CLCT ARTIFICIAL INTELLIGENCE NEWSLETTER

Welcome

By Jennifer Heeg
&
Courtney Lange

Editors

WILLIAM & MARY LAW SCHOOL

Welcome to the first edition of William & Mary's Center for Legal & Court Technology (CLCT) Artificial Intelligence (AI) newsletter. As a student publication supervised by CLCT faculty and staff, we anticipate publishing one issue per semester covering ways in which AI-related issues may affect courts, attorneys, and the general practice of law.

CLCT's primary mission is to improve the world's legal systems through appropriate technology. At CLCT, students, faculty, and staff conduct research to determine the likely effects of technology on the legal professions and practice. We welcome your feedback and suggestions for future newsletter topics. (AInewsletter@wm.edu)

Overview

Problematic AI Symposium

On February 10, 2023, CLCT hosted the Problematic AI Symposium, with support from the Coastal Virginia Node of the Commonwealth Cyber Initiative (COVA CCI) and the University of Montreal's CyberJustice Laboratory. Panelists from industry, government agencies, and academia discussed the state of AI from the perspective of errors and risks associated with the technology. The panelists put forward their ideas and debated risks unique to self-developing algorithmic structures as well as approaches to identifying and mitigating those risks. Far from stopping at the technical water's edge, industry representatives presented their organization's evolving internal policies and frameworks for addressing ethical and legal issues. Government agency representatives discussed the current status of regulatory approaches, including the recently-released NIST's AI Risk Management Framework. The Summary Conference Paper can be accessed on CLCT's website.

ChatGPT

image by vecstock

In 2022, OpenAI, one of the most prominent AI companies, made available its new version of the Chat Generative Pre-trained Transformer (ChatGPT) to the public. The company focuses on developing tools to utilize AI to generate content, including visual works (through tools such as DALL-E and DALL-E 2) and written works (through ChatGPT). ChatGPT is a Large Language Model (LLM) whose latest version was trained on an "internet scale" data set, being sufficiently large to be a "representative sample of all English speakers on the [I]nternet," as noted here. It has a wide range of capabilities, including serving as a chatbot to assist with online services and composing legal documents (e.g., contracts and pleadings). Although this LLM has powerful use cases, ChatGPT introduces various legal concerns, especially in the areas of intellectual property. The AI model's actual utility in legal practice is open to significant debate and controversy. As this newsletter goes to press, the FTC is investigating ChatGPT's inaccuracies.

Problematic AI Symposium

From Principles to Principals—Establishing Safe Harbors & Regulating the Ethical Development of Artificial Intelligence Systems

By Jeremy Bloomstone

Across the public and private sectors, stakeholders have increasingly recognized and coalesced around foundational principles that should guide the creation and adoption of Artificial Intelligence (AI) systems. While the terminology differs subtly across industry standard-bearers, they all converge around the following concepts: fairness and inclusivity; robustness and reliability; privacy and security; and transparency and explainability. Corporate initiatives to translate these value statements into effective governance structures, impact assessments, and deployment practices are proactively shaping the regulatory environment and legal landscape for addressing the rapid innovation and expansion of these predictive and generative AI tools. With the introduction of new practical frameworks for managing AI risk and developing analysis of the cascading policy proposals focused on risk-based mitigation and evaluation, the tide is quickly rising on the need for regulation.

In the market for AI analytics, tools and systems connect diverse actors from vastly different sectors of the global economy. Governments seek to leverage AI-powered insights in their administrative and governance functions. The relationship between innovative, profit-driven vendors and demanding competition-driven customers will soon strain traditional concepts of contractual and product liability. Ethical responsibility as bespoke development and mass deployment of AI resources will also continue to accelerate. Regulation in this space should address two fundamental questions: who bears responsibility when AI systems become problematic in decisions, applications, and operations; and when and where should liability attach in the AI lifecycle?

