

ALSTON & BIRD



Help! My Business Wants to Start Using ChatGPT!

February 2023

Help! My Business Wants to Start Using ChatGPT!

By Dan Felz, Wim Nauwelaerts, Paul Greaves, and Josh Fox

Corporate legal departments are increasingly receiving requests from business clients to use ChatGPT or similar AI-powered tools in their operations. These requests can be urgent, with business clients demanding enablement from Legal. This advisory briefly details what “generative AI” tools like ChatGPT are. It then provides an overview of key legal considerations, including by looking forward to upcoming AI-specific legislation in the EU and the U.S. Finally, it outlines potential ways for corporate counsel to think about enabling engagement with this new technology.

1. What is “Generative AI”?

ChatGPT is one of a suite of AI-powered technologies that is being dubbed “generative AI.” These are tools that can take a prompt or query from a user (the “input”) and respond to it with a type of “output” that resembles what a human would create. These tools are referred to as “generative” because they do not rely on a database of preformulated answers or responses that they can retrieve to address user input. Instead, they have been trained to “recognize” a user’s input and to “generate” a response entirely on their own.

Some of the more well-known examples of generative AI include:

- [ChatGPT](#). ChatGPT is – in simplified terms – a powerful chatbot. It is a “large language model” powered by a neural network that can (a) receive natural-language input from a user, and (b) provide natural-language output that resembles how a human would respond. ChatGPT is operated by the company [OpenAI](#).
- [Generative AI for Images](#). There are also generative AI tools that autonomously create images. OpenAI operates [DALL-E](#), a tool that creates images from a natural language description. The tool “[Stable Diffusion](#)” is similar, creating images automatically in response to natural-language inputs from users.
- [Generative AI for Coding](#). Generative AI tools also assist with creating computer code. OpenAI’s [Codex](#) uses an AI model to generate computer code in response to user input. Codex also powers Github’s “[Copilot](#)” functionality, which suggests code to programmers in real time in response to the code they are already creating.

2. Common Generative AI Use Cases

Generative AI tools are not restricted to any particular use case. But, requests to corporate legal departments seem to be presently coalescing around several specific use cases:

- [Coding Assistant](#). ChatGPT is [reportedly](#) a serviceable coder. It may not be able to create “ready-to-deploy” code, but it can take a goal (“Write a C++ script that does X”) and generate a workable first draft. Generative AI tools like ChatGPT, Codex, or Copilot can potentially save hours of coding time per job – and thus be valuable to architects, developers, and others who are tasked with shipping products and features on deadlines.
- [Content Creation](#). Generative AI can quickly create a wide variety of content. For example, ChatGPT can write draft copy for Sales, Marketing, or Comms. Similarly, image generators like DALL-E or Stable Diffusion can quickly generate a series of images that could be used to mock up initial versions or features of mobile apps, games, or other visual-heavy products or services.

- Document Drafting. ChatGPT can draft documents upon request. For example, it could draft policies for HR. It also drafts legal documents if requested (even though these come caveated, stating that a lawyer should be consulted). Again, ChatGPT’s output is not necessarily ready-to-use, and its quality has not yet been comprehensively reviewed. But, it could conceivably be viewed as a sort of virtual “draughtsman” able to generate initial drafts for review.
- Customer Support. AI-powered chatbots are already a feature of customer service. Some AI-powered chatbots still rely on a limited set of preformulated answers. ChatGPT does not; it generates answers that do not exist prior to a query, based specifically on what it is asked. Customer support departments are thus evaluating whether ChatGPT can be used to improve operations. Some companies may be considering how to integrate ChatGPT directly into customer-facing interactions. Still, even if ChatGPT isn’t used to generate responses directly to customers, some companies may be considering how it can streamline support operations. For example, ChatGPT could quickly review the history of a customer’s prior interactions with a company – then summarize this in 3-4 bullet points for a customer service rep.

