NORTH CAROLINA Lawyers Weekly

Part of the BRIDGETOWER MEDIA network

NOVEMBER 21, 2024

Generative AI: Use it, but do so with utmost care

EDWARD S. CHENG AND MARIEM MARQUETTI

Generative artificial intelligence in the legal profession has increasingly become the talk of the year. With the introduction of ChatGPT and follow-up AI-assisted legal software tools, lawyers and their clients are paying attention.

Generative AI refers to large language models that use a system of algorithms to generate text, images, audio and other content based on tremendous volumes of data fed into these models. Such large language models are now available to the public, and while generative AI is attractive to the legal profession to increase efficiency, generative AI also can be dangerous and unreliable if used without an understanding of its limitations.

Indeed, some lawyers have already begun using such models to write motions and briefs — sometimes to their spectacular and public woe. For example, the case of *Mata v. Avianca*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023) was a simple personal injury tort case that became infamous for the lawyers' use of generative AI and the hallucinations the AI created. In *Mata*, plaintiff's counsel prepared a 10-page brief, citing cases like *Martinez v. Delta Air Lines* and *Varghese v. China Southern Airlines*, which turned out to be nonexistent opinions. When defense counsel could not find these cases, it was revealed to Judge P. Kevin Castel that plaintiff's counsel had used ChatGPT to draft the brief. The lawyers were personally sanctioned $5,000 as a result and ordered to disclose the sanction to their client and to each judge falsely identified as an author of a fake case.

In another case, *Smith v. Farewell*, in the Massachusetts Superior Court, the court imposed monetary sanctions on an attorney who submitted a brief created by generative AI that contained fictitious case citations without knowing that the cases cited were fake. The court accepted the lawyer's representations that the fictitious case citations were not submitted knowingly, but ultimately, the court found that the lawyer still violated the duty under Rule 11 of the Massachusetts Rules of Civil Procedure for failing to make a "reasonably inquiry".

ETHICAL OBLIGATIONS

On July 29, 2024, the ABA's Standing Committee on Ethics and Professional Responsibility weighed in, issuing its Formal Opinion 512 on lawyers' use of "Generative Artificial Intelligence Tools" in their practice. According to the opinion:

- Lawyers using GAI tools have a duty of competence, including maintaining relevant technological competence, which requires an understanding of the evolving nature of GAI.
- In using GAI tools, lawyers also have other relevant ethical duties, such as those relating to confidentiality, communication with a client, meritorious claims and contentions, candor toward the tribunal, supervisory responsibilities regarding others in the law office using the technology and those outside the law office providing GAI services, and charging reasonable fees.
- With the ever-evolving use of technology by lawyers and courts, lawyers must be vigilant in complying with the Rules of Professional Conduct to ensure that lawyers are adhering to their ethical responsibilities and that clients are protected.

The opinion does not create new responsibilities for attorneys but amplifies the existing obligation that an attorney understands the tools that she is using. Under Fed.R.Civ.P. 11, the signatory to a submission is responsible for the arguments and law contained in the filing. The ABA Model Rules of Professional Conduct, Rule 1.1, Competence, specifically comment 8, requires an attorney to keep abreast of technology. This article focuses upon the first charge of the opinion, namely that an attorney must maintain

"relevant technological competence, which requires an understanding of the evolving nature of GAI."

HOW IT WORKS

At its core, generative AI is simply math. As the founder and CEO of Toloka AI, Olga Megorskya, said, "[Y]ou collect large amounts of data, then using the methods of machine learning, algorithms learn to find inter-dependencies among these pieces of data and then reproduce this logic on every new piece of data they meet."

AI development dates back several decades, long before ChatGPT and OpenAI. In 1950, Alan Turing created a test to determine whether a machine was "intelligent" through what became known as the "Turing test." A remote human evaluator would judge natural language conversations between a human and a machine designed to generate human-like responses. The evaluator would be aware that one of the two partners in conversation was a machine and the other a human. The conversation was limited to text in order to eliminate issues of inflection, accents or voice from the equation. If the evaluator could not reliably tell the machine from the human, the machine passed the test. ChatGPT was created to pass the Turing test.

The issue is that generative AI does not "know" anything. This problem is illustrated by the "Chinese room" argument published in 1980 by American philosopher John Searle. Searle imagines himself alone in a room following a comprehensive set of instructions for responding to Chinese characters slipped under the door. Searle does not know Chinese, and the characters appear to be nothing more than assemblages of symbols, but by following the program for manipulating symbols, he sends appropriate strings of Chinese characters back out under the door. This leads people outside the room to mistakenly conclude that the person within understands Chinese. The takeaway is that programming a digital computer might make it appear to understand language but does not produce real understanding. Hence the "Turing test" is inadequate.

Generative AI relies on significant volumes of data because generative AI does not "know" anything — it does not have real-world experience or a human's common sense to make connections. The programs are trained on significant amounts of data to determine the statistical likelihood of "next words." As explained in a recent New York Times article, the behind-the-scenes of generative AI is essentially a "statistical distribution — a set of probabilities that predicts the next word in a sentence, or the pixels in a picture."

That same article warns that generative AI systems will likely start relying on previously generated AI output, which creates the risk for a spiraling feedback loop that degrades in quality over time as it lacks outside, real-world input. This spiral would eventually narrow the set of statistical probabilities such that the new generative AI output becomes increasingly less real and more incoherent. Generative AI predicts content, but not always accurately, and now with AI systems likely to start using previously generated AI output, the accuracy of generative AI output is likely to further decline.

