

Publications

Key Contract Terms and Conditions for AI Products and Services Part 2 – Commitments, Disclaimers, Regulatory, Privacy, and Risk Allocations

Related Attorneys

Craig R. Auge

Related Services

Copyrights

Data Strategy, Privacy and Security

Intellectual Property

Patents

Technology Transactions

Trade Secrets

Related Industries

Colleges and Universities

Financial Institutions

Health Care

Insurance

Manufacturing

Private Equity

Restaurants, Food and Beverage

Retail and Consumer Products

CLIENT ALERT | 9.7.2023

*This is the second installment on key contract terms and conditions for AI products and services and is more from a customer-side vantage. A **first installment** covered data ownership and licensing. Scope note – this is **not** about AI preparing contract forms.*

When entering contracts for artificial intelligence (AI) products and services, providers and customers should apply well-known legal concepts to lesser-known AI elements. Parties experienced with SaaS agreements will find some familiar landscape.

But they also should be prepared for unique issues and more complexity – and potential risks and liabilities – associated with an array of types and sources of data that train, fuel, guide, emanate from, and modify generative AI-based models, solutions, and systems (AI Solutions).

For convenience, “provider” references in this Client Alert could be the actual developers of the foundational model for an AI Solution or application developers that build on that foundational model or otherwise make it available (e.g., as a part of larger offerings from Microsoft).

Data Types and Sources

- **Training**

The algorithm that is a part of the AI Solution is exposed by the provider to massive data sets to “train” it, typically before being made available to customers. *Think of it as “initial training data.”*

Subsequently, the initial training data is often fine-tuned, improved, and optimized by more data added as a new layer, sourced from data contributed by the provider, by the customer, by both of them jointly, or by or through a third party from which

they have obtained it. *Think of it as “enhanced training data.”*

- **Prompting**

Customers provide the AI Solution with prompts, instructions, queries, data (including, for example, Internet-of-Things device data), or other input. *Think of it as “input prompts.”* It is possible that input prompts become a part of enhanced training data.

- **Generating Outputs**

In response to input prompts, the AI Solution generates responsive output – such as other data, text, images, video, audio, new code, or other materials or content. *Think of it as “output.”* Output could become a part of enhanced training data.

Commitments and Disclaimers

- **Use**

A provider may license use of the AI Solution to a customer for limited purposes and with restrictions. A provider concerned about unintended uses or associated risks may go beyond the software license’s and SaaS agreement’s common “*for internal business purposes*” to address specific parameters, such as permitted types of input prompts or what output may not be used for.

- **Performance**

Like other cloud and platform agreements, uptime and accessibility service levels (“SLAs”) may be appropriate.

As with a software license or SaaS subscription, an AI provider may commit that its AI Solution is designed to perform certain functions or for a use case, operate as a tool for particular business needs, or integrate with particular devices. However, due to potential variability of input prompts and (what providers will point to as) lack of control over output, AI Solution providers are more likely to make performance commitments as general and vague as possible.

A unique facet of some AI Solutions is that they should improve over time – due to more enhanced training data.

- **Quality of Output**

Generative AI Solutions sometimes produce “hallucination” output, which may appear correct or accurate but is not. By contract, providers may warn of that possibility, alongside broad express disclaimers that they do not warrant the accuracy of results or how customers decide to use output.

Beyond common disclaimers of express warranties not otherwise in the agreement and implied warranties, a provider may also include cautionary acknowledgements that customer will carefully review and validate output before relying on it.

- **Customer's Data Benefiting Others**

If a customer's input prompts or output are contributed back as enhanced training data to improve the AI Solution generally, customer may ask provider for a broad liability disclaimer or indemnity, in connection with provider's further exploitation with others.

Practice Tips

* Customer due diligence can be essential.

* Unlike well-established software or SaaS solutions that have recognizable brand names or version numbers, or which are described at vendors' websites or in documentation, newer and rapidly evolving AI Solutions may lack description. And without some description, establishing failed commitments and entitlement to remedies are made that much harder. Contracts should ideally describe, at least briefly and higher level, the AI used (e.g., natural language processing), sources of training data, and use cases.

* While challenging to draft upfront provisions to measure AI Solution improvement over time (e.g., levels of key performance indicators) or periodically raise the performance bar, critical or niche AI Solutions may justify the effort.