3. Legal Considerations When Using Generative AI.

While business clients may be seeing benefits from using generative AI, corporate counsel tends to focus on legal risks that may arise from permitting enterprise use. The risks of generative AI are still being discovered, so this advisory cannot present an exhaustive, closed-ended list of considerations that may be relevant to counsel. At present, however, reporting has identified a number of relevant considerations. Some of the more salient are:

- (a) Who has the Contract with the Generative AI Provider? Is the Contract Workable? Generative AI providers tend to make their products easy to sign up for and use. For example, to use ChatGPT, one needs only to go to [OpenAI.com](https://openai.com) and create an account – after that ChatGPT can be used for free, subject to [OpenAI’s Terms of Use](#). This ease of sign-up and use of standard terms can blur who has a contract with the generative AI provider. Employees can sign up for generative AI tools, without procurement or purchasing functions knowing about it, bypassing corporate due diligence and contracting processes. Use may be partially for business purposes, partially for personal use. This issue occurs across generative AI providers – the tools are readily available to employees who create accounts outside of corporate contracting procedures.
- (b) Confidentiality. Generative AI is “self-learning” technology. That means that what is input into the AI tool – along with outputs and follow-up responses – can be ingested into the tool’s AI model to improve its operation. Employees’ inputs into generative AI tools may thus become part of the tool itself, and thus start reappearing to other users outside your organization. This improvement-centered model creates confidentiality concerns. Companies should assume that everything that is input into a generative AI tool, and everything that the tool outputs in response, will be available to others outside your company.
 - ChatGPT may serve as an example. [Per OpenAI’s TOUs](#), both the inputs into ChatGPT and the outputs from ChatGPT may flow back into ChatGPT “to improve our models.” Of course, the TOUs also indicate there is an “opt-out” option – organizations can “opt-out of having Content used for [model] improvement” by sending an “organization ID” to OpenAI’s support email address.

However, there appears to be as yet no publicly-available information on the effect and extent of this opt-out, so placing organizational reliance on the opt-out may be premature.

- Confidentiality concerns are already reportedly driving corporate decisions. Amazon [reportedly prohibited its employees](#) from inputting confidential information – including code – into a generative AI tool. Apparently, Amazon was beginning to receive output that resembled Amazon’s internal, confidential code.

(c) Intellectual Property. Generative AI intersects with three axes of intellectual property risks.

- First, it is difficult to discern whether AI-generated output contains or resembles IP that belongs to third parties. The training data for generative AI tools has not been disclosed. It is assumed that a significant portion of the training data was available on the internet. Thus, it could be – as an example – that a generative AI tool’s answer to a prompt contains unlicensed excerpts from a copyrighted work, or an image owned by someone else. [Getty Images has announced](#) it is suing Stability AI (the company behind Stable Diffusion) for allegedly unlawfully copying and using “millions” of Getty-owned images to train its AI – suggesting Getty believes Getty-owned works (or derivatives thereof) may be in Stable Diffusion outputs. If AI output resembles important IP of other companies, it is unlikely they will let your organization use it simply because a generative AI tool happened to ingest it during training.
- Second, if AI tools are used to create code, their output may contain open source software (OSS). OSS can pose a host of IP challenges. A number of OSS licenses require attribution to original authors. Other OSS licenses have more draconian consequences, particularly if OSS code is modified and integrated into other code; these OSS licenses can require publication of the new and modified open-source code – with the proprietary additions – to the world.
- Lastly, it is not clear what IP rights users of generative AI tools can claim over creations of AI-powered tools. Per [OpenAI’s TOUs](#), as between companies that use ChatGPT and OpenAI, OpenAI “assigns to you all right, title, and interest” to ChatGPT-generated output. This is designed to settle ownership between OpenAI and ChatGPT users – but note, it does not mean ChatGPT users can necessarily claim copyright or patent rights to ChatGPT’s output. Indeed, “authorship” and “inventorship” rules are still being worked out for AI-created works or inventions. Companies should not automatically assume they will be able to register IP protections for works or inventions that are created or assisted by AI.

(d) Cybersecurity. Generative AI’s “everything that users input can be ingested to improve the model” approach can present security risks.

- As stated above, employees can readily create accounts with generative AI providers and start using their tools, all via a standard web browser. This potentially creates a new data loss risk vector. Existing data loss prevention tools may not detect if employees input restricted data into generative AI tools. These tools may thus potentially expose restricted data outside the organization.
- Additionally, generative AI can be used to create code. But, there does not appear to be any published research on whether their code contains vulnerabilities or malicious elements. This risk can be mitigated by running AI-created code through a security scan – but that assumes that employees using generative AI coding tools are putting AI-created code through secure software

development processes. If AI-powered coding occurs outside of secure software development cycles, it could possibly become a vector for introducing vulnerable code into the organization's source code.

