taxnotes federal

Volume 180, Number 5 ■ July 31, 2023

The Ethics of Generative AI in Tax Practice

by Benjamin Alarie and Rory McCreight

Reprinted from *Tax Notes Federal*, July 31, 2023, p. 785

The Ethics of Generative AI in Tax Practice

by Benjamin Alarie and Rory McCreight

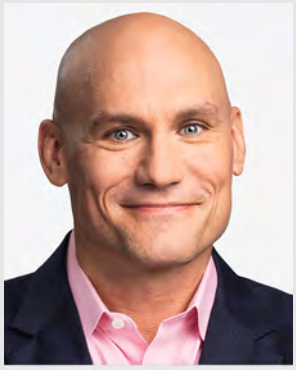

Benjamin Alarie

Rory McCreight

Benjamin Alarie is the Osler Chair in Business Law at the University of Toronto and the CEO of Blue J Legal Inc., and Rory McCreight is a lead analyst at Blue J Legal. They thank Ethan Wilkinson for his support in the preparation of this article.

In this article, Alarie and McCreight examine the ethical concerns arising from the proliferation of generative artificial intelligence tools in tax research and provide recommendations and guidance for legal professionals using AI tools.

I. Introduction

The integration of generative artificial intelligence is beginning to have a profound effect on various industries, including knowledge industries such as law and accounting. Professional services firms are capitalizing on AI technologies for a host of applications, such as tax research, memo drafting, contract analysis, due diligence, document review, predictive analytics, and the list goes on. Recently, AI's sharply increasing competence, most evident with OpenAI's GPT-4, has made it a valuable tool for firms looking to improve their operations and deliver more effective and efficient tax planning and advisory services. In-house tax departments are starting to explore using generative AI to assist them in their work as well.

The AI tools that have become particularly popular recently in the media and within tax practice are based on large language models (LLMs), of which GPT-4 is the leading example. LLMs like GPT-4 are the backbone of applications like ChatGPT and can recognize, predict, translate, summarize, and generate language. ChatGPT is a conversational web-based application of an LLM that uses as its training data text drawn from all over the internet, as well as the user's input, to generate an output. Companies like OpenAI have popularized LLMs by making them easily accessible through web-based chat interfaces.[1]

Despite the power of these new tools, their implementation by professional services firms has been highly heterogeneous, ranging from strident bans to enthusiastic interest, adoption, and even direct investment in their development. The potential equity and efficiency gains of AI in the legal field are extraordinarily promising, as one of us has recently written about in a new book.[2] Still, the adoption of AI tools by professional services firms raises critical concerns regarding data privacy and ethical usage.

This installment of Blue J Predicts provides an in-depth exploration of the ethical concerns arising from the proliferation of generative AI

[1] Molly Ruby, "How ChatGPT Works: The Model Behind the Bot," Towards Data Science, Jan. 30, 2023.

[2] Abdi Aidid and Benjamin Alarie, *The Legal Singularity: How Artificial Intelligence Can Make Law Radically Better* (2023).

tools used for legal research and, especially, for tax research.[3] We examine the challenges concerning the quality and accuracy of output, the potential for biased answers, the lack of verifiability, liability considerations, and privacy risks. We also explore the regulatory measures, technological advancements, and professional solutions being pursued to address these challenges, along with practical recommendations to help tax professionals effectively mitigate risk and safely use AI tools in their work.

II. Challenges

A. Quality and Accuracy

One potential challenge posed by LLMs is the quality and accuracy of the output. A recent news story out of New York has been making waves in the legal community, highlighting the dangers of overreliance on generative AI. Two New York lawyers face potential sanctions after submitting a legal brief in court that cited cases that do not exist. The lawyers defended themselves by saying that they used the generative AI platform ChatGPT to help create their legal submission and did not know that it could fabricate the judicial opinions that it provided as legal precedent. One of the lawyers, Steven A. Schwartz, said he "greatly regrets" relying on the chatbot, explaining that he had never used the platform for legal research before and was "unaware that its content could be false."[4]

This is an example of an AI "hallucination" in which AI generates untrue information that is not backed up by real-world data.[5] The cause of AI hallucinations is multifaceted and not easily overcome with today's available LLM-based systems. While these models have shown impressive results in tasks such as language translation and information retrieval, their ability to produce referenceable and accurate legal analyses is still under scrutiny.