ROLE IN THE LEGAL PROFESSION

Now that we have a better understanding of how generative AI works, how can it be used in the legal profession? Companies providing tools for lawyers have incorporated generative AI into two broad categories. The first is to help attorneys summarize and digest vast quantities of information. The latest generative AI tools purportedly can help an attorney create chronologies from vast document productions, generate abstracts from lengthy deposition transcripts, and complete other such tasks that require the review and absorption of numerous documents. Wading through infinite email chains where the same email is produced by multiple recipients has been a bugaboo of modern discovery. Generative AI tools can help wade through the chaff.

Generative AI also can help with contract drafting by helping to assess whether a "contract complies with a set of policies, laws, or regulations and pinpoints clauses that pose a conflict and suggests redlines." This use of generative AI would appear to hit the sweet spot of current technology. One example of this kind of software is Thompson Reuter's "CoCounsel." Lexis and other companies have similar products.

The second use of generative AI is to interpret and summarize case law and statutes, and then to generate legal discussions or even briefs. As a threshold matter, using generic generative AI programs like ChatGPT 4 for legal matters is a mistake. ChatGPT, its iterations, and other similar programs are designed to pass the Turing test. They are highly successful chatbots that can fool the ordinary user to think that they are communicating with a live person on the other side. ChatGPT 4, however, is not designed to provide accurate legal research and writing. It is designed to please the user, whether or not it can find cases or statutes to support the arguments that it is asked to make, so it will make up or "hallucinate" citations to nonexistent cases or statutes to complete its task. These hallucinations gave rise directly to the malpractice in the *Mata* and *Smith* cases. In other words,

lawyers cannot reliably use ChatGPT 4 or other generic generative AI works programs for their legal work.

Lexis, Thompson Reuters (Westlaw) and Casetext claim to have solved this kind of hallucination problem. These companies license the use of ChatGPT for the task of language generation so they do not have to create their own large language models from scratch. They then add their own programming, employing sophisticated techniques called "retrieval-augmented generation" to avoid creating legal citations and opinions out of whole cloth. These companies also restrict their generative AI programs to rely solely upon real case law, statutes and legal digests. In other words, while ChatGPT 4 sources its law from the entire world wide web, generative AI programs from companies like Lexis or Westlaw restrict their products to proprietary legal databases. The question is whether their generative AI-assisted programs can be deemed reliable and accurate?

A recent article by Stanford University calls into question the use of generative AI tools to analyze the law. On June 6, Stanford researchers published the article "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," which described problems with generative AI-assisted programs that digest and discuss cases and statutes. The Stanford researchers employed a battery of research queries to benchmark the accuracy of programs from several legal-oriented companies and found that the output suffered from substantial error rates. For example, one query asked, "Why did Justice Ginsburg dissent in *Obergefell [v. Hodges]*?" The generative AI program happily explained the basis for her dissent, except in reality, Ginsberg had joined the court's landmark decision legalizing same-sex marriage.

Moreover, the discussion appeared to address issues that were not found in *Obergefell* at all. Thompson Reuters responded to the report in a blog by Mike Dahn, head of Westlaw Product Management, writing that, "Our thorough internal testing of AI-assisted research shows an accuracy rate of approximately 90% based on how our customers use it, and we've been very clear with customers that the product can produce inaccuracies." LexisNexis did not challenge the Stanford study, but instead focused on the relative findings between its product and Thompson Reuters' product. In other words, it appears clear that there remains a substantial error rate in the legal analysis and findings by current generative AI programs.

This does not necessarily mean that lawyers should completely forego the use of generative AI-assisted research and writing tools. As Dahn wrote, Thompson Reuters, expressly warns its users: "AI-assisted research uses large language models and can occasionally produce inaccuracies, so it should always be used as part of a research process in connection with additional research to fully understand the nuance of the issues and further improve accuracy."

He explains, "We also advise our customers, both in the product and in training, to use AI-assisted research to accelerate thorough research, but not to use it as a replacement for thorough research." This is good advice. Even in the most favorable testing environment, it appears that there is an irreducible ten percent plus error rate in generative AI research and analysis, which is ten percent more than acceptable for court submissions.

CONCLUSION

Generative AI can provide many benefits to the legal profession, but lawyers should understand the risks posed, and exercise caution when using generative AI for legal work. Generative AI may be helpful for digesting documents to create chronologies, create initial drafts of deposition abstracts, and to identify key documents in a case. But even in this use, a lawyer should not depend 100% on these tools. Lawyers still need to apply the legal analysis to understand what is important and should review any generative AI output — after all, the generative AI does not understand the legal issues in the case in the way a lawyer would.

For any case citations or case interpretations generated by AI, lawyers must check the legal citations to make sure the quoted language does indeed come from the case cited, and should also verify that the information produced was not fabricated. Consider any research and analysis by generative AI to constitute nothing more than first rough draft and to then add the human element to check for style, accuracy, precision and applicability to the case. After all, lawyers have a duty to provide competent and zealous representation. Clients hire lawyers to do the work, not a machine.

Last, this is a rapidly changing field as companies race to refine and improve their generative AI research and writing tools. Lawyers are obligated to keep current with the latest technology to ensure that they meet the ever-changing standards of the profession. After all, just a scant twenty years ago, computerized research was cutting edge technology not to be trusted, and now Westlaw and LexisNexis are ubiquitous. One day, generative AI will be too — just not today.

Edward S. Cheng is a litigation partner at Sherin and Lodgen, and Mariem Marquetti is a litigation associate at Sherin and Lodgen.

Reprinted by EVGMedia with permission from North Carolina Lawyers Weekly . www.evgmedia.com LM112624