- (e) Privacy and Data Protection. Most companies considering generative AI will likely be subject to U.S., EU, or UK privacy and data protection laws (or a combination of them). These require several considerations when using generative AI.
- Personal data input into generative AI tools may become part of the AI model itself, and start appearing to other users. This can have privacy compliance impacts. It is unclear whether data can be deleted from generative AI models – which could impact individuals' rights to request deletion of personal data. In some cases, U.S. and EU/UK laws can require affirmative consent to process "sensitive" data, suggesting it should not be input freely into generative AI tools.
 - Privacy laws typically require companies to classify their vendors as processors, independent controllers, or joint controllers. Depending on which role the vendor plays, certain contractual terms are mandated by law. Which role is appropriate for generative AI tools? Many current generative AI tools are silent on the issue, while emerging enterprise solutions suggest a "data processor" model.
 - Higher-risk use cases of generative AI may trigger requirements to carry out a data protection impact assessment (DPIA) under the EU/UK GDPR, or to carry out a "data protection assessment" under U.S. state privacy laws.
 - U.S. and EU/UK data protection laws regulate "automated decision-making" that results in "legal or similarly significant effects." Any use of generative AI for high-impact decisions affecting individuals or small businesses could potentially implicate these rules.
- (f) Accuracy and Reliability. Generative AI, like any new and evolving technology, should not yet be considered reliable. None of the generative AI tools on the market have made their training data public, meaning there is no indication of whether training data sets themselves display accuracy or reliability. It could be that training data included a broad set of data generally available from the internet, with the resulting wide swings in quality that one finds online. Further, ChatGPT training data reportedly all predates the year 2021, meaning answers will pre-2021 data. It cannot be determined how often generative AI tools will provide accurate or inaccurate answers. However, at present, it seems settled that generative AI tools will at times provide incorrect answers to queries. This suggests a layer of human review remains necessary at this stage.
- As an example, OpenAI CEO Sam Altman tweeted: "ChatGPT is incredibly limited, but good enough at some things to create a misleading impression of greatness. It's a mistake to rely on it for anything important right now." [ChatGPT's website](#) states: "ChatGPT sometimes writes plausible-sounding but incorrect ... answers." Thus, OpenAI appears to encourage ChatGPT users to add a layer of human review to ChatGPT outputs.
 - As another example, Google recently introduced its "[Bard](#)" feature in search as a ChatGPT alternative. But, during its first public demonstration, Bard [incorrectly stated](#) that the James Webb Space Telescope was the first to photograph an exoplanet.

- (g) Unpredictability. [Recent reports](#) indicate that generative AI, if pushed by users, can take on a “persona” whose interactions with users become disturbing – such as declarations of love or insults. In fairness, both of the linked reports involved reporters intentionally attempting to coax unexpected responses from a ChatGPT-powered search engine. A company’s internal users will typically not be doing that – instead, they simply want generative AI tools to create useful code, text, images, or documents. Still, caution is warranted in making generative AI output directly available to consumers or the public. Members of the public may attempt to “hack” AI-powered tools to make them say embarrassing or disturbing things (and indeed, there already appears to be a subreddit dedicated to doing this.)
- (h) Impact & Bias. As a general matter, AI can take on biases inherent in training data sets. OpenAI [publicly states](#) that although efforts are made to make ChatGPT refuse inappropriate requests, ChatGPT will sometimes “exhibit biased behavior.” Bias in AI is an increasing focus for regulators. In the U.S., the FTC, for example, [has issued advisories](#) requiring companies to accountably ensure AI they use is not biased or discriminatory. [New York City requires](#) AI used in recruiting or employment decisions to undergo a “bias audit”. On the other side of the Atlantic, the [UK ICO](#) has GDPR-related guidance on addressing risks of bias and discrimination in AI systems. This may again counsel against making generative AI outputs directly available to the public, or using AI outputs in ways that impact consumers, prior to an enterprise-level agreement and validation having been put in place.
- (i) Compliance Programs. Companies in regulated industries often maintain compliance programs designed to adhere to – for example – financial services regulatory regimes, anti-bribery regimes (such as the U.S. Foreign Corrupt Practices Act), and similar. Generative AI tools may not be subject to the required level of monitoring these programs require for auditability. For example, if an employee “discusses” potential securities trading strategies a generative AI chatbot, these could be communications that would otherwise need to be monitored and retained in accordance with financial services regulatory requirements.
- (j) Lawyer-Client Privilege. Lawyers using generative AI tools run the risk of waiving legal privilege in some jurisdictions. If, as noted above, information input into generative AI tools may become available to others, privilege may be waived by using otherwise privileged information in a generative AI tool. Legal departments must weigh the risk of waiving privilege before using ChatGPT or similar tools in a manner that includes placing privileged information in a prompt.
- (k) Liability may Fall to the Corporate User. At present, generative AI providers generally limit their own liability, while requiring indemnity from users. As an example, [OpenAI’s TOUs](#) limit OpenAI’s liability to \$100 USD in direct damages (or 12 month’s fees – but ChatGPT is often used for free). Users indemnify OpenAI for claims arising from or relating to “your use of the Services,” including from the “Content” output by generative AI tools. These terms are not necessarily unusual for a no-cost service. But, they may not be terms corporate counsel are used to accepting for enterprise solutions used by their business.