Dennis Hirsch, Professor of Law and Director of the Program on Data and Governance at the Moritz College of Law at The Ohio State University, suggests industry executives view responsible management and the ethics of AI through the lens of corporate sustainability, as opposed to rote compliance. This recognition of responsibility influences corporate decision-making and shapes companies' AI development processes to reduce regulatory risk, build and sustain trust, retain employees, improve quality and competitiveness, and demonstrate company values. But studies have also shown that looming regulations on AI also affect corporate decision-making in ways that reduce the risk tolerance of managers and the internal priority for ethical product development and adoption of AI. Hirsch's analysis of the developing trends in law and policy, as well as management strategies for responsible AI management, lead to a forward-looking conclusion: future regulatory proposals should prioritize shielding the incentives and impulses to innovate while advancing procedural mechanisms for holding developers accountable to the principles they proclaim publicly and avoiding burdensome and ineffective obligations.

As pressures mount for sustained innovation in AI, legislators and regulators in the US and EU might turn to check-the-box style compliance measures, which fail to reflect the

active and innovative governance many companies already leverage in the design, development, and deployment of AI technologies. Creating safe harbors from liability for companies who commit to a reporting, disclosure, and monitoring scheme would move the needle beyond self-regulation. Such safe harbors also recognize and capitalize on the investment and leadership of AI developers in committing to ethical practices. Functionally, this approach would instill accountability for adhering to the ethical principles. Organizations are already publicizing while also leveraging audits and impact assessments at various stages of the AI development lifecycle. Organizations would be forced to share liability for all contracting parties and stakeholders involved in designing, developing, and monitoring any AI solution or system brought to market.

Fundamentally, a safe harbor approach ensures the process of protecting principle and would allow organizations some flexibility in how they structure their oversight. Industry leaders like Cisco, Hewlett Packard Enterprise, IBM, Microsoft, and Google have already set-up internal procedural infrastructure with review boards, processes for identifying and escalating uniquely risky projects, and evaluating in real-time potential and system failures. Recognizing and distinguishing structural harms and acute personal injuries from AI decision-making is critical to this scheme because these harms and injuries require diverging regulatory approaches. But ensuring regulators, key stakeholders, and (the public?) have access to information throughout the AI design, development, and deployment lifecycle can help not only curtail potential abuses arising out of AI deployment but also inform awareness of market participation and the potential need for sector-specific regulation.

The discussion throughout CLCT's Problematic AI Symposium highlighted the competing perspectives, public concerns, and geopolitical pressures to calibrate legislation and regulation in this evolving space. A first step should be formalizing accountability, responsibility, and liability for corporate best practices and incentivizing their wider adoption by other organizations seeking to develop ethical AI technologies or responsibly integrated AI components into their business operations.

Data Poisoning

By Daniel Wicklund

The Problematic AI symposium focused on whether and how AI systems ought to be used, given the near certainty that present technology will produce errors, often in unpredictable ways at unpredictable times. As CLCT Director and William & Mary's Professor of Law Fred Lederer noted in his symposium introduction, the alternative to the use of AI usually is relying on fallible human decision-making.

The Symposium's first panel discussed why AI produces erroneous results. Assuming that the AI algorithm is prepared correctly, the challenge in producing accurate AI results is correctly training the model. Training requires a large amount of accurate data. Consequently, inadequate training tends to produce errors, raising the possibility of biased AI conclusions. Furthermore, data poisoning—the intentional or unintentional addition of bad data into otherwise benign training sets for AI—increases the likelihood that AI will make incorrect decisions. Accordingly, tool developers must verify the underlying AI system

training data at the very beginning, especially in contexts like healthcare, which involves sensitive personal health information and large risks. In diagnostic health care tools that utilize AI, poor training data risks harm arising out of a false positive diagnosis (creating over-treatment) or a false negative diagnosis (preventing the administration of necessary treatment).