Regulatory and Privacy

- **Regulatory Environments**

In the United States, regulatory initiatives are percolating at both the federal and state levels. Laws are catching up that prohibit or require safeguards for types of AI and particular uses. Some apply horizontally (across the spectrum of AI) and others vertically (targeted to specific technologies or industries). For instance, several U.S. state privacy laws now contain restrictions surrounding automated decision-making and profiling, including rights for consumers to "opt out" of those functions.

- **Personal Information**

Regulated data and in particular personal information – as a part of training data, input prompts, and output – create layers of potential liabilities. For example, if a customer submits input prompts containing health information governed by HIPAA or personal information subject to GDPR or U.S. state privacy laws, is a second use as a part of output or (if contributed back to the AI Solution) as enhanced training data permitted? Or permitted pursuant to a Business Associate Agreement or Data Processing Agreement?

Where possible, parties may also commit contractually to using synthetic or de-identified data or data minimization techniques.

A customer will want to impose certain data privacy and protection requirements that it is subject to on the provider receiving customer's input prompts or storing output.

- **Potential Bias**

Customers covered by anti-discrimination regulations in making hiring, lending, or other decisions will want the above-referenced types of disclosures on how the AI Solution was trained and operates to determine if biases were or can be introduced that could lead to inaccurate or unfair decision making. This could occur if the training data was biased or the algorithm as designed calculates in a potentially unfair manner.

Providers concerned about unintentionally taking on liability may include a customer acknowledgement that automatic decisions are not determined solely by output, but instead require customer's human review and final decision.

Practice Tips

* Because "applicable law" likely differs for a customer and a provider and may be unclear or yet to mature, a provider may want to try to narrow the common representation that it will perform in compliance with applicable laws to "as may be applicable to provider as an AI Solution services provider as described in the agreement" and "as of the effective date of this agreement." A customer may want to resist this narrowing or name the particular law, such as "comply with HIPAA."

Risk Allocations

- **Determining if AI Was Used**

Where a customer is not intentionally signing up to use AI, but wants to assess possible reputational and legal risks if its provider is using AI in services being provided or in creating deliverables, a customer can add a provider representation that "no AI was or is used," forcing disclosures.

- **IP Infringement**

In software licenses and SaaS agreements, providers customarily indemnify and defend a customer if their software or service infringes a third party's intellectual property rights. Similarly, an AI provider may step up and cover infringement stemming from the AI Solution's algorithm.

But AI Solutions spawn many possible infringement risks, compounded by the derivative nature of output that incorporates, builds from, or modifies training data. A provider may resist covering infringing initial training data, citing copyright fair use or the inability to obtain sublicensing permissions from its upstream data sources (particularly true if web-scraped from the internet).

Most challenging is output – which could be infringing due to initial training data, enhanced training data, input prompts or what a customer does with it. Providers will argue that liability for copyright infringing output, traced back to input prompts chosen by customers, rests with customers.

- **Lacking Rights to Use**

A kindred risk to IP infringement is breach of contract, germinating from training data and input prompts obtained by providers and customers pursuant to licenses with third parties – who granted

certain limited rights to use only to those providers or customers, not for them to further share or for other secondary uses, such as training or querying an AI Solution.

If open-source license terms govern portions of code or content (and, in effect, flow down to output), a customer will want to be aware if output it distributes is infected with “copyleft” terms, requiring customer to make its associated proprietary code or content available.

- **Risk Profiles and Gaps**

There other potential risks depending on the type of AI and use cases that parties can address by contract – product liability, false advertising, and defamation, to name just a few.

With the more nascent AI Solutions and use cases (and depending on negotiating leverage), the “residual risk bucket” for those risks that are neither apportioned to provider nor to customer may simply be larger than for other types of transactions.

Parties may also include in their agreements required types and amounts of insurance, to gain some shelter from liabilities.

Practice Tips

* Risk allocations coalesce in liability caps and exceptions, exclusions of certain types of damages, and alternative caps. With the newness of generative AI and variety of uses, there is no pre-set “map” of what is commercially reasonable or market, but an understanding of the applicable AI model, types of data, and use cases provides a “compass” as to what direction to head when negotiating caps and exceptions.

When using AI products and services, be sure to consult your Vorys attorney.

Prior installment on key contract terms and conditions for AI products and services – data ownership and licensing considerations.