

## Key Issues in Generative AI Transactions

25 Apr 2023

### Technology Transactions

#### Client Alert

Over the past year, we have seen a dramatic increase in the adoption of AI technologies across industries. Because transactions involving AI technologies can resemble those involving traditional software, like SaaS agreements, parties often assume that their expectations from those standard agreements about what is reasonable and “market” should also apply to AI-related technology transactions. And, in some cases, this approach is appropriate. When a vendor supplies hospitals with a hosted medical imaging platform, the fundamental nature of the transaction remains much the same, regardless of whether the platform includes AI-assisted diagnosis functionality. However, it is important not to overlook the unique issues raised by certain AI technology transactions, particularly those that involve generative AI models (“Generative Models”). This client alert offers an overview of those issues.

#### Training Data

Generative Models typically undergo two stages of training, initial training and fine-tuning.

Initial training (as the name suggests) comes first and exposes the Generative Model algorithm to massive datasets to help it gain a broad “understanding” of the data and the relationships between the data. In most cases, the provider of the Generative Model (the “Model Provider”) will have completed this initial training before the model is made available to customers and end users (“Model Customers”). For their initial training, many Generative Models (particularly **large language models**) are trained on data scraped from the Internet. There have been **years of litigation** on the permissibility of web scraping (independent of scraping for training purposes) and there are **ongoing cases** specifically addressing whether training on copyrighted materials without permission qualifies as a fair use.

Fine-tuning typically involves a much smaller and more specialized dataset that is designed to refine the Generative Model, so it is better able to generate outputs for specific use cases.

Fine-tuning may involve data provided by the Model Provider or by the Model Customer. Some Model Providers offer customers the ability to fine-tune an instance of a Generative Model for the customer’s specific use case. In these cases, sometimes the Model Customer provides the fine-tuning data, sometimes the Model Provider gets the data from elsewhere, and sometimes the parties work together to create the data.

Agreements concerning Generative Models should clarify the ownership and permissible uses of this data. For instance, in cases where the Model Customer provides data for use in fine-tuning, the agreement should clarify whether the Model Provider is permitted to use the data only for the Model Customer or also to train the generally available version of the Generative Model. Parties should also consider how to allocate liability arising from use of training data. Generally speaking, it may be reasonable for each party to be responsible for training data that it provides or sources and to indemnify the other party from claims arising from such data.

That said, where a Generative Model is initially trained on large swaths of data scraped from the Internet, the Model Provider may be reluctant to indemnify the Model Customer for claims arising from such data, particularly given that the law around such training is unsettled, as noted above. The Model Provider may argue that, at least in the case of large language models, this type of training is a necessary aspect of developing the technology and that there is no practical way for the Model Provider to affirmatively secure rights to the data. In any event, even without any contractual risk allocation, the Model Provider will likely bear most of the risk arising from use of such data for initial training, given that the Model Provider will be the party sourcing the data (*e.g.*, through scraping) and performing the activities that could give rise to claims (*e.g.*, making copies and derivative works).

#### User Prompts

User inputs (“Prompts”) are the questions, queries, and other inputs that users enter into Generative Models in order to generate outputs.

## Contacts

#### Aaron P. Rubin

[arubin@mof.com](mailto:arubin@mof.com)

(415) 268-7522

(415) 268-6809

#### Heather M. Whitney

[hwhitney@mof.com](mailto:hwhitney@mof.com)

(415) 268-7522

(415) 268-6076

#### About Morrison Foerster

We are Morrison Foerster — a global firm of exceptional credentials. Our clients include some of the largest financial institutions, investment banks, and Fortune 100, technology, and life sciences companies. The Financial Times has named us to its list of most innovative law firms in North America every year that it has published its Innovative Lawyers Reports in the region, and Chambers Asia-Pacific has named us the Japan International Firm of the Year for the sixth year in a row. Our lawyers are committed to achieving innovative and business-minded results for our clients, while preserving the differences that make us stronger.