During the Symposium, Blake Anderson, Cisco's Principal Engineer, emphasized securing data from the beginning of the AI development process to mitigate risks of data poisoning attacks. Currently, developers often exercise little oversight of training data. One proposal from the IEEE Journal of Biomedical and Health Informatics seeks to help secure data and increase transparency in the event of poisoning attacks. It suggests "periodically constructing a model using the training dataset, evaluating its accuracy on the validation dataset, and raising an alarm in case of any suspicious change in the accuracy metric." In doing so, users can make accurate determinations with less risk of error. This method has been successful in preventing three types of poisoning attacks: attacks that do not have access to the AI learning sets, attacks that are not aware of the underlying algorithms the AI uses, and attacks which focus on AI with algorithms that create predictable outcomes (e.g., using attacks to generate fake clinical reports for evaluating breast cancer on clump thickness). However, these algorithms still rely on a base dataset, which may make these evaluation methods unable to detect underlying data biases that may skew the results.

The European Union proposed an enforcement scheme, the AI Act, which, among other things, sets forth standards to avoid unacceptable risks undertaken by companies employing AI in healthcare contexts. This proposal will also aid in preventing data poisoning in high-risk areas, such as healthcare, by imposing obligations on AI tool developers. In particular, the regulations require "high risk" industries to:
- implement a quality assurance system,
- keep up-to-date technical records,
- undergo initial assessments and update the assessments after modification of the system,
- register the AI system with the government,
- collaborate with "market surveillance authorities," and
- "inform the provider or distributor about any serious incident or any malfunctioning."

This proposed regulation also requires the AI tools themselves to "use high-quality training, validation and testing data," "draw up technical documentation & set up logging capabilities," "ensure appropriate degree of transparency and provide users with information on capabilities and limitations of the system and how to use it," "ensure human oversight," and "ensure robustness, accuracy and cybersecurity."

However, in the United States, there is little regulation of AI. There are proposed risk management guidance with the AI Bill of Rights and the recently-announced NIST framework, which could become the basis of a regulatory framework in the future. Many of the principles in the AI Bill of Rights are similar to those in the proposed European Union regulatory framework. For example, to protect the public from harm, the AI Bill of Rights suggests the use of consultation, testing, risk identification, and mitigation in a proactive and ongoing manner, ongoing monitoring, and clear organizational oversight.

The Federal Trade Commission produced similar guidance as to what companies can do to mitigate risks from AI use. One important principle—increased transparency— eliminated racial bias in an AI healthcare algorithm in at least one instance.

The Federal Reserve has issued the Guidance on Model Risk Management, designed to be used by banking organizations to assess AI risk stemming from adverse consequences in financial decision-making. The Guidance focuses on providing "critical analysis by objective, informed parties that can identify model limitations and produce appropriate changes" to improve the validity and accuracy of these decisions and emphasizes that "where models and model output have a material impact on business decisions … a bank's model risk management framework should be more extensive and rigorous." Ultimately, the federal government needs to enact effective laws to meet the ongoing challenges brought by AI. Government needs to continue to pursue a proactive approach to mitigate a range of risks, including the risk arising out of data poisoning while promoting the ethical use of AI technologies.

Symposium Wrap-up

Workshop findings included:
- The Problematic AI topic also needs to consider Problematic Humans and how they use AI;
- Corporations already govern and manage AI risk, but there are few regulations, standards, and best practices from which to draw;
- What baseline do we use to assess AI's performance; an expert human, or a theoretical ideal?
- Intellectual Property challenges to model training datasets exist and are being contested in the courts;
- General regulation is needed but also should be supplemented with domain-specific regulation (e.g., Healthcare, Finance, etc.) ;
- There are domains where AI is useful, but there are also domains where AI may need to be prohibited, and the domains are yet to be decided;
- There will be cultural differences in perspectives of AI governance; how does the global community deal with these differences?

A 2024 Symposium to continue this conversation is being planned.

