Proskauer

Agentic Artificial Intelligence: Looking Ahead to Potential Practical and Legal Issues When AI Gets Autonomous

New Media and Technology Law on **May 8, 2024**

Generative AI has been most synonymous in the public mind with “AI” since the commercial breakout of ChatGPT in November 2022. Consumers and businesses have seen the fruits of impressive innovation in various generative models’ ability to create audio, video, images and text, analyze and transform data, perform Q&A chatbot functions, and write code, to name a few functions.

Still, so far AI chatbots have been mostly contained to the confines of a chat window. Although a basic generative AI tool could recommend top-rated hotels for your weekend getaway plans, it will (as of the date of this article) lack the capability to secure a reservation or otherwise transact with the outside world on your behalf. With the advent of agentic AI, that might be changing.

This is Part I of a two-part article on agentic AI. In this part, we will give an overview of agentic AI technology, discuss some basic development concerns with such systems, and then close with a discussion about how the concept of reliability may evolve in terms of contracting practices between customers and developers/hardware makers.

In Part II, we will dive into the legal issues with respect to user liability for agentic AI-assisted transactions and outline the open questions about how the principles of agency law may apply in these AI-assisted transactions (and how the law may need to evolve with this new technology).

Overview: What is Agentic AI?

So-called large action models (“LAMs”) are AI systems that understand digital interfaces typically designed for humans and learn to execute human actions autonomously (or with limited human supervision) within these digital environments. Indeed, this type of technology could solve one of the problems that has plagued robotic process automation: The need to reprogram the robot every time changes are made to an interface with which it interacts. More flexible Agentic AI systems could enable increased automation and worker productivity in certain types of industries and assist those who lack digital literacy.

In the consumer and ecommerce space, an AI agent might be able to interact with apps or websites, add items to a shopping cart and check out in accordance with pre-registered preferences and payment options, fill out and submit a form, or RSVP to an event. Of course, the basic idea for such an automated process is not new (the Amazon Alexa device, for example, offers users voice-activated, automated ecommerce), but agentic AI systems are different in that they are not limited to a single website or interface. Instead, LAMs are trained to read and understand interfaces generally, giving them the ability to autonomously finalize transactions on a range of apps and platforms.

Quietly – amid the flurry of attention around generative AI – the first wave of AI gadgets has arrived. For consumers, the promise of such gadgets is a mixture of expedience and technological wizardry – the ability to navigate apps or complete basic transactions without the need to unlock a phone, stare at a screen or open multiple apps. For example, the Rabbit R1 device, which developer rabbit inc. states is “enabled by recent advances in neuro-symbolic programming,” is designed to be able to navigate apps and perform tasks for the user; the recently released Humane Ai Pin attaches to a user’s shirt and acts as an AI-powered digital assistant, responding to touch and voice and featuring a laser projection on a user’s palm; and various smartphones and other tech gear now feature an AI assistant. Of course, we won’t be able to fully gauge the functionality and effectiveness of these devices and features until they become more widely adopted (and go through several versions and updates).

Still, it remains to be seen whether human users would feel comfortable trusting 100% autonomous transactions. From botched flight reservations to unintentional spending sprees, it is not hard to see how glitches and bugs in LAMs could have serious real-world consequences (not unlike malfunctioning autonomous vehicles, which of course have their own form of agency).

With agentic AI technology mostly still in development, it is impossible to completely predict the myriad of legal issues that will follow, but our two-part article will examine a few.

Development Concerns

In December 2023, OpenAI released "Practices for Governing Agentic AI Systems," an outline of some general parameters and open questions around the responsible and safe development of agentic AI systems.

The paper suggested certain measures for testing agentic AI effectiveness, including:

- **Goal complexity**: How difficult would a system's goals be for a human to achieve reliably, quickly and safely? How can developers, system deployers and users effectively evaluate the agentic system's level of reliability?
- **Environmental complexity**: How complex are the environments in which a system can achieve its goal? Can a system be more "legible" to the user or system deployer (e.g., presenting a "ledger of actions taken by the agent" to assist review)?
- **Adaptability**: How well can the system adapt and react to change or unexpected events? What real-world failures are beyond reasonable expectation or training?
- **Independent execution**: To what extent can the system achieve reliable results with no or limited human supervision?

To both individual and commercial users of agentic AI, the most pressing concern is reliability – "Will this system work properly and without constant oversight?" As noted in the OpenAI paper, the technology is still in its infancy and, therefore, it can be difficult for system developers to evaluate an agentic model's success across a range of conditions and use cases. Even if an agentic AI system is able to perform the various individual tasks, the paper also states that it can be difficult to evaluate whether it can "chain" these actions together in a real-world environment where change can be hard to identify or anticipate and where one error in a single subtask could cascade into failure.

Contracting Issues

Thinking ahead, one might ask how questions of reliability will be handled in relevant terms of use or other license agreements with developers. Possible emerging trends include:

Warranty Disclaimers. Given the likely reliability limitations, agentic AI developers will probably assert broad reliability disclaimers for their tools. For example, the Human Ai Pin's terms of use state the developer is "NOT RESPONSIBLE FOR ANY DAMAGE THAT MAY RESULT FROM THE PRODUCTS, THE SERVICE, OR YOUR USE OF THE PRODUCTS OR SERVICE." Disclaimers like this might be analogized to generative AI disclaimers that model output can be nonsensical or false (e.g., because of "hallucinations") – though in this case, unlike with most generative AI, there may be no chance for human review to mitigate this risk.

Indemnification. Unlike autonomous vehicles, which have multiple cameras and sensors that automatically respond to road conditions, emerging agentic AI devices generally rely on user prompts – voice, touch or motion – to act. As such, agentic AI developers may not only disclaim responsibility for their products' actions, but even seek indemnification against claims arising out of their use. Again, the Humane Ai Pin's terms are instructive: The user must indemnify the developer against any claim related to "use of the [pin] by [the user] or under [the user's] account." This sort of risk-shifting, together with the reliability issue, may make first generation agentic gadgets too risky for many users.

To be sure, it would be a game-changer if a vendor was able to offer some sort of reliability commitment around an agentic AI system. However, the current generative AI licensing market provides a helpful precedent as to why this likely will not happen in the technology's early days. In the past year, major generative AI providers have added certain customer-friendly provisions to their commercial terms, such as rolling out intellectual property protection for generative AI outputs, but have generally avoided commitments regarding reliability of generated outputs. Thus, given that agentic AI is still a developing technology, it is similarly unlikely that reliability warranties will become part of an agentic AI developer's commitments in the near term.

In Part II of this article, we move from considering agentic AI legal issues between developers and users, to legal issues involving the websites, apps, and other third parties and platforms with whom the AI agents might interact. Stay tuned!

View original.

Related Professionals

- **Peter J. Cramer**
  Associate
- **Caroline E. Rimmer**
  Associate