Legal 500
Country Comparative Guides
Hot Topic | Artificial Intelligence

Artificial Intelligence

Contributor

Debevoise & Plimpton

Avi Gesser

Partner | agesser@debevoise.com

Matt Kelly

Counsel | makelly@debevoise.com

Jarrett Lewis

Associate | jxlewis@debevoise.com

Melissa Muse

Associate | mmuse@debevoise.com

Artificial Intelligence

The pace of Artificial Intelligence ("AI") adoption by companies in the United States has increased rapidly over the last year. This is attributable, in large part, to dramatic leaps in the capabilities of certain text-based Generative AI ("GAI") tools, with the most well-known example being OpenAI's ChatGPT.

The increasing use of AI – and its rapid integration into our personal and professional lives – raises a wide range of important challenges and questions for sectoral regulators, policymakers, and the public. At the same time, companies deploying AI face an array of risks under existing and upcoming legal regimes.

In this Chapter, we discuss four common use cases for AI today and present some of the most significant legal risks associated with those use cases, along with ways that companies may mitigate those risks.

Coding and Software Development

Perhaps the most common GAI use case for businesses is coding and software development. This includes using GAI to create new code, translate existing code (as well as code snippets between programming languages), and check existing code for bugs or unexpected behavior. Much of the software development process already involves borrowing pieces of code from existing libraries and repositories. Therefore, developers have been able to easily integrate GAI tools – particularly those that supply or recall snippets of code for particular functions or purposes – into existing coding workflows. Software developers frequently report that using GAI tools such as ChatGPT or GitHub Co-Pilot makes them more efficient and productive.

Risks

Quality Control. Code may be judged objectively by whether it achieves a particular function without errors or without undesirable or unsustainable levels of resource consumption. It also may be judged on more subjective criteria, including elegance and relative efficiency.

AI-generated or AI-evaluated code raises risks in both respects. For instance, because consumer-facing GAI tools are trained based on large databases of written content (which includes un-curated collections of code and code snippets), they are capable of replicating significant coding errors that may have existed in the training data or creating new coding errors because the training data is ill-suited for the use case. As a result, code produced by GAI tools can be inefficient or inappropriate as a means to address a particular problem. Lack of human involvement in the code generation also makes the output more difficult to review and validate. Subjective quality deficiencies can therefore contribute to diminished functionality in the near- or long-term.

These performance and operational risks can become legal risks to the extent they impact essential business functions or a firm's ability to meet service-level commitments to customers or contractual counterparties. This is particularly true for firms in highly regulated industries, where liability can attach based on performance and quality-control shortcomings.

Cybersecurity. Code snippets provided by GAI may replicate security vulnerabilities that appear in the training data or may create new vulnerabilities due to the particular weaknesses of AI-generated code for novel contexts or network environments. These vulnerabilities may be overlooked or not addressed due to a lack of direct human oversight or over-reliance on the performance of the AI model – particularly to the extent a firm's adoption of GAI has led to reductions in process, experience, or headcount among its software engineers.

If this kind of over-reliance on GAI leads to unanticipated cyber-vulnerabilities or exposure of sensitive data or personal information, it can lead to legal liability based on expected standards of care in the relevant industry or jurisdiction.

Intellectual Property. Some of the most important legal risks related to coding and software development stem from the evolving landscape of U.S. intellectual property law. These include:

- Open Source Licensing. Many developers use open-source software for code development, including AI-related code. Developers also borrow snippets from repositories of open-source code to integrate into their own creations. However, certain open-source software solutions and open-source repositories come with onerous license terms that might “infect" code developed with the assistance of the open-source software and materials. Use of open-source solutions for code development, therefore, can carry legal risks based on those license terms. These risks can range from being sued for failing to attribute the open source to the developer to having your proprietary code base subject to publication in the event of the most onerous “copyleft" licenses.
- Copyright. While the United States Copyright Office has noted that technological tools can be part of the creative process, it remains unclear how courts will evaluate whether an AI-assisted creation, including code, is copyrightable. If employees use GAI to generate significant amounts of code (or other content), and to the extent a company cannot prove a sufficient level of human contribution, GAI-produced content may not be protected by copyright.There is also a risk that the GAI, and any content it produces, may be deemed a derivative work of copyrighted materials used to train the GAI models themselves. If that view prevails, content generated by GAI may be deemed an infringing work – particularly if the resulting content looks substantially similar to copyrighted training data or contains watermarks or other trademarks of that data.
- Trade Secrets. A critical element of trade secret law, which tends to be a common recourse for software developers, is that the owner of a trade secret takes reasonable measures to ensure the secrecy of the protected information.If an employee uses a GAI tool for review or debugging and does so without taking adequate measures to protect the code provided to the tool as an input from being added to the model, there is a risk that other users may gain access to the same data, compromising its confidentiality, and potentially supporting an argument that the company (and its employee) failed to take reasonable steps to preserve the code's confidential status.

