Westlaw Today
powered by Reuters

What is generative AI and how does it work?

By Resa K. Schlossberg, Esq., Jenness E. Parker, Esq., Pramode Chiruvolu, Esq., and Matthew P. Majarian, Esq., Skadden, Arps, Slate, Meagher & Flom LLP*

JULY 13, 2023

Key points

- Generative AI systems have already found widespread application in the business world and their capacity to disrupt a broad range of industries is now apparent.
- The technology has enormous potential, but comes with many risks, including the potential for copyright and privacy infringements, contractual violations, the disclosure of trade secrets and untrustworthy outputs.
- Directors weighing opportunities for a company to use, build or contribute content for these AI platforms need to understand in broad terms what generative AI systems are, as well as the legal issues they pose and what steps companies can take to mitigate the risks.

What is generative AI and how does it work?

OpenAI's ChatGPT platform is reportedly the fastest-growing consumer application in history. It uses generative AI models, which produce new content based on training on vast quantities of data. The system underlying ChatGPT reportedly trained for months on hundreds of billions of words pulled from the Internet.

Boards should be aware of the risks associated with models trained on data from the Internet, including the intellectual property, privacy and contractual risks of relying on the outputs of such models.

Through this process, OpenAI's systems and similar "large language models" have achieved near-human level abilities to answer questions, write poetry, compose essays and even perform in the 90th-99th percentile across a wide range of college, graduate and post-graduate exams.

But boards should be aware of the risks associated with models trained on data from the Internet, including the intellectual property, privacy and contractual risks of relying on the outputs of such models.

How generative AI is being used today

ChatGPT and other text-generating models are far from the only uses of generative AI.

Companies across different industries are using the technology:

- *Financial services firms* are leveraging generative AI to streamline backend operations, bolster cybersecurity, support service chatbots, accelerate software development, enhance fraud detection and provide personalized financial advice.
- *Entertainment companies* are using text-to-image generators to create art for storyboards and visual content, including special effects, for films and video games.
- *Pharmaceutical researchers* are using generative AI to better understand the structure of proteins and design them specifically for medicines. For example, Canadian researchers trained an AI system on images of known proteins to generate new proteins with specific difficult-to-replicate protein folding.
- *The materials science industry* uses the technology to compose new materials with the desired physical properties.
- *In healthcare*, AI is being applied to electronic health records systems, and generative AI systems are also being used to produce synthetic data (*i.e.*, fictitious data that mimics real-world data without personally identifiable information) to allow data sharing and analysis otherwise restricted by privacy laws.

Boards will need to understand the risks these innovations bring as this technology expands across all sectors.

Frequently asked questions about the use of generative AI

1. What are the risks when inputting information into third-party generative AI platforms?

Generative AI platforms' terms of use often permit them to use inputs to improve their models and monitor system usage (including for compliance purposes), and some terms of use grant even broader rights to AI platforms to use and sublicense any inputs for any purpose.

But if the information input is owned by a third party, you may breach confidentiality obligations to them. Furthermore, it may be hard to anticipate the impact of supplying information for the model.

Thomson Reuters is a commercial publisher of content that is general and educational in nature, may not reflect all recent legal developments and may not apply to the specific facts and circumstances of individual transactions and cases. Users should consult with qualified legal counsel before acting on any information published by Thomson Reuters online or in print. Thomson Reuters, its affiliates and their editorial staff are not a law firm, do not represent or advise clients in any matter and are not bound by the professional responsibilities and duties of a legal practitioner. Nothing in this publication should be construed as legal advice or creating an attorney-client relationship. The views expressed in this publication by any contributor are not necessarily those of the publisher.

THOMSON REUTERS®

The platform's use of your information to improve its model could result in that information being incorporated into a training dataset published by the platform provider, or used to train a model that ultimately discloses your information in response to another user's prompt.

For instance, if your employees ask a generative AI system to debug confidential software source code, that source code could be used to train an improved model that releases the code in some form to subsequent users.