Blogs by panelists:
- Peter Chapman: Can humans catch up to AI? https://articleoneadvisors.com/can-regulation-catch-up-to-ai/

ChatGPT

IP Concerns with ChatGPT

By Mike Papakonstantinou

Transformative technologies can disrupt society and challenge existing norms, including how to reconcile legal rights in a new era of innovation. Just as previous digital technology advancements presented Intellectual Property (IP) challenges during the Web 1.0 and Web 2.0 revolutions, ChatGPT reintroduces the issue of how society must resolve

existing IP rights, including copyright, patent, and trademark, in the present disruptive technological age.

Copyright

Overview

ChatGPT implicates various IP copyright issues, including in the areas of infringement and ownership. As noted in previous articles, OpenAI used voluminous datasets to develop ChatGPT. Some sources include Wikipedia, various digitized books, and sources across the web. However, critics allege that the input data used to train ChatGPT-like systems may constitute copyright infringement at a very large scale because it was used without the permission of the copyright holders. In addition, even though Wikipedia already provides licensing options to utilize its copyrighted text, the licenses typically require some form of attribution or linking. At this time, it appears that ChatGPT does not cite Wikipedia or other sources regularly or accurately.

In the United States, copyright laws grant the owner of a copyright the exclusive right to reproduce or distribute their work, to prepare a derivative work, and other associated rights. However, without large amounts of quality training data, these AI systems simply perform poorly or are ineffective.

Under the fair use doctrine, any usage of copyrighted works to develop or enhance ChatGPT may be acceptable, because the fair use doctrine permits non-licensed uses of copyrighted works in some scenarios. A common example of fair use is in an educational context, such as when a professor distributes a relevant article to students for the purposes of teaching or facilitating scholarship. While courts consider various factors to determine fair use, the inquiry ultimately may turn on how "transformative" the use is. Transformative uses add something new, with a different character or purpose, and do not substitute for the original use of the work. Consequently, even though OpenAI commercialized ChatGPT and it no longer remains a research tool, it may still qualify as a fair use. The copying of data for AI training is non-expressive, making the resulting output highly transformative of the original works. The purpose of generative AI (synthesizing new information) is highly different from the original purpose of the original works (expressing an idea of the work's author).

Analysis

Generative AI systems, which can create novel content, may infringe the rights granted to a copyright owner in various ways. For example, when a company like OpenAI creates a database with training data, it is likely making a copy of the work for inclusion in the database, which may violate the right of reproduction. In addition, if ChatGPT creates a new work based on its training data and a user's prompt, the resulting new, generated work may be a derivative work. In this case, there may still be copyright infringement because a copyright owner of the original work (from the training data in this example) has the exclusive right to a derivative work, subject to any fair use defense claims.

Furthermore, if ChatGPT outputs a non-de minimis amount of copyrighted text, such as from a textbook or novel, in response to a user's prompt, then that may constitute distribution. Even if the system quoted text and properly cited the copyrighted work ("attribution"), attribution alone cannot protect an infringer from copyright infringement liability. There seem to be some safeguards built into ChatGPT to prevent the widespread

dissemination of copyrighted materials. For example , at least one outlet reported that ChatGPT rejected direct prompts to output copyrighted works.

Moreover, works generated by ChatGPT may present novel legal questions of authorship and ownership. To receive copyright protection (and thus create an ownership right in the copyright), a work must meet the authorship requirement. Based on current U.S. Copyright Office policy and case law, the Office likely will reject an application for copyright if a machine or algorithm is listed as the author. However, the question of what copyright protection exists for a work created by AI with human involvement is a different legal issue. This inquiry may turn on the involvement of the human in the authorship of the work. Margaret Esquenet, a partner at IP boutique law firm Finnegan, opined that AI-generated works might either constitute a work in the public domain or a derivative work of materials in the training data.