One of the key challenges in ensuring the quality and accuracy of legal research generated by LLMs is the need for careful grounding, training, and fine-tuning. The complexity and nuance of legal language and legal concepts are grounded in an understanding of legal principles, which may not be fully captured by foundational LLMs. The training data used to develop these models must be comprehensive, accurate, and unbiased, which can be difficult to achieve. To mitigate this risk, it is essential to carefully evaluate the quality and accuracy of the data used to train LLMs.

Indeed, ensuring the quality and accuracy of legal research generated by LLMs presents a significant challenge. Even with high-quality training data, errors can still occur during the generation of outputs because of limits on computing power, limitations in the length of the input and output text (the "token" length), the complexity of the language used in legal reasoning, and the simple fact that this technology is still new and constantly developing.[6] Therefore, robust systems must be implemented to monitor and validate LLM outputs to ensure their accuracy. Recognizing the nascent nature of this technology, legal professionals bear the responsibility of thoroughly reviewing and validating the analysis generated by LLMs before relying on it or providing it to clients. This approach not only safeguards against potential inaccuracies but also allows lawyers to harness their creative problem-solving skills when navigating complex legal issues.

The potential for generative AI to produce biased answers is closely intertwined with the issue of poor quality or hallucinated answers. The accuracy of the output generated by LLMs heavily relies on the quality of the input data, which can be flawed or biased, resulting in inaccurate and misleading outcomes. Biases can stem from various sources, including the data set itself, divergences within the data set, or biases introduced during the training process, which can compound over time. An IRS audit study conducted recently revealed a significant disparity, with Black taxpayers between 2.9 and

[3] For an earlier discussion outlining the rise of generative AI being used for tax research, see Alarie et al., "The Rise of Generative AI in Tax Research," *Tax Notes Federal*, May 29, 2023, p. 1509.

[4] Kathryn Armstrong, "ChatGPT: US Lawyer Admits Using AI for Case Research," BBC, May 27, 2023.

[5] Alarie et al., *supra* note 3.

[6] "Why Large Language Models Sometimes Make Mistakes," ZeBall, May 3, 2023.

4.4 times more likely to face audits compared with non-Black individuals.[7] This kind of bias in the data set can perpetuate discriminatory patterns, leading to biased answers.

It is essential to recognize that biases in generative AI outputs can be attributed to the quality and composition of the data used for training, as well as the presence of inherent biases within the legal system.[8] These biases can manifest in discriminatory patterns and perpetuate inequality. Therefore, mitigating bias in generative AI tools requires rigorous scrutiny of the training data, careful selection of unbiased data sets, and continuous monitoring and evaluation to address and rectify any biases that may arise.

A recent study has found that using AI-generated content to train algorithms causes model accuracy to degrade over time and may ultimately lead to model collapse. The study defines model collapse as "a degenerative process where models start forgetting improbable events over time."[9] The model becomes detached from reality as its data set becomes corrupted by AI-generated data. This phenomenon raises serious concerns about the quality and accuracy of the LLMs over time, especially if the data being generated is legal information. If AI models generate incorrect information that is then used to train other AI models, the errors could be absorbed and amplified over time, and become difficult to trace back to their origins.

B. Verification and 'Explainability'

Not being able to fully understand the cause of hallucinations leads us to a related concern with large-scale AI systems: the black box problem. The black box problem is the inability of people, including those working on the AI system, to see how these large AI systems generate their responses.[10] AI tools often fail to explain their outputs and may be unable to cite genuine sources that would support the responses given. Not only does this make verification of the AI-generated responses difficult, but it also undercuts a tax professional's ability to confidently trust and rely on AI.