Mitigating Risks

Pre-Production Quality Assurance Testing. Consider how best to maintain or enhance existing risk-based

quality assurance testing for code developed with GAI's assistance. For particularly high-risk codebases (including those that support products or solutions that will be web-accessible), consider specific measures to support detection and remediation of vulnerabilities. Finally, consider how best to ensure that records of the measures taken are maintained.

Iterative Development and Quality Control. Consider running AI-developed solutions in a sandbox or alongside tools developed via more established means for some period before full deployment. This can help ensure that code developed with GAI is working as anticipated and can also provide opportunities to contextualize or fine-tune GAI models for particular use cases or enterprise IT environments. To help ensure business continuity, consider also maintaining earlier versions of models or other software solutions that can be quickly operationalized when needed.

Documenting Human Contribution. GAI will rarely be used to produce an entire, fully functional codebase for a particular solution or software. Extensive human contribution is required to make use of the GAI-generated code. To later support potential arguments in favor of copyright protection, consider how best to document the extent of human contribution when using GAI to develop code.

Enhanced Risk Acceptance Processes for Open-Source Solutions. Until questions regarding enforceability of open-source license conditions are sufficiently resolved, consider implementing enhanced procedures for risk review and acceptance of use cases involving open-source software or code. Particularly for high-risk or business-critical uses, it may be preferable to avoid open-source uncertainty altogether and adopt solutions subject to commercial licenses or strictly permissive open source licenses.

In any event, ensure that all developers are aware of the risk of "copyleft" open source licenses and only utilize open source software and code that may be subject to such licenses in ways that do not risk disclosure of proprietary code.

Customer Voice Analytics

Customer-facing businesses are increasingly using AI to analyze voice data on customer calls for various purposes, including identity verification, targeted marketing, fraud detection, and improved customer service.

For example, AI voice analytics can detect whether a customer is upset and should be connected with an experienced customer service representative or whether the person on the phone is not really the person they purport to be. These tools also can be used to assist customer service representatives in de-escalating calls with upset customers by making real-time suggestions of phrases to use that only the customer-service representative can hear, as well as to evaluate the employee's performance in dealing with a difficult customer (e.g., did the employee raise her voice, did she manage to get the customer to stop raising his voice, etc.).

The use of AI for voice analytics raises a number of legal risks arising not only from the sensitive content that may be communicated via AI-monitored calls, but also from the sensitive personal information that can be gleaned from recordings of a customer's speaking voice.

Risks

Consumer Privacy Laws. Depending on the data collected, the use case, and the location of the company and its customers, certain privacy laws may require companies to disclose to customers the purposes for which they are using their data, which may include voice data and audio recordings. Some uses of AI voice analytics may not be covered by existing privacy notices. Also, depending on the data collected and the use case, consumer privacy laws may require consent from the customers, mandate that customers be provided with opt-out rights or the right to limit the use of that data, and/or limit the ability of the company to share personal data with third parties, which may include a vendor providing the voice analytic services. Data subjects in these jurisdictions also may have rights with respect to this data – such as the right to access and deletion.

Biometric Privacy Laws. Some uses of voice analytics involve matching customer voices to a particular person so that they can be re-identified the next time they call. This can be helpful for marketing, authentication, and fraud detection. But because these use cases often involve identification of an individual based on their voiceprint, they may trigger various requirements and potential liability under biometric data laws, such as Illinois' Biometric Information Privacy Act ("BIPA"). These laws often require that customers be notified that their voiceprints are being collected and provide their express consent for such collection. And penalties for non-compliance can be steep. BIPA, for instance, provides a private right of action for violations, with statutory damages ranging from $1,000-$5,000 per violation.

Anti-Discrimination Laws. To the extent that uses of AI voice analytics impact important decisions about customers or employees based on their tones of voice (e.g., to assess how effective a customer-service representative is at calming down an upset customer), companies should consider whether these tools have been tested for disparate and discriminatory impacts based on protected characteristics (including race, gender, age, ethnicity, and disability). If the tools have not been appropriately and accurately trained using a wide variety of speech types, there is risk of legal liability if it can be shown that one or more protected classes will be treated worse than other groups based on speech patterns or other cultural attributes.