Because of the generality of this update, the information provided herein may not be applicable in all situations and should not be acted upon without specific legal advice based on particular situations. Prior results do not guarantee a similar outcome.

Which party owns Prompts is a question that may arise in negotiations. Typically, as between the Model Provider and Model Customer, the Model Customer generally retains any ownership rights it may have in its Prompts. But a word of caution: Model Providers should avoid language that suggests that they are assigning to the Model Customer any preexisting rights the Model Provider may have in Prompts. If a Model Customer copies patent claims from the Model Provider's patents and uses them as a Prompt, surely no reasonable person would think the Model Provider intended to assign its rights in those patents. The Model Customer had no rights in the Model Provider's patents beforehand and its inputting the patent claims into the Generative Model should not give it those rights either.

Another issue that often arises in transactions involving Generative Models is confidential treatment of Prompts. Model Customers may want their Prompts to be treated as their confidential information for a variety of reasons, including to maintain trade secret protection. For trade secrets to maintain their status as trade secrets, companies must take reasonable steps to protect them from disclosure. Model Customers may be concerned that inputting Prompts containing trade secrets into a Generative Model will undermine their trade secret status, particularly if the Generative Model is trained on the Prompts and subsequently reproduces them in a later output or the Model Provider otherwise discloses the Prompts.

While these concerns are understandable, Model Providers often do not know what kinds of data customers are inputting into the model and are not in a position to do a case-by-case analysis of Prompts to determine which Prompts may contain trade secrets or other confidential information. Given this, taking on onerous confidentiality and data security obligations (and potential liability) for trade secrets or other confidential information contained in Prompts may mean, in practical terms, that the Model Provider has to comply with those obligations for every Prompt by every Model Customer. Moreover, most commercially available Generative Models are not trained on Prompts from enterprise customers and Model Providers often provide non-enterprise users an opt-out right if they do not wish to have their Prompts used for training (frequent alarm about this possibility notwithstanding).

Another concern Model Customers may have is whether Prompts containing information about new inventions would constitute a public disclosure (specifically, a "printed publication" under **35 U.S. Code § 102**) that would prevent patenting the inventions. Again, Generative Models are generally not trained on enterprise customer Prompts, though the possibility that a Model Provider could access and disclose a Prompt is a theoretical possibility. Even in that case, though, the Prompt would not constitute a public disclosure of the invention unless it describes the invention in enough detail for someone skilled in the field to replicate or use it (*i.e.*, enablement). The good news for Model Customers is that, even if their Prompts were to somehow constitute a public disclosure, they have a **one-year grace period** from the date of disclosure to file a patent application (at least under U.S. patent law).

In an attempt to avoid the above issues, where feasible Model Providers may tell Model Customers to simply not include in their Prompts confidential information, like PII, or information about inventions for which patent applications have not yet been filed. Some use cases may require Model Customers to input highly sensitive confidential information, like proprietary source code. In those cases, the parties may need to negotiate specific terms to address the treatment of such information contained in Prompts.

## **Generative Model Outputs**

The defining feature of Generative Models is that they generate outputs ("Outputs") in response to Prompts. In commercial transactions, these Outputs run the gamut; **language tutoring, code, the identification of new therapeutic targets, the saved voice of a person with ALS**, all of these are possible.

In most cases, agreements between Model Providers and Model Customers should expressly address ownership of Outputs.

The threshold question, though, is whether Outputs constitute intellectual property that can be owned at all, and that question remains unanswered. As **we have written previously**, the Copyright Office has taken the position that only works authored by humans can be copyrightable. However, there is a lack of clarity around what a human using a Generative Model must do to qualify as the "author" of an Output—even after the Copyright Office published **guidance** on this question in March. Earlier this year, the Copyright Office **revoked a**

**copyright** in a comic book with images the author, **Kris Kashtanova**, created using a Generative Model, taking the position that Kashtanova was not the “author” of the images. (*Disclaimer: Morrison & Foerster represents Kris Kashtanova.*) Recently, we **submitted** a copyright registration application on behalf of Kashtanova for “**Rose Enigma**,” an image Kashtanova generated using a Generative Model with ControlNet. We hope that the Copyright Office’s response will provide greater clarity on this issue.