In response to problems with accuracy and opacity, some judges have issued practice directions and orders stating that lawyers must certify that generative AI drafted no portion of a filing, or that a human checked the accuracy of any language that AI crafted. For example, a federal judge in Texas, Brantley Starr, has mandated that lawyers certify that they have not used AI for drafting case filings without a human ensuring their accuracy because of concerns about AI-generated misinformation and bias.[11] Courts in other jurisdictions, including the Supreme Court of Canada, have already followed suit or are considering doing so, indicating that we can anticipate future comparable guidance around the world.[12]

C. Liability

Other concerns that legal professionals face when using LLMs are accountability and liability. As LLM use evolves, it becomes crucial for legal professionals to determine who is responsible for the output when an error is made, when the output of a model causes harm, or when the model uses data inappropriately. These models could also be found to be in breach of privacy or copyright laws. Recently, two authors filed a lawsuit against OpenAI, alleging a breach of copyright law because of the unauthorized training of their novels in the model.[13]

If a lawyer relies on the model's output for a recommendation that later turns out to be incorrect, it can be difficult to determine whether the error was caused by the model itself, the data that was fed into the model, or some other factor. Ultimately, the lawyer will be held responsible when something goes wrong.

[7] Lauren Loricchio, "Black Taxpayers Much More Likely to Be Audited, Report Finds," *Tax Notes Federal*, Feb. 6, 2023, p. 896.

[8] Aidid and Alarie, *supra* note 2.

[9] Ilia Shumailov et al., "The Curse of Recursion: Training on Generated Data Makes Models Forget," arXiv (May 31, 2023).

[10] Lou Blouin, "AI's Mysterious 'Black Box' Problem, Explained," University of Michigan-Dearborn News, Mar. 6, 2023.

[11] Jacqueline Thomsen, "US Judge Orders Lawyers to Sign AI Pledge, Warning Chatbots 'Make Stuff Up,'" Reuters, May 31, 2023.

[12] Cristin Schmitz, "SCC Considers Possible Practice Direction on Use of AI in Top Court as More Trial Courts Weigh In," Law360, July 7, 2023.

[13] Ella Creamer, "Authors File a Lawsuit Against OpenAI for Unlawfully 'Ingesting' Their Books," *The Guardian*, July 5, 2023.

According to OpenAI's use policy, any information or data uploaded to ChatGPT, and any output generated by ChatGPT, are under the ownership of the user, for as long as it complies with OpenAI's terms of use.[14] This means that ChatGPT assumes no liability for any outputs and provides them "as is."

D. Privacy Risks

When considering the adoption of generative AI in legal practice, law firms must understand the potential privacy risks associated with this technology. The use of generative AI tools, such as LLMs, is not without vulnerability when it comes to maintaining the confidentiality and integrity of sensitive data. There are multiple ways in which generative AI tools can pose risks to data privacy:

- Data breaches: Much like other web-based tools, without adequate security measures, these tools may be vulnerable to data breaches, potentially resulting in unauthorized access to or disclosure of sensitive user information.
- Inadequate anonymization: If these tools access personal or sensitive data for training or generating outputs, and the anonymization techniques are insufficient, there is a risk of re-identification, compromising individual privacy.
- Inadequate data retention and deletion practices: If these tools retain user data for longer than necessary or fail to delete data properly, they may increase the risk of unauthorized access or unintended use of personal information.

The possibility of sensitive information being made public in a data breach is a serious concern for tax professionals. In a recent instance, OpenAI identified and promptly addressed a bug that briefly compromised user privacy by making some chat history visible to other users.[15] The risk of a data leak is amplified when the AI tool uses its interactions with users to further train the LLM. If tools like ChatGPT are trained on personal or confidential information, the AI could inadvertently repeat this information to another user.[16]

Within the legal field, the primary types of high-risk data include confidential information, such as that covered by attorney-client privilege, and personally identifiable information (PII). Other types of confidential information, such as commercially valuable news that has not yet been made public, should also be considered high-risk.