4. What about upcoming AI-Specific Laws?

European Union. The EU is working on finalizing a draft AI Act, which regulates “AI Systems” used for particular “High-Risk” purposes. Like the GDPR before it, the AI Act will reach beyond the borders of the EU, and apply to companies in the U.S. (and elsewhere) that place or put into service AI systems on the EU market, or that use an AI system’s output in the EU. While the draft AI Act has not been finalized (and there are outstanding questions about how “general purpose AI Systems” will be regulated), companies should exercise caution before using generative AI tools for any purpose that may be considered “High-Risk” under the draft AI Act. These include:

- Recruitment or selection of staff, or evaluating candidates;
- Making decisions on staff promotions and terminations, for allocating tasks to, or monitoring and evaluating performance and behavior of staff;
- Management and operation of critical infrastructure; and
- Purposes connected with establishing the creditworthiness of individuals.

Additionally, if a company is considering incorporating generative AI into a chatbot that directly interacts with individuals (e.g., for customer services purposes), the AI Act will require that individuals are provided with “AI transparency,” which means that they should be informed that they are interacting with an AI System (unless it is obvious from the circumstances). It may be advisable for companies to start considering how to implement this kind of “you are speaking with a bot” transparency in a customer-friendly manner.

Lastly, companies should consider that the EU is also working on a draft AI Liability Directive. This Directive would make it easier for claimants to sue a company where damage is allegedly caused by the company’s use of a High-Risk AI System (as defined in the draft AI Act). To facilitate claims, it would empower courts in the EU to order the company in question to disclose “relevant evidence at its disposal about a specific high-risk AI system that is suspected of having caused damage” to plaintiffs in some circumstances. While U.S. companies would be used to extensive discovery, this could expose companies doing business in the EU to U.S.-style discovery when AI-caused harm is at issue. Companies may wish to align their U.S. and EU recordkeeping practices to account for this.

USA: For an overview of forthcoming AI regulation in the U.S., see our A&B Advisory titled [“AI Regulation in the U.S.: What’s Coming, and What Companies Need to Do in 2023.”](#)

5. So What Do I Tell My Business?

Despite the risks, corporate counsel may find themselves in a situation where their business is pushing strongly towards obtaining access to generate AI tools. The reasons can be manifold, such as competitive pressures, internal metrics, or productivity needs. Two possible paths forward are as follows:

- (a) **Acquire an Enterprise Solution.** Counsel does not have to work on the assumption that their company will have to use generative AI tools like ChatGPT as they are presently offered. Enterprise solutions are being developed. As an example, Microsoft has announced it is offering an [“Azure OpenAI Service”](#) that grants enterprise customers access to OpenAI’s suite of generative AI tools (ChatGPT, Codex, and DALL-E) within an Azure environment. Enterprise solutions are more likely to memorialize confidentiality, privacy, security, and IP terms responsive to needs of enterprise customers in an enterprise-level agreement.

If internal business clients cannot make a compelling business case that their needs for generative AI are urgent, the more prudent approach may be telling the business to wait – while starting procurement processes to obtain an enterprise solution.