Similarly, and as we have also **previously discussed**, the Federal Circuit made headlines last summer when it affirmed the U.S. District Court for the Eastern District of Virginia’s holding that an AI model cannot qualify as an “inventor” under the Patent Act; only humans can. On April 24, the Supreme Court **denied** plaintiff Stephen Thaler’s petition for a writ of certiorari. The circumstances under which a human can accurately claim to be the “inventor” of an otherwise patentable invention generated by a Generative Model are, as in the case of authorship of copyrightable subject matter, unclear.

The upshot of the above is this: parties to transactions involving Generative Models should keep in mind that there are no clear answers yet regarding whether and in what circumstances Outputs can be “owned” at all in the intellectual property sense.

Moreover, even if Outputs are protectable intellectual property, allocating intellectual property rights appropriately between the parties can be more challenging than it first appears. In traditional transactions where vendors create customized content for customers, vendors often assign all right, title, and interest in that content to the customer, but that may not be appropriate in the generative AI context. Outputs may, for instance, incorporate or be substantially similar to third-party material, including material contained in the data the Generative Model was trained on. Obviously, the agreement between the Model Provider and the Model Customer cannot allocate ownership of intellectual property that belongs to third parties. Therefore, the agreement should at least carve out such third-party material from any terms purporting to allocate ownership of intellectual property as between the Model Provider and the Model Customer.

Parties should also consider what happens if Outputs are substantially similar to the *Model Provider’s* preexisting intellectual property. Imagine a case where the Model Customer is able to get a Generative Model to disclose parts of the Generative Model’s own source code; the Model Provider presumably would not be willing to assign *those* intellectual property rights to the Model Customer. To deal with this problem while also providing comfort to Model Customers, the parties may choose to omit any express assignment language from their agreement and simply rely on background intellectual property law to allocate ownership of Outputs between the parties. Or, if the Model Customer insists on having the Model Provider assign whatever intellectual property rights it may have in Outputs, the Model Provider will want to carve out from the assignment any preexisting intellectual property of the Model Provider and anything derived from such preexisting intellectual property.

While these carve-outs for third-party material and preexisting intellectual property of the Model Provider may help address the Model Provider’s concerns, such carve-outs can create additional issues for the Model Customer. In particular, the Model Customer will not necessarily know whether Outputs do or do not contain third-party material or preexisting intellectual property of the Model Provider, so there may be uncertainty as to what Outputs are owned by the Model Customer. There is no easy solution to these ownership issues, but on balance it may be more reasonable for the Model Customer to be responsible for doing the necessary due diligence (*e.g.*, code scans, clearance searches, etc.) to determine the provenance of Outputs, to the extent it is possible to do so. After all, the Model Customer is the one generating the Outputs through its Prompts, and the Model Provider will typically not even be aware of the specific Outputs being generated.

The issues with Outputs do not end there. Putting to the side preexisting intellectual property, in theory a Generative Model could create virtually identical outputs for two or more Model Customers. Model Providers should make Model Customers aware of this possibility and clarify that, while Outputs generated by the Model Customer may belong to that customer, that does not mean that the Model Provider is guaranteeing that other customers will not generate similar or identical Outputs using the Generative Model (and, if they do, those other customers’ Outputs may belong to them).

In cases where Model Customers own Outputs, the question arises as to what rights, if any, the Model Provider retains in them—for example, the right to train a generative AI model using them. As noted above, many Generative Models are provided under terms that do not permit the Model Provider to use Prompts for training (at least for enterprise customers) and the same approach may make sense in many cases for Outputs. But, of course, this can be negotiated on a case-by-case basis.

Another important issue in any transaction involving Generative Models is allocation of liability for claims arising from Outputs. The issue of “hallucination,” where a Generative Model confidently provides incorrect information, has been much in the news lately. Such hallucinations could result in a variety of claims, anything from defamation to false advertising to product liability. In addition, a number of lawsuits have been filed alleging that Outputs infringe preexisting copyrighted works. If such claims arise from use of a Generative Model, who should be responsible, the Model Provider or the Model Customer?