III. Solutions

The following sections give an overview of regulatory, technological, and professional solutions to some of the ethical challenges discussed. They also provide recommendations and guidance for legal professionals using AI tools. The recommendations do not purport to be exhaustive but merely orient legal professionals to ask the right questions and promote future exploration.

A. Regulatory Solutions

It is essential for the legal profession to take a proactive approach to the regulation of AI to ensure that it is used in a manner consistent with professional standards and that it does not place clients at risk of harm. Regulators and professional bodies like the American Bar Association have already taken steps to shape best practices for legal professionals interacting with generative AI. The European Union's release of its proposal for a law on AI is the first by a significant regulatory body. In the United States, Congress has held hearings with AI leaders, like Sam Altman, CEO of OpenAI, and will soon follow the EU with its own set of guidelines.

The Model Rules of Professional Conduct, created by the ABA, are guidelines that set the standards of professionalism required of lawyers. They serve as a model for individual states, which regulate the conduct of lawyers practicing within their jurisdictions. Although states can adopt or modify the model rules, they must follow the standards set by the ABA. These model rules should be considered in the context of legal professionals using generative AI.

[14] OpenAI, "Terms of Use" (Mar. 14, 2023).

[15] OpenAI, "March 20 ChatGPT Outage: Here's What Happened" (Mar. 24, 2023).

[16] Mahdi Assan, "Notes on LLMs and Privacy Leakage," The Cyber Solicitor, Mar. 10, 2023.

Model Rules 1.1 and 1.6 are particularly relevant. Model Rule 1.1 states that a lawyer should provide competent representation to a client, which requires the necessary legal knowledge, skill, thoroughness, and preparation for the representation.[17] The commentary on this rule indicates that lawyers should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology.[18] The rising use of AI software in daily legal practice could be subject to many ethical risks, as described earlier. Therefore, lawyers are expected to cultivate a general understanding of those technologies, which will enable them to consult with experts when designing, adopting, and using new AI software applications in their practice.

Model Rule 1.6 primarily focuses on the preservation of client confidentiality. It states that a lawyer should not reveal information concerning the representation of a client unless the client gives informed consent. This is a fundamental principle of the attorney-client relationship. Also, a lawyer should make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information concerning the representation of a client.[19] Any breaches could result in sanctions, disbarments, and loss of trust. Consequently, with the rise of AI technologies, law firms face the challenge of upholding this confidentiality in an increasingly digital environment.

In 2019 the ABA passed Resolution 112 addressing AI, which was more proactive than reactive.[20] The ABA urged courts and lawyers to grapple with the "emerging ethical and legal issues related to the usage of artificial intelligence ('AI') in the practice of law." The ABA wanted to get ahead of the developing technology and ensure that the law was ready to protect against various forms of abuse and misuse of the technology.

With the rapid and extensive proliferation of ChatGPT, the ABA's concern is no longer one about the future. AI is a present reality, and the ABA has reacted accordingly by adopting Resolution 604. Resolution 604 urges all those creating and using AI to follow three main guidelines:

> 1. Developers, integrators, suppliers, and operators ("Developers") of AI systems and capabilities should ensure that their products, services, systems, and capabilities are subject to human authority, oversight, and control;
>
> 2. Responsible individuals and organizations should be accountable for the consequences caused by their use of AI products, services, systems, and capabilities, including any legally cognizable injury or harm caused by their actions or use of AI systems or capabilities, unless they have taken reasonable measures to mitigate against that harm or injury; and
>
> 3. Developers should ensure the transparency and traceability of their AI products, services, systems, and capabilities, while protecting associated intellectual property, by documenting key decisions made with regard to the design and risk of the data sets, procedures, and outcomes underlying their AI products, services, systems and capabilities.[21]

For tax professionals, the second point is the most relevant. While most legal professionals will not be responsible for developing AI systems, they will be using them to serve their clients.