Employee Monitoring Statutes. To the extent that voice analytics are used for employee training or to assess employee performance, they may fall under one or more of the employee monitoring statutes, which require companies in certain jurisdictions to notify their employees when they are monitoring their activity.

Data Retention. Several privacy and cybersecurity laws, such as the New York Shield Act, GDPR, and BIPA, specify that personal information (potentially including voice recordings and biometric identifiers) should be stored no longer than necessary to achieve a legitimate business purpose or to meet some other legal or regulatory requirement. Some privacy laws also require companies to communicate the retention period to data subjects.

Mitigating Risks

Reviewing Laws for Compliance. Knowing the jurisdictions in which voice analytics are used, and which laws apply, is important for reducing potential liability. This includes knowing when potential carveouts may apply, such as those related to GLBA-covered entities or GLBA-covered data.

Updating Privacy Documents. Legal and reputational risks associated with voice analytics often can be

reduced though enhanced notices to customers. Companies should review their privacy notices and policies and consider updating them to the extent that voice analytics uses are not fully disclosed.

Obtaining Consent for High-Risk Uses. The level of consent needed to collect and use voice data for certain high-risk purposes is currently being tested by plaintiffs and regulators and expanded by new legislation and legislative proposals in the U.S. If a company is currently deploying voice analytics for high-risk purposes, it should consider whether to seek customer consent until the relevant issues are resolved through court decisions, new laws, or clear regulatory guidance. However, note that in several international jurisdictions and in some U.S. states, there may be no choice but to obtain consent.

Avoiding Creating Voiceprints. Many benefits of voice analytics (e.g., assessing if a customer is very upset) can be obtained without creating a voiceprint that allows for identification of an individual based on their voice. Because creating and storing voiceprints triggers additional regulatory obligations under various biometric and privacy laws, consider whether it is strictly necessary for the particular business objective or whether it can be avoided altogether.

Not Sharing Data with Third Parties. Several privacy and cybersecurity laws require additional notice, consent, and security requirements when companies share sensitive customer data, such as voiceprints, with third parties. Therefore, such sharing should be avoided where possible.

Not Using Data for Evaluations. Voice analytics applications are often used to train and coach customer service representatives. These applications take on significant additional regulatory compliance obligations and risks if they are also used in employee evaluations. Consider whether such additional uses are worth the regulatory and reputational risk, especially without meaningful human oversight.

Assess Risks Associated with Ethnic and Racial Differences. If AI voice analytics are used to modify service levels or offerings (e.g., to decide whether a customer should get a certain discount), careful consideration should be given to whether these tools have been tested to make sure that they do not treat people differently based on protected status, including based on race, gender, age, ethnicity, or disability.

Extracting Value from Collections of Client Documents and Other Third-Party Data

AI vendors and consultants are encouraging companies to undertake AI-powered “Big Data” projects designed to extract value of gain insights from existing collections of client documents and other categories of third-party data.

A common example involves collection and analysis of a company's library of vendor agreements to summarize key terms and conditions for strategic or other purposes. For instance, in the event of a data incident, GAI could be used to take a large volume of existing contracts and quickly identify how many have breach notification obligation clauses, what the deadlines for notification may be under those clauses, and whether there are any additional obligations related to cyber breaches, such as cooperation on notification or indemnities.

While these use cases have obvious appeal, they can raise unique risks based on contractual terms.

Risks

Contractual Restrictions. Client or third-party data that companies want to use for these Big Data projects can be subject to two main kinds of contractual limitations. First, there are often confidentiality restrictions that prevent the company from sharing client documents with a third-party AI consulting company or generative AI provider. Second, there may be limitations placed on permissible uses of the data, which may not include the Big Data project. These restrictions can limit the ability of a company to share data with third parties involved in implementing or developing AI solutions . Contractual use limitations may also constrain a company's ability to utilize that data for testing or training an AI application. Other contractual terms, including those relating to limits on the means-of-performance, supervision, transparency, quality, and the ability to leverage subcontractors, also may create legal risk if not appropriately managed.