Moreover, AI systems typically store information on an external server. If the security of that server is breached, the user's information could be disclosed publicly, which potentially could be devastating to the user, the owner of the information or both.

2. Who owns the output (or results) of generative AI systems?

While the terms of use of generative AI platforms typically grant ownership of outputs to the user, whether the technology can generate *any* protectable intellectual property rights remains unsettled. Citing established U.S. law that human authorship is required for a work to be copyrighted, the U.S. Copyright Office recently canceled the copyright registration of an AI-produced graphic novel. It reasoned that, because of the unpredictable nature of the image generation, the human author could not be considered the "mastermind" of the work.

In March 2023, the Copyright Office published guidance stating that that registrability of works including AI-created content depends on factors such as how the AI tool operates and how it is used to create the final work. Complex written, visual or musical works generated from simple human prompts by an AI system are not registrable, it said. In copyright registration forms, applicants now must disclose any AI-generated material in a work and explain the human author's contribution.

Regardless of whether generative AI outputs qualify for copyright protection, companies also need to consider whether their use may infringe third-party intellectual property rights. In a suit by Getty Images against Stability AI (developer of the text-to-image platform Stable Diffusion), for example, Getty claims that the output of the defendant's image generation platform often contains a modified version of a Getty Images watermark, creating confusion as to the source of the images.

And most publicly available terms of use of generative AI systems expressly disclaim liability for third-party copyright infringement, leaving end users to take the risk that outputs they might incorporate into their products or publications are infringing.

Companies will therefore need to document the use of AI- and non-AI-generated content to ensure that their products can be copyrighted, and be alert to the possibility that the output of the models could infringe on the rights of others.

3. How trustworthy are the results of generative AI?

Today's technologies are far from perfect. Because these systems are trained to generate responses that appear similar to the training data based on probabilities, they are prone to "hallucinations," where the system generates inaccurate content and presents it as fact — often convincingly to non-experts.

Such inaccuracies could impact business outcomes or create liability issues if, say, false information is communicated to the public. That could result in reputational or operational damage, and even defamation claims.

AI text-to-image generators allow users to create amazingly realistic but fictitious images, such as the "deep fake" photos of Pope Francis wearing a white puffer jacket that went viral. The possibility of misuse of generative AI to spread disinformation and misinformation is a concern not just to organizations, but to society as a whole.

Bias also continues to be an important concern because generative AI systems may perpetuate or amplify biases in the training data or in the algorithms or user prompts. This is particularly concerning if the technology perpetuates or amplifies biases based on legally protected characteristics such as race, gender or sexual orientation.

Regulators are already focused on this issue. The EU and U.K.'s General Data Protection Regulation and the California Consumer Privacy Act both give individuals the right to opt out of "automated decision-making" where AI may be used to build profiles and make individual decisions, such as extending employment or product offers.

4. Can companies face liability for training generative AI systems?

State-of-the-art generative AI is trained on vast amounts of data (including text or images). Boards considering deploying generative AI should understand if the training process violates copyrights, privacy requirements and/or contractual restrictions.

Copyright. Absent an express license, training generative AI may violate copyrights in the works included in the training data. Copyright owners have already filed infringement suits against generative AI providers in the U.S. and U.K.

There may be a "fair use" defense for a limited and "transformative" purpose, such as commenting on, criticizing or parodying a work. But fair use turns on the facts of a specific case and U.S. courts have not addressed the defense in the AI context, and other jurisdictions could come to different conclusions.

Privacy. An AI model trained on sensitive or personal information might unintentionally generate outputs that reflect this information. Even if the details are not explicitly in the training data, the systems might learn associations between individuals and sensitive attributes like race, gender or health status, potentially leading to privacy breaches. Individuals also may not be aware that their personal information is being used to train these systems, and they may not have given consent or been given an opportunity to opt out.

These types of privacy concerns recently led Italy's data protection authority to briefly halt the use of ChatGPT in Italy while OpenAI responded to inquiries regarding privacy risks.