In addition to potential legislation, federal AI regulation may be in the pipeline at the Federal Trade Commission. In 2022 the FTC issued an advance notice of proposed rulemaking to address commercial surveillance and data security, which included a discussion on automated decision-making systems.[22] Public comments were invited regarding the kind of

[17] ABA, "Model Rule 1.1: Competence" (2023).

[18] Rafael Baca, "Model Ethics Rules as Applied to Artificial Intelligence," Law Practice Today, Aug. 14, 2020.

[19] ABA, "Model Rule 1.6: Confidentiality of Information" (2023).

[20] ABA, "Resolution 112" (Aug. 12-13, 2019).

[21] ABA, "Resolution 604" (Feb. 6, 2023).

[22] FTC, "Commercial Surveillance and Data Security Rulemaking" (Aug. 11, 2022).

transparency companies should provide to consumers regarding AI, prohibitions on unfair or deceptive AI uses, and certification of AI as meeting accuracy, validity, and reliability standards.

The EU's Artificial Intelligence Act proposes different rules for AI systems depending on the level of risk they pose.[23] Under the legislation, AI systems that are deemed to present an unacceptable risk, such as those that could manipulate people's behavior or classify people based on personal characteristics, are prohibited. High-risk AI systems, like those used in critical infrastructure or law enforcement, are to be assessed before being put on the market and then periodically throughout their lifetime. Limited-risk AI systems must comply with minimal transparency requirements, allowing users to make informed decisions about their use.

From social engineering to surveillance, the EU's proposed guidance addresses the high-level concerns with AI use on a multinational scale. It does not, however, address the concerns that AI and LLMs pose to tax professionals.

B. Technical Solutions

As we navigate the intersection of AI and tax law, our attention turns to how technology designers can craft tools that both enhance efficiency and safeguard against the potential pitfalls outlined above.

Understanding AI tools, in accordance with the ABA's Model Rule 1.1, is vital for law firms and privacy-centric businesses. Generative AI offers significant time and cost savings, yet firms must carefully manage the data input to optimize their potential without compromising privacy. What follows will provide an overview of these essential considerations, giving practical insights for legal professionals using AI tools.

1. Proprietary general-use LLMs like ChatGPT.

Using ChatGPT in your web browser is one of the quickest ways lawyers can start using generative AI for tax research as it allows access to OpenAI's proprietary LLM. OpenAI's models are trained on large amounts of internet data, including Wikipedia and the Internal Revenue Code, and can provide some kinds of answers to tax questions out of the box. However, because ChatGPT is trained on much more than tax law, the answers it provides for complex tax questions are not always reliable, such as in the case of the aforementioned lawyers relying on cases that ChatGPT hallucinated.

That being said, the best way to avoid being led astray by poor information from general LLMs, like ChatGPT, is to use these platforms for purposes they are better equipped to handle. Because OpenAI cannot completely guard against hallucinations in tax and legal research work, using these platforms for more general tasks or to help summarize or revise existing work — instead of using it to generate new tax or legal research — is a safer way to harness the power of these tools and to maintain quality standards.

A major privacy concern with ChatGPT is that under the default settings — when you input a prompt — all your data ingested into the public ChatGPT model becomes part of OpenAI's data repository and could be incorporated into the responses ChatGPT generates for other users. In response to concerns about data leakage, OpenAI has introduced the ability to turn off a user's chat history. This feature allows users to interact with AI systems while ensuring that their personal information remains anonymous and untraceable. Once chat history is turned off, OpenAI claims that it will not use conversation history to train its models and that the data will be deleted after 30 days. Also, OpenAI has advertised that they are working on a business subscription that will give users more control over how their data is used.[24]

To avoid contravening the ABA ethical guidelines, lawyers should still avoid putting confidential information into ChatGPT, even if they have disabled chat history. While the risk of data leaking out through the chatbot is ostensibly lowered, the information is still stored for 30 days and could remain vulnerable to data breaches.