Cybersecurity. As part of Big Data AI projects, companies often take large volumes of client data that are in a secure location and aggregate them in a centralized data lake or data warehouse either on premises or in the cloud. Some of these locations may present additional cybersecurity risks, especially if the company is not very familiar with cloud architecture and proper configurations.

In addition, many Big Data projects involve consulting partnerships with third-party AI companies, who may themselves have cybersecurity vulnerabilities that could jeopardize the security, confidentiality, or availability of client or other third-party data.

Privacy Risks. To the extent that the documents used for a Big Data project include personal information about clients, employees, or other individuals, there may be privacy-related limitations on the ability of the company to use the documents for these purposes, as privacy law may limit the company's ability to share that information with third-party AI companies, without first establishing the proper notices or consents.

Mitigating Risks

Proactive Contract Review. Review existing contractual agreements to understand whether there are any restrictions on how the data governed by contracts may be used. For future business arrangements, consider the potential use cases for data you are purchasing or licensing and ensure your contractual terms allow for such use cases, especially when data might be used to develop AI applications.

Cybersecurity and Privacy Management. Information security personnel and privacy professionals should be involved in planning and implementing Big Data and AI projects to ensure that security and privacy concerns are identified and addressed. Consider also using internal or outside penetration testers to ensure that any repository housing data that supports an AI initiative is subject to appropriate risk-based security testing and hardening. Also consider conducting cybersecurity diligence on AI providers to ensure that they can be trusted with sensitive company data or provide them with only non-sensitive sample data, synthetic data, or data that has been pseudonymized, anonymized, or tokenized. Furthermore, consider implementing requirements for employees and vendors to delete confidential company data as soon as possible.

Real-Time Transcription

The most popular remote video meeting solutions now offer the ability to easily generate complete audio

and video recordings of meetings. GAI tools are being integrated with these platforms to allow them not only to transcribe discussions during recorded meetings but also to generate near-real-time meeting minutes, including abbreviated summaries of discussion topics, follow-up task lists, and agendas for future meetings.

Risks

Accuracy. The quality and accuracy of live captioning and transcription of multi-participant conversations will depend on a variety of variables, including background noise, the number of speakers, the volume and clarity of each speaker's voice, and particularized cultural lexicons and dialects. Quality control issues can be compounded to the extent the transcripts are then fed as inputs to a GAI tool that then generates meeting summaries or task lists for participants based on the transcription.

Data Retention. For firms in highly regulated industries, records of conversations at meetings (even abbreviated summaries of conversations) may be deemed business records that are required to be retained under regulatory books and records recordkeeping requirements. Companies faced with a reasonable likelihood of litigation may also need to identify and preserve meeting artifacts that could be relevant to subject-matter at issue in the litigation. In either case, the failure to retain relevant records could subject a company to significant liability for rules violations or possible judicial sanctions.

Production and Admissibility of Meeting Artifacts. Transcripts or auto-generated summaries of video meetings would likely be discoverable and subject to production in connection with regulatory investigations or civil litigation. Although questions will arise regarding the reliability and admissibility of these AI-generated artifacts, it is highly likely that (in at least some cases) they will be admitted into evidence.

Consent. In many jurisdictions in the U.S., it is illegal to record a conversation without notice and consent of all involved parties, particularly if the recording party is an attorney. These same laws would in many cases prevent recording or analysis of meetings without the same required consents. Finally, as mentioned previously, certain privacy laws grant individuals rights to their personal data, such as the right of deletion. If a full transcription of a meeting is taken, it is likely that the transcription and anything derived from it is unusable by a company if the speaker exercises their right of deletion with respect to the transcription.

Mitigating Risks

Policies. Companies should have clear policies as to which meetings can and cannot be recorded, what consents are needed ahead of time from the participants (including how those consents must be recorded), and how meeting artifacts will be retained over what period of time.

Disabling via Administrative Settings. Firms in highly regulated industries should strongly consider whether to disable automatically generated meeting artifacts – particularly meeting summaries and task lists that may not reflect the actual topics discussed or the takeaways assigned in the course of the meeting, as even incorrect records (once created) may need to be retained.

Content Review Procedures. To the extent firms do choose to enable these tools, they should consider

whether to adopt routine review procedures for sensitive meetings or meetings that are likely to become the subject of external or regulatory interest.

Contributors

Avi Gesser
Partner
agesser@debevoise.com

Matt Kelly
Counsel
makelly@debevoise.com

Jarrett Lewis
Associate
jxlewis@debevoise.com

Melissa Muse
Associate
mmuse@debevoise.com