And in several cases where personal data used to train AI models was gathered or used in violation of privacy policies, the Federal Trade Commission has required "algorithmic disgorgement" — the permanent deletion of all models improperly trained on the

personal data. After years spent developing AI training datasets and training models, companies facing algorithmic disgorgement need to start from scratch.

Breach of contract. Where content is gathered by crawling or scraping websites, website owners might contend that their websites' terms of use were violated.

What should a board do?

Careful oversight remains critical. For a board, that entails taking reasonable measures to implement and oversee risk management and compliance controls.

Although courts in Delaware, whose law governs most large companies, have yet to weigh in on AI issues, some guidance can be found in a case involving a data breach that exposed customers' personal information. Stockholders alleged that directors violated their duty of oversight.

Although the breach stemmed in part from significant lapses (including the use of a simple generic password to secure critical data), the Delaware Court of Chancery ruled for the directors. It noted that the board had charged two committees with monitoring the company's data security processes, that those committees were well-functioning and met regularly, and that the committees set up appropriate reporting structures.

Boards should carefully consider how best to oversee a company's use of generative AI.

That may entail:

- Establishing monitoring and compliance systems and paying ongoing attention to them, perhaps through a committee empowered to evaluate technology-related risks.
- Paying particular attention to "mission critical" issues involving the use of generative AI.
- Discussing with advisers issues on which the board should receive regular reports and identifying what "red flags" (*i.e.*, indications of potential operational deficiencies) may arise and how best to respond.
- Documenting in board minutes and materials the monitoring system reports to the board, and both directors' and officers' oversight efforts, so the company can respond to books-and-records demands and defend itself.
- Evaluating how generative AI may be used to enhance a company's oversight systems and processes, for example, by automating reports or by creating monitoring or analysis tools to spot potential deficiencies.

Jurisdictions including the U.S., U.K., European Union and China are grappling with the question of whether and how to regulate the technology. Boards will also need to stay abreast of those developments.

About the authors




(L-R) **Resa K. Schlossberg** is a partner in **Skadden, Arps, Slate, Meagher & Flom LLP**'s New York office and head of the firm's intellectual property and technology transactions practice. She advises clients in numerous industries on a wide range of IP-related issues, focusing on life sciences and technology-intensive representations. Her research experience and knowledge of molecular biophysics and biochemistry help her understand the scientific and technical aspects of cases. **Jenness E. Parker** is a partner in the firm's Wilmington, Delaware, office who litigates corporate disputes in trial and appellate courts. She represents public and private companies, their directors and advisers in derivative and class actions, corporate governance disputes, M&A-related actions, books and records matters, appraisal proceedings, advancement and indemnification actions, other Delaware statutory matters, federal securities laws, and complex contractual disputes. **Pramode Chiruvolu** is a counsel in the firm's Palo Alto, California, office, where he represents clients in licensing, services, manufacturing, logistics, restructuring, financing and other IP, technology and commercial transactions. He also advises on IP strategy, privacy and cybersecurity issues and on matters involving emerging technologies, including artificial intelligence, digital health and biotechnology, the internet of things and 5G networks. **Matthew P. Majarian** is an associate in the firm's Wilmington office. He advises and represents individual, corporate and alternative-entity clients on a broad range of matters, including fiduciary duty litigation, governance issues, books and records actions, complex commercial disputes, and transactional structuring and related litigation. This article was originally published May 24, 2023, on the firm's website. Republished with permission.

This article was published on Westlaw Today on July 13, 2023.

*** © 2023 Resa K. Schlossberg, Esq., Jenness E. Parker, Esq., Pramode Chiruvolu, Esq., and Matthew P. Majarian, Esq., Skadden, Arps, Slate, Meagher & Flom LLP**

This publication was created to provide you with accurate and authoritative information concerning the subject matter covered, however it may not necessarily have been prepared by persons licensed to practice law in a particular jurisdiction. The publisher is not engaged in rendering legal or other professional advice, and this publication is not a substitute for the advice of an attorney. If you require legal or other expert advice, you should seek the services of a competent attorney or other professional. For subscription information, please visit legalsolutions.thomsonreuters.com.