[23] *See, e.g.*, the EU's 2021 Artificial Intelligence Act.

[24] OpenAI, "New Ways to Manage Your Data in ChatGPT" (Apr. 25, 2023).

2. Third-party AI applications.

Third-party AI applications can leverage the proprietary models created by companies like OpenAI while improving on the limitations of the web-based ChatGPT. To bolster the accuracy and verifiability of information generated by AI, tools like Blue J Legal's new generative AI tool, Ask Blue J, employ two key strategies: augmenting training data with additional context and providing citations and links to source materials for the AI's answers so that the user can audit the system's accuracy and reliability.

Ask Blue J is trained on an extensive tax content library, enabling it to provide precise answers to tax-related queries within specific contexts. When a user poses a question, Ask Blue J identifies the most relevant tax content that aligns with the scenario set forth in the query. To obtain an answer, the identified tax content is securely transmitted to a general-purpose LLM via an application programming interface. APIs are a way for different software applications to communicate and share services and data in a structured and secure way.

Ask Blue J also illuminates the AI black box by citing and providing links to resources for its answers.[25] In contrast to ChatGPT, which simply provides an unverifiable response to user queries, Ask Blue J provides the most important documents that it relies on for its answers. Users can open these documents and find — highlighted — the exact section that Ask Blue J is citing for its response. This feature allows users to access these documents, identify sections cited by Ask Blue J, and further explore related resources, enhancing the quality and depth of their research. It allows users to verify the responses, thereby providing confidence in the response and the system.

In accordance with OpenAI's data use policies, third-party applications, like Ask Blue J, that use their own API are opted out of data being used for training. By default, OpenAI will not use data submitted by customers via its API to train or improve its models, unless a user explicitly decides to share their data. Any data sent through the API will be retained for abuse and misuse monitoring purposes for a maximum of 30 days, after which it will be deleted (unless otherwise required by law).[26]

When using third-party applications, OpenAI's policy will satisfy many lawyers and tax professionals, but it is important to note that third-party applications using proprietary models may have opted into data training, or may use your data in other ways. Users must review the terms and conditions of service providers to learn what happens to the data — including client information — entered into the tool. This understanding is essential for assessing whether client information can be used with a tool without raising confidentiality concerns.

Blue J has passed stringent and independent auditing, adhering to the American Institute of CPAs' System and Organization Controls (SOC 2) requirements. This compliance ensures that information remains confidential, accessible, and secure. Also, Blue J's approach to AI does not require PII. PII collection is limited to the data necessary to manage authentication and authorization to use the Blue J platform (email address and name).[27]

3. In-house trained or fine-tuned models.

Conceptually, it is possible for law firms to avoid ever sending data to external proprietary models by developing private in-house AI tools. One setup could include an LLM on in-house private servers on which all data is stored. Generative models have largely been confined to larger tech companies because training them requires massive amounts of data and computing power. But once a generative model is trained, it can be fine-tuned for a particular content domain with much less data.

A potential advantage of training or fine-tuning a model on internal servers is that law firms would not need to anonymize PII. This could offer advantages in terms of time savings and potential customization. For example, law firms could generate customized legal documents that already include the relevant client information.

[25] Alarie et al., *supra* note 3.

[26] OpenAI, "API Data Usage Policies" (June 14, 2023).

[27] Blue J Legal, "Information Security Program at Blue J" (2022).

Another legal risk for in-house AI is the potential to violate privacy laws. If an internal database includes personal information, businesses' training models must comply with relevant privacy laws, such as the General Data Protection Regulation in the EU or the California Consumer Privacy Act. This includes the right for information to be forgotten, which may include training data. Thus, if a firm uses client data in training a model and the client requests that its data be deleted, firms may have obligations under these acts.