(b) **Limited, Temporary Internal Use on Risk-Adjusted Basis.** Corporate counsel may be faced with situations where business clients insist that an urgent need to use generative AI tools immediately exists. In such situations, steps can be taken to mitigate – although not eliminate – the risks involved, which could be presented to the business clients to determine if they desire to proceed. Below are measures that could be considered to help address the legal considerations discussed above. Of course, this list is not (and cannot be) comprehensive or complete, given the emerging nature of generative AI. Instead, this list can be viewed as a “toolkit” for corporate counsel to address potential legal risk in enabling business use of generative AI:

- Preapproved Users and Use Cases Only. Generative AI tools may not yet be fit for general approval in organizations. Limiting how, when, and why employees can use generative AI is one step businesses can take to address risk.
 - As a practical example, companies wanting to test Microsoft’s Azure OpenAI Service must complete a [detailed application form](#) – with Microsoft reserving the right to deny applications that appear to be for personal use, or for a “scenario we are not yet supporting because of higher potential for misuse.” Microsoft’s approach of requiring pre-approval for generative AI use cases may be advisable for companies generally.
 - Organizations could consider having employees make a documented business case for their need for generative AI, along with the specific users who would need to use it. If approved, the organization could maintain records of which users are allowed to use which generative AI systems, and for what specific purposes.
- No Personal Data or Confidential Information, Period. Until enterprise solutions with confidentiality and privacy protections arrive, counsel should likely continue to prohibit any personal data or confidential information from being input into generative AI tools. Even if a user can show a strong business need for a specific use case, inputting such information into generative AI could still expose it outside the organization.
- Internal Guidance on Generative AI Tools. Counsel could consider placing internal policy guardrails around the use of generative AI tools. As an example, Microsoft already publishes a “[Code of Conduct](#)” for its Azure OpenAI Service. To mitigate some of the risks highlighted in section 2, such guardrails may include, for example, requiring that a human assesses and revises the output of generative AI tools (a “No Unsupervised Generative AI in Our Organization” rule). When generative AI is used to write copy, draft documents, or produce code, its “work” should always be subject to human review. This accords with recommendations from the generative AI providers themselves, such as OpenAI.
- CISO Involvement is Likely Mandatory. Generative AI tools also impact security. The CISO likely needs to weigh on what uses should and should not be permitted; what functions should be permitted to use generative AI; and what controls can be implemented to monitor and secure the use of generative AI tools.

- Secure Software Development is Doubly Important. If using generative AI for coding, it would be prudent to place it early in the secure software development process. AI-generated code should be (i) manually reviewed, (ii) put through open-source and security scanning processes, and (iii) subjected to standard testing and deployment processes.
- For ChatGPT: Get an Organization ID and Opt-out. [OpenAI permits](#) users to associate their account with an “organization ID.” Once this is done, the users can use the organization ID to “opt out of having [the organization’s] Content used for improvement” of OpenAI’s AI algorithms. Of course, as stated above, the effect and extent of this “opt-out” has not yet been publicly reported, so organizations may not be able to place full reliance on it. But, it potentially mitigates risk to use OpenAI’s opt-out function.
- Determine Whether Claims Arising from Anticipated Generative AI Use Cases Would be Insured. If a generative AI tool inadvertently supplies a company with another company’s copyrighted code, and a claim arises, is there an insurance policy that covers the claim? What if sensitive personal data is inadvertently input into a generative AI tool, can’t be deleted, and begins to reappear in other users’ searches? Counsel may want to confirm that – for each use case the business is submitting – the company has insurance in place that would at least help defend a claim if something were to go wrong.
- Legal Should Not Own the Decision to Use Generative AI. Legal can outline the risks and potential mitigations, but any decision to use generative AI tools outside of a full enterprise solution should be made by the business itself.

About the Authors

Alston & Bird’s [Privacy, Cyber & Data Strategy Team](#) is continuing to monitor developments surrounding AI. Please contact us if you have any questions.

			
<p>Daniel J. Felz Partner +1 404 881 7694 daniel.felz@alston.com View Full Bio</p>	<p>Wim Nauwelaerts Partner +32 2 550 3709 wim.nauwelaerts@alston.com View Full Bio</p>	<p>Paul Greaves Senior Associate +32 2 550 3791 paul.greaves@alston.com View Full Bio</p>	<p>Josh Fox Associate +1 404 881 7869 josh.fox@alston.com View Full Bio</p>