrepresentation could have been prevented through ordinary care like cite-checking, however, the lawyer could still face sanctions, claims or other consequences, as did the lawyers in *Mata*.

15. As noted above, GAI technologies can create false photographs, audio recordings and videos that look or sound very real. These sorts of “deepfake” files have the potential to be used as false evidence squarely within the prohibitions of Rules 3.3 and 3.4.

16. A lawyer may, at her own expense, retain a copy of the client file. See D.C. Legal Ethics Opinion 273 (1997); D.C. Legal Ethics Opinion 250 n. 2 (1994).