The trade-off of in-house data storage is that it requires a large capital investment in development costs and maintenance, and can be biased if training data is not updated. This may prove unattractive to firms. In the future, companies like Blue J may offer more bespoke, in-house tools that are fine-tuned with a firm's data, which would be the perfect middle ground between LLMs built in house and publicly available models, like ChatGPT. These tools would offer greater PII protection, and their responses would be specified to the firm. They would be trained with the firm's own secondary sources as well as file memo templates stripped of all client information.

IV. Mitigating Risk

There are several ways that tax professionals can mitigate some of the risks involved in using AI tools. Perhaps the most important responsibility is to verify their outputs. Lawyers and accountants must be professionally diligent and review the material that these platforms provide. Moreover, whether professional services firms choose to limit or allow the use of generative AI, robust security measures, clear policies, and a culture of awareness can also help alleviate risk. To encourage the ethical use of AI tools and maintain data privacy, firms can adopt the following strategies:

- Third-party protection: Firms should verify the policies of third-party AI applications. The policies implemented for in-house use should also extend to contractors and other third parties that may use generative AI.
- Employee education and awareness: Keeping employees informed about the risks associated with generative AI is crucial. This could include updating employee handbooks, agreements, and policies to address the use of generative AI, and conducting training programs. Incorporating AI tools into the workflow is not just about technology acquisition — it's about organizational transformation. Therefore, law firms must have a comprehensive plan to train employees about company expectations, boundaries for appropriate use, and the identification of potential violations. Maintaining a robust policy on the use of generative AI and ensuring communication and adherence at all levels of the organization are crucial.
- Client consent: With increased digitization in law firms and the potential use of AI in legal research, obtaining informed consent from clients becomes critical. Discussing the processes and implications of using AI tools, especially regarding data privacy, should be a regular part of a firm's engagement with clients.
- Anonymization: Anonymize text so that sensitive information never reaches OpenAI or other service providers in the first place. For example, client names and other PII can automatically be replaced with placeholders.[28]

Tax professionals should also consider some practical self-limitations, whether or not formalized in corporate policies. For example, it would be prudent to forgo mentioning the company name or other company-specific or identifying information or any nonpublic or proprietary information in chats with generative AI.

To address privacy concerns regarding web-based AI tools, there are third-party solutions available. Companies like PrivateGPT have developed data privacy tools designed to safeguard privacy for organizations. These tools offer an additional layer of protection by identifying, removing, and replacing PII from ChatGPT requests. Importantly, these tools ensure

[28] Anthony M. Insogna et al., "Trade Secrets and Generative AI: Protective Measures in an Evolving Technological Landscape," Jones Day Insights, June 2023.

that the data provided by the organization is not used in training the AI model, thereby mitigating privacy risks.[29] By leveraging those solutions, companies can enhance data privacy and maintain confidentiality while still benefiting from the functionality of AI tools, like ChatGPT.

V. Conclusion

As generative AI systems, including tools based on LLMs like GPT-4, become increasingly integrated into the legal profession, it brings both opportunities and challenges. While AI tools can enhance efficiency and assist legal professionals in tasks like legal drafting and tax research, they also raise important ethical concerns and data privacy risks.

The challenges associated with generative AI in the legal field include concerns about the quality and accuracy of the output, the issue of biased answers, the lack of "explainability" or verifiability, and questions of liability when errors are not caught by the responsible professionals and hurt clients. These challenges necessitate careful training, ongoing monitoring, and validation of AI-generated information, as well as the development of safeguards to prevent damaging mistakes or data breaches and to ensure client and user privacy.

To address these concerns, regulatory, technological, and professional solutions are being developed. Regulators, AI developers, and professionals are taking steps to shape best practices and directives for the ethical use of AI in professional services. As the law and accounting professions navigate the use of generative AI and its ethical implications, it is crucial to strike a balance between embracing the benefits of AI technology and safeguarding client interests and data privacy and maintaining professional responsibilities. By addressing the challenges and implementing appropriate solutions, tax professionals can leverage the power of generative AI to better serve their clients while upholding their ethical duties. ■

[29] Private AI, "PrivateGPT: The Privacy Layer for ChatGPT" (2023).