

MAYER | BROWN

AI SUMMIT 2025

Thought Leadership Materials

Panel Session 4: Contracting of AI

MAY 30, 2024

AI-SPECIFIC REPRESENTATIONS IN TECH M&A

AUTHORS: RICHARD M. ASSMUS, NINA L. FLAX, ARSEN KOURINIAN

Navigating the acquisition of any company which makes substantial use of artificial intelligence (AI) requires a nuanced understanding of both its technological intricacies and legal complexities. As the landscape of AI continues to evolve rapidly, we expect to encounter myriad representations and warranties aimed at specific issues in intellectual property rights, data rights, and regulatory compliance relating to AI. In this Legal Update, we delve into core representations in IP and data privacy, the concerns they cover, and certain representations and warranties we have seen in transactions specific to AI—which, to date, seem to be largely duplicative to previously existing, standard representations, but with AI-specific language.

INTELLECTUAL PROPERTY

There are a few key representations regarding intellectual property:

- *Ownership*: A representation that the target company owns what it claims to own.
- *Non-Infringement*: A representation that the conduct of the business of the target company does not infringe on any third-party intellectual property rights.
- *Sufficiency*: A representation that the target company has available to it (through ownership or licensing) all of the intellectual property it needs to operate its business.

In considering AI companies, the core intellectual property representations will cover—with a few notable exceptions—many of the relevant issues. For example, if the target company lacked the rights to use the data that it used to train its model, then the non-infringement representation would be breached, absent an appropriate disclosure. Nonetheless, we see representations specific to AI—defined broadly to mean machine learning, reinforcement learning, deep learning and other artificial intelligence technologies:

- Concerning data used to train, refine, or improve the AI technology, a representation that the target company has both (i) obtained all required licenses or consents to collect and use all such inputs in the course of its business; and (ii) complied with all use-restrictions governing its collection, and the use of such inputs. These representations are informed by the tidal wave of US training-data litigation we've seen in recent years against major AI tool providers, not by any gap in the standard buy-side representation regarding non-infringement.
- A representation that the target company has implemented procedures designed to (i) ensure that any of its AI technology is reasonably reproducible; and (ii) enable the company to substitute or replace such AI technology with something reasonably similar. These representations address

concerns raised by the White House's Blueprint for an AI Bill of Rights. In our view, they reflect a concern that AI will be subject to regulatory disruption, and are additive to standard representations on compliance with law.

- Very broadly, a representation that the target company has not used any generative AI in a manner that adversely affects the ownership, validity, enforceability, or protection of any of its owned intellectual property. While speaking to intellectual property rights themselves, this could be drafted in a manner that could be seen as overreaching, and not, in our view, easily made, in light of the ever-changing AI landscape.

The most recent NVCA Model Documents (updated in April 2024) include a few representations specific to “Generative AI Tools,” which they define as “generative artificial intelligence technology or similar tools capable of automatically producing various types of content (such as source code, text, images, audio, and synthetic data) based on user-supplied prompts.” In relation to intellectual property rights, investors are asking the target company to represent that:

- It has not “used Generative AI Tools to develop any material Company-Controlled Intellectual Property that the Company intended to maintain as proprietary in a manner that it believes would materially affect the Company’s ownership or rights therein.” This representation reflects a concern that the target company owns what it claims to own, and would similarly be covered by the existing core representations. It also reflects a concern that AI-produced IP may not be protectable for lack of a human author or inventor. This representation is also overly broad, considering the definition of Generative AI Tools. For example, if a company uses an email service that (1) uses AI to draft messages based on prompts; (2) uses a hybrid language generation model to show wording suggestions; (3) includes a feature that can generate options for a prompt response to an email you receive; or (4) uses machine learning to categorize emails into different “tabs,” it has arguably used Generative AI Tools to create all emails. It is unclear if all such emails would be material, though they would certainly be Company-Controlled Intellectual Property, which is also defined very broadly to include “all intellectual property rights [owned by the company], whether registered or unregistered, that are recognized in any jurisdiction of the world.”
- It has not included any sensitive personal information, trade secrets, or other confidential/proprietary information in any prompts or inputs in tools that utilize prompts or inputs to improve the AI tool. This representation addresses the protection of this company information, which is typically covered by the standard representation that the target company has taken commercially reasonable steps to protect and maintain the confidentiality of its trade secrets.
- It has used AI tools in material compliance with license terms, consents, agreements, and laws. This representation goes to non-infringement and contractual and legal compliance, though also raises concerns with accuracy, given the evolving AI regulatory landscape.

DATA PRIVACY

As with intellectual property, there are a few key representations that will cover many AI-related issues, such as:

- *Compliance with Laws:* A representation that the target company maintains policies regarding data privacy, protection, and security that comply with applicable laws, and that its business has been conducted in compliance with applicable laws.
- *Data Treatment:* A representation that the target company has created, collected, used, stored,

maintained, processed, recorded, distributed, transferred, received, imported, exported, accessed, manipulated, or otherwise taken actions with respect to data, in compliance with not only applicable law, but with applicable contracts and/or pursuant to rights in its favor, as well as with its own privacy policies.