A Model Customer who does not have a sophisticated understanding of how generative AI works may see the Generative Model as simply a way to obtain content. And the Model Customer may have a general expectation that when it obtains content from a content provider, the content provider should take responsibility for that content through representations and warranties, indemnification, and other contractual mechanisms to allocate risk.

But viewing the Model Provider as merely a content provider is an oversimplification of how a Generative Model works. In reality, it is quite difficult for a Model Provider to take on this kind of liability for Outputs. Certainly, Model Providers can incorporate guardrails into their models; for example, a Model Provider could prevent its model from generating images of well-known copyrighted or trademarked characters. But designing a Generative Model that is incapable of producing any potentially infringing, defamatory, false, discriminatory, or otherwise problematic Outputs is very difficult, if not impossible.

Moreover, as noted above, Model Providers do not ultimately control which Prompts its Model Customers, or Model Customers’ downstream users, enter. Yet, those Prompts go a long way toward determining what Outputs the Generative Model generates. (Indeed, in most cases where Generative Models have generated Outputs that include what appears to be third-party content, users went to great lengths to push the model into generating that potentially infringing content.) At the same time, though, while Model Customers and their downstream users can control what Prompts they enter, they do not control whether the Generative Model takes their innocuous Prompt and provides a potentially problematic Output in response. Nonetheless, in practical terms, it may be more reasonable for the Model Customer to be the party that is ultimately responsible for deciding whether the Outputs are suitable for use.

One topic we have seen raised is whether statutory safe harbors—**Section 230** of the Communications Decency Act, specifically—apply to Outputs. (Justice Gorsuch raised this issue during oral arguments in *Gonzalez v. Google* this past term.) Section 230 is the federal statute that provides broad immunity to online platforms and other providers and users of “**interactive computer services**” for liability arising from third-party content, such as posts and comments from users. For example, if a user posts defamatory comments about another user on a social media platform, the defamed user can pursue a claim against the other user but, because of Section 230, the platform will not be liable for publishing the comment. Some have argued that Section 230 immunity should extend to Outputs because Outputs are derived from third-party content, namely the training data. The difficulty with this argument is that Section 230 immunity does not apply in cases where the platform operator itself is responsible *in whole or in part* for the illegality of the content at issue. A plaintiff is likely to argue that a Model Provider is responsible at least in part for the Outputs of the Generative Model.

### **Other Issues**

Sometimes, a Model Provider may provide Model Customers with access to a third-party Generative Model (or an API for a third-party Generative Model), including as part of a larger services offering. In such cases, the Model Provider and Model Customer should review the terms that apply to the third-party Generative Model to confirm that the Model Customer’s intended use case is consistent with those terms. Some Generative Models are distributed under terms that prohibit commercial use (*e.g.*, **AlexaTM 20B**) or the creation of derivative works (*e.g.*, **AlexaTM 20B** and **Lyra-Fr 10B**), while others impose specific use restrictions, intended to

enforce the responsible use of Generative Models, which must be flowed down to end users (e.g., **Claude** and models licensed under a **Responsible AI License** [RAIL]).

Finally, jurisdictions around the world are implementing laws and regulations targeting the use of Generative Models and the Generative Models themselves. Parties to transactions involving these technologies need to consider how to address compliance with applicable laws and regulations. While it is straightforward to say that each party is responsible for its own compliance in connection with its performance under any agreement, in some cases the Model Customer's use of the Generative Model may trigger compliance obligations on the *Model Provider* that the Model Provider would not otherwise have had.

### **Concluding Thoughts**

By staying up to date on the rapidly evolving legal landscape surrounding Generative Models, both Model Providers and Model Customers will be better equipped to successfully navigate the complexities of these transactions and enter into agreements that make sense.

© 2023 Morrison & Foerster LLP Client Alert [www.mofo.com](http://www.mofo.com)