In addition, the U.S. Copyright Office has specified its intentions to focus on legal uncertainties involving technology and copyright in 2023 in light of rapid technological developments. Recently, the Office published a notice indicating that AI-assisted works are eligible for copyright protection if there is sufficient human authorship. One example provided by the Office is when a human arranges or selects AI-generated content in a sufficient manner that the overall work meets the authorship requirement. In February 2023, the Office also determined that individual AI-generated illustrations and images utilized by a human author in a novel did not receive copyright protection, though the text (which was entirely human-written) and the overall work did receive a copyright. Although the Office intended to clarify these murky legal questions with its recent actions, there is still confusion regarding the boundaries of copyright eligibility for AI-assisted works. Attorneys predict that courts still will need to make determinations in individual cases and to provide more clarity on the Office's guidance.

Even though OpenAI's terms of service indicate that the requestor/end-user ultimately receives "right, title and interest" to the resultant work, it still may not meet the standards for authorship to receive copyright protection. ChatGPT end-users should understand that the system's outputs may not result in works that are eligible *currently* for copyright protection under existing laws and policies. This may change as the Office and courts provide more clarity within the generative AI domain. More guidance is necessary to protect human-made creations and to inform how to protect AI-assisted creations as well.

Patent

Overview

ChatGPT implicates issues in patent law as well. Given the huge training data set, algorithmic sophistication, and vast computing resources, it is possible for ChatGPT to create inventions. However, the Federal Circuit last year held that AI could not be the inventor of a patent, as an inventor must be a natural person. This ruling notably contrasts with South Africa's patent office decision, which granted a patent to the AI system for the same invention.

Analysis

The patent at issue in both the American litigation and the South African patent application listed AI as the sole inventor, but AI-augmented inventions generated by ChatGPT present different legal issues. For example, it is unclear how central natural

persons were to the underlying invention when humans use generative AI systems. If ChatGPT does most of the inventive work, then the person that prompted the inventive output "may not be able to take the oath required by the patent office that they are the rightful inventors," as noted here. If natural persons did not contribute to the invention, then ChatGPT likely alone created the resultant invention, which, under last year's Federal Circuit decision, bars patentability. The United States Patent and Trademark Office (USPTO) has solicited input from stakeholders on the issue of AI and inventorship. Although the USPTO cannot overrule the Federal Circuit's holding, it can suggest any potential changes to federal copyright and IP laws. Mark Lemley, a Stanford Law professor and Lex Luminia's counsel, opined that Congress must amend the statute because "AI is engaging in significant inventive activity…the PTO and the courts have to pretend that activity was done by a human, or conclude that the invention isn't patentable at all because it was done by an AI."

In addition, one requirement for patentability is non-obviousness based on a "person having ordinary skill in the art" (PHOSITA), as noted here. A person using ChatGPT to invent increases likely the knowledge and skill level of a PHOSITA due to the system's vast training data across numerous technical fields and computing resources, essentially merging various "arts" into one pool of data. Increases in what knowledge and skill a PHOSITA has raise the threshold for what is non-obvious, potentially inflating the barrier to patentability. Patent law, consequently, must clarify what non-obviousness entails in the context of AI-augmented inventions. Without such clarity, individuals and companies will operate in uncertainty as they pursue patent protection for AI-assisted creations.

Trademark

Overview

Unlike copyright and patent law, it is "legally irrelevant" who or what creates a trademark. Consequently, tools like ChatGPT can help individuals and companies in generating marks eligible for federal trademark protection. Also, practitioners already utilize AI tools to augment their trademark practices, such as using automated systems and AI-assisted methods for client counseling. Some AI tools, such as Corsearch, check the USPTO trademark registry to spot existing trademarks that can pose a bar to registration for a new product name or brand. Such tools allow for quick processing to facilitate review by a practitioner, who ultimately provides legal counsel to clients. Corsearch uses various factors, such as risk of similarity between a desired trademarks and registered trademarks. Other existing AI tools help trademark owners automatically comb the web to flag potentially infringing products or counterfeit goods to enforce trademark rights.