Using the same example, if the target company lacked rights to use the personal data that it used to train its model, then the data-treatment representation would be breached, absent an appropriate disclosure, and the target company would violate the compliance with laws representation.

We are not seeing any representations specific to AI, other than the representations regarding data described above. While the general compliance with law representation would likely address this—since existing data privacy and other subject-area regulators (e.g., the Equal Employment Opportunity Commission and Federal Trade Commission) oversee AI in the United States—it could make sense to use an AI-specific representation in the future as well. The issues an AI-specific representation might address would pertain to the governance aspect of AI use and/or development, such as: a representation that the target company has adopted and implemented written policies, procedures, and technical documentation and logs, based on commercially reasonable AI governance laws, frameworks and guidelines that address risk management, impact assessments, data governance, risk mitigation measures, common AI ethics principles (transparency and explainability, nondiscrimination, fairness, accuracy, human-centered, human-involved, robustness, privacy-enhanced, safety, security, and contestability), continuous monitoring, and accountability.

AUTHORS

PARTNER

RICHARD M. ASSMUS

CHICAGO +1 312 701 8623

RASSMUS@MAYERBROWN.COM

PARTNER

NINA L. FLAX

NORTHERN CALIFORNIA (PALO ALTO & SAN FRANCISCO)

+1 650 331 2070

NFLAX@MAYERBROWN.COM

PARTNER

ARSEN KOURINIAN

LOS ANGELES +1 213 229 5141

AKOURINIAN@MAYERBROWN.COM

Mayer Brown is a global legal services provider comprising associated legal practices that are separate entities, including Mayer Brown LLP (Illinois, USA), Mayer Brown International LLP (England & Wales), Mayer Brown Hong Kong LLP (a Hong Kong limited liability partnership) and Tauil & Chequer Advogados (a Brazilian law partnership) (collectively, the “Mayer Brown Practices”). The Mayer Brown Practices are established in various jurisdictions and may be a legal person or a partnership. PK Wong & Nair LLC (“PKWN”) is the constituent Singapore law practice of our licensed joint law venture in Singapore, Mayer Brown PK Wong & Nair Pte. Ltd. More information about the individual Mayer Brown Practices and PKWN can be found in the [Legal Notices](#) section of our website.

“Mayer Brown” and the Mayer Brown logo are the trademarks of Mayer Brown.

Attorney Advertising. Prior results do not guarantee a similar outcome.

Rethinking Tech Contract Terms For Governance Of AI Use

By **Marina Aronchik and Samuel Hartman** (November 27, 2023, 5:09 PM EST)

Artificial intelligence — the combination of computer science and data to solve problems, including through the use of algorithms that attempt to make predictions or classifications based on input data — is booming.

AI tools like ChatGPT and DALL-E have captured the public consciousness,[1] and the world's largest technology companies, including Google LLC,[2] Microsoft Corp.[3] and Amazon.com Inc.[4] have announced significant investments in AI technology.

Because the AI technology marketplace is dynamic and rapidly evolving, so too are the relevant legal terms for deals involving AI. In this article, we discuss several high-value terms lawyers should carefully consider and address in a deal that involves AI.

The Need for New or Different Terms in AI Deals

Existing deal terms and traditional considerations in technology deals are often inadequate for AI.

Businesses might expect existing agreements — including, on the procurement side, those with major technology providers like Google, Microsoft and Amazon — to be sufficient to cover AI.

Similarly, one might assume that existing contract templates and negotiation playbooks can be tweaked to cover AI-specific deal points. Yet, providers may — and often do — introduce new contractual terms for AI technology, including through links to online terms or a registration process requiring customers to consent to separate legal terms in order to use AI products.

It is also possible — though in our experience, less common — for a provider to insist that a contract for AI be entirely separate from the existing enterprise deal.

The use of AI technology by an enterprise requires technology lawyers to revisit, and in many cases, reimagine existing terms across the full spectrum of relevant contracts, ranging from procurement agreements and data licenses to customer contracts.

There are several examples of the terms and conditions governing the provision and use of AI that should be considered and addressed to mitigate the risks attendant to AI technology.



Marina Aronchik



Samuel Hartman

Rights to AI Input, Training Data and AI Model Improvements

The concepts of AI input including prompts, training data and model improvements — or the data processed in the AI tool and results of this processing, and the related allocation of intellectual property and other rights — are similar to the traditional constructs of customer data, usage data and foreground IP (as distinguished from background or preexisting IP) but transformed in an AI context.

In both cases, the contracting parties need to have a clear understanding of the contractual terms that govern the use of data provided to the AI tool by end users, or otherwise collected or processed by the AI tool. In an AI deal, however, the inputs into an AI tool, and the data training the AI model in such a tool, continuously refine and improve the model and thereby become inextricably linked.

For that reason, the issues of confidentiality, data rights and restrictions, and IP rights — in both data and improved AI models — are more complex and interrelated.

The analysis and associated risk assessment of these terms needs to be performed carefully on a use-case-by-use-case basis, identifying any inconsistencies and ambiguities in the proposed approach given the AI technology and relevant use case, and considering the value, risks and restrictions associated with each category of data and technology.