Analysis

ChatGPT can help attorneys by generating ideas for trademarks. For example, Ashley G. Kessler, a trademark attorney at Cozen O'Connor, recently described using ChatGPT to identify brand names for a client. Although Generative AI is helpful, the technology is not a complete replacement for attorneys in this context. For example, of the ten names generated by ChatGPT, two conflicted with existing registrations at the USPTO.

Also, some names were already in use based on web searches, indicating that some entities already had unregistered, common law rights to those names. As a result, tools like ChatGPT can augment the brainstorming process by generating possible brand names, even if the present state of the technology does not generate perfect outputs. One potential improvement to ChatGPT could be "incorporating access to the USPTO's trademark registry," so that the system could check the registry and generate a name that is eligible for protection.

One potential issue with harnessing ChatGPT in trademark practices is client confidentiality. OpenAI's policies indicate that their systems collect data provided by the user, including the prompts. If an attorney prompts ChatGPT for trademark ideas, that means that the system will collect a client's confidential information, which may be an improper disclosure of client information by the attorney. A potential remedy is for practitioners to utilize a tool specifically made for attorneys that operate similarly to ChatGPT. For example, Allen & Overly recently announced a chatbot to assist attorneys with various legal tasks. This kind of tool can provide similar functionality for generating trademark ideas as ChatGPT, but potentially with enhanced data protection stemming from the tool's usage.

Furthermore, OpenAI recently released the ChatGPT API, which does not use customer data or queries to improve its models. As a result, tools that harness the API likely do not present the same client-confidentiality concerns as directly querying ChatGPT. Large law firm DLA Piper also announced its adoption of a tool powered by the latest version of ChatGPT, potentially signaling that prior confidentiality issues may now be mitigated. Although attorneys can harness effectively AI outputs to provide more creative trademark counsel to clients, they should be vigilant in minimizing risks of breaching inadvertently client-confidentiality.

Revolution or Routine? The Possible Impacts of ChatGPT on Legal Practice

By Andrew Heiser

Introduction

ChatGPT's debut has given rise to a great deal of discussion of its effects on a variety of areas of life, from its ability to help write social media posts to its potential to change how academic tests are taken. It is no surprise, then, that there is already debate about its possible impact on the legal profession and the practice of law. Opinions have ranged from speculation that ChatGPT will eventually replace lawyers to complete dismissals of ChatGPT's impact on the profession. The truth appears to be somewhere in the middle. ChatGPT will not replace lawyers, but several potential uses in the legal profession are already being explored. Some, however, have suggested that technical limitations and ethical and privacy concerns may limit its immediate implementation in the legal field.

ChatGPT as Legal Counsel

The rise of ChatGPT has led to speculation that it (or similar technology) might act as an alternative to human lawyers sometime in the relatively near future. One legal startup went so far as to put forward a plan to have ChatGPT (and other text generators) "represent" a client in traffic court by dictating their response through headphones, with supposed plans to eventually test the system before the U.S. Supreme Court. However, the plan was dropped after it allegedly ran into legal threats from multiple state bars, and it is unclear whether the plan was ever permissible under current court rules.

More successful experiments have involved ChatGPT taking legal exams, but even these suggest serious limits compared to human lawyers. ChatGPT was able to pass a series of legal exams provided by professors from the University of Minnesota, performing comparably to law students whose answers were also graded in a blind trial. Notably, though, it only received a C+ level on average on these exams, which would result in academic probation for a real law student. Its overall performance was very erratic: it would perform very well on one essay question, only to be completely off topic on the next. And while it showed a strong grasp of some basic legal issues, its application of legal rules was often superficial, and it struggled with issue-spotting in more open-ended prompts. In another study, ChatGPT passed the evidence and torts portions of the bar exam. However, it failed to pass the bar exam as a whole, and furthermore did generally more poorly than human exam-takers on average.

Taken together, these observations challenge the idea that ChatGPT could act as a legal counsel in the near term. Indeed, it is arguable that the kind of information retention and production that ChatGPT excels at isn't a key part of the modern lawyer's toolset, despite being a skill tested by both the bar and law school exams. Lawyers already rely on technology to look up specific rules, statutes, and regulations in the modern, complex legal environment. The real skills of practice involve beyond legal research and writing, such things as client interaction, identifying new avenues of research, tackling complex problems, and knowing when to leverage certain research tools.

ChatGPT Supplementing Legal Practice

Despite the limitations discussed above, some legal commentators believe that ChatGPT might be a useful tool for assisting lawyers in a variety of practice areas. A paper co-written by Andrew Perlman of Suffolk University Law School and ChatGPT (cited as OpenAI's Assistant) identified four possible uses of ChatGPT for lawyers: assisting in legal research by scanning large amounts of text data automatically; generating automatically legal documents, such as contracts and briefs; providing general legal information to the public; and supplementing a lawyer's legal analysis by providing its own advice and suggestions.

A test of some of these capabilities in the context of family law seemed to show some promise. ChatGPT was able to handle successfully various tasks, including offering sound advice on a request for production involving cryptocurrency, drafting a workable email to an opposing council, creating a usable financial information document for a client to fill out, and analyzing data to spot financial anomalies.

The National Law Review suggests several other ways ChatGPT can aid a law firm with other areas of its business, such as writing ad copy, planning a marketing strategy, helping in search engine optimization, and planning the design of its website. While these

uses might not be directly related to legal practice, having ChatGPT facilitate these tasks would potentially open up more of a lawyer's time to work on legal tasks.

Limits of ChatGPT as a Legal Tool

Limited uses of ChatGPT, including those mentioned previously, have been subject to criticism and scrutiny. A test of ChatGPT's capabilities by lawyers from Haynes and Boone LLP for Law360 identified several issues. While ChatGPT was generally accurate in answering questions about Texas procedural law, it was completely wrong in areas of substantive law, at times even fabricating incorrect information. For example, when the lawyers asked it to draft a motion to dismiss for a clearly erroneous complaint, they had to go through a lengthy process of trial and error to find an input that gave a result that addressed all the relevant issues. Even then, the writing was at a substandard level. They concluded overall that while there are uses for AI tools in legal practice, submitting unedited ChatGPT-produced content would violate a lawyer's duty to provide competent and diligent representation. The lawyers noted that when they asked ChatGPT itself about its ability to provide legal services, the system generated a warning that its knowledge is limited, and thus it can give misleading or incorrect answers.

Others have raised more general doubts about the implementation of ChatGPT in legal practice. Bloomberg Law has noted that the black-box nature of deep learning systems like ChatGPT makes it difficult, if not impossible, for a user to understand the underlying model and how it produces its results. For the lawyers, the duty of professional competence to understand these technologies and explain them to their clients may present a hurdle to the use of AI tools like ChatGPT. Furthermore, ChatGPT's use of large corpus of text data for training data raises confidentiality concerns. While ChatGPT professes to be confidential, OpenAI states its staff may review conversations to monitor whether Open AI's terms of use are violated by users. Even if a law firm could secure a confidentiality agreement with OpenAI, the potential of a data breach means that using ChatGPT still raises concerns regarding confidentiality.

Conclusion

There is no disputing that ChatGPT represents an exciting new frontier in AI. But even this new frontier falls far short of some of the media sensation and hype about AI technologies in general. ChatGPT will not replace lawyers, at least in the short term. It currently does not encompass all the necessary skills that define the legal profession. And while it might be useful as a tool in legal practice in some circumstances (e.g., performing legal research and drafting initial documents), the technology's current shortcomings and ongoing concerns about confidentiality and transparency may well limit its widespread adoption in the near term. ChatGPT is a fascinating, useful tool, and lawyers and other legal professionals should certainly follow its development and widespread adoption to consider its potential uses for our field. However, we should also be cautious about adopting hastily it into our legal toolset without further research and consideration of its limitations.