Any negotiated changes to the terms impacting these constructs should be traced through the agreement, such that, for example, by obtaining IP rights in a particular AI model improvement, a customer does not lose protection of the warranties that otherwise apply to the AI tool. Or, conversely, by foregoing ownership rights in an improvement to the AI model, a customer does not also grant broader-than-intended rights to the corresponding AI input.

If a customer is not able to secure appropriate contractual protections in a given area relative to the anticipated use cases, technology lawyers may be able to collaborate with technical teams to identify operational mitigation measures.

If these measures are insufficient, any remaining concerns may require narrowly tailored adjustments to the underlying use case, such as limiting the data sets that are exposed to the AI tool.

Rights to AI Output

The allocation of rights to the output of AI models raises issues closely linked to those discussed above for AI input, training data and model improvements.

Customers' contractual rights to AI output are often limited or, worse, ambiguous, which is particularly problematic given the uncertainty of IP protection that may accrue to the AI output under existing IP laws. The starting point should be to resolve any ambiguity on this important point and consider whether the express allocation of the ownership or use rights and any corresponding limitations are appropriate for the relevant use case.

There is also a risk of third-party challenges to the customer's negotiated rights to AI output, in part because AI output is often based on or derived from vast data sets obtained from a variety of sources — including publicly available data or data of other users of the AI tool — and therefore subject to a variety of use restrictions.

To assess the risk of these potential claims, you would need to conduct relatively extensive due diligence on relevant AI technology, such as: (1) the manner in which the applicable AI model was trained; (2) the data absorbed by the trained AI model; (3) the sources of such data; and (4) the confidential nature of such data and other restrictions on its use.

From the customer's perspective, it is important that these diligence disclosures be properly reflected in the relevant contract as representations, warranties and covenants, including in connection with a noninfringement warranty described in more detail below.

The AI Noninfringement Warranty

Based on extensive IP challenges and related litigation,[5] the noninfringement warranty is a key and often difficult issue in a deal involving AI, with coverage of AI models and their improvements, and AI output at the top of the list of concerns, together with the allocation of responsibilities for defense and indemnification of infringement claims.

Customers should seek to negotiate targeted provisions to address this important topic, with a particular focus on potential claims by third parties relating to the use of their content or other data to train the model, or the AI model or its improvement.

Thorough due diligence or, if and when available, trustworthy third-party certifications of compliance with third-party consents, licenses and other restrictions in the use of training data and other pertinent aspects of development and monitoring of the underlying model, will also help businesses assess the likelihood of adverse claims.

In addition to these measures, given the current landscape of IP challenges and related litigation, users of AI technology should consider supplementing any contractual protections relating to AI output with infringement searches — mirroring any existing processes for IP reviews in connection with creation and use of new data or materials — and potentially corresponding allocation of costs of these searches with the provider.

The AI Performance Warranty

In an AI context, an ordinary performance warranty that the AI tool complies with documentation may present a major challenge because: (1) some AI models, by their nature, are constantly evolving based on continuous training; and (2) the requirements for an AI solution may, depending on the industry and relevant use cases, be grounded in one or more of the existing frameworks and standards, such as those based on the concepts of responsible AI,[6] or other similar businesswide standards and governance processes that your organization may adopt for AI technology.

Rather than relying on existing documentation, a technology lawyer should collaborate with the relevant stakeholders to identify a clear list of parameters by which the parties will measure whether the AI tool or technology meets a contractual standard. To do so, you may consider setting a quantitative target or functional requirement for the AI tool or the output that it generates.

For example, performance warranties can be based on availability, i.e. uptime, or predictive power of the AI tool, a specified percentage in the accuracy, precision, or consistency of the answers, or an increased speed of response to customer questions.

With respect to accuracy and precision of AI output in particular, while a performance warranty is helpful, it may be prudent for an organization to implement a separate verification process or supplement the AI tool with a separate accuracy-checking solution.

As your organization establishes and advances AI governance efforts, including by implementing the requirements of trusted AI legal frameworks, AI agreements should take into account developed standards and policies.

Trusted AI Legal Frameworks and Compliance With Laws

The simplicity of an ordinary course of compliance with laws representation and its related indemnity belies the complexity of regulatory changes in a growing number of jurisdictions.

Key among them are trusted AI legal frameworks emerging in the leading proposed regulations, including the AI Act in the European Union,[7][8] the White House's Blueprint for an AI Bill of Rights[9] and the National Institute of Standards and Technology's Artificial Intelligence Risk Management Framework in the U.S.,[10] the U.K. Information Commissioner's Office Guidance on AI and Data Protection,[11] and China's Draft Measures for the Management of Generative Artificial Intelligence Services.[12]

Perhaps surprising in their consistency, these emerging frameworks tend to be modeled after the EU's AI Act, contemplating a risk-based approach for regulating AI, with compliance requirements driven largely by categorization of each AI use case into one of four established categories — prohibited, high-risk, medium-risk or low-risk — with high-risk use cases triggering the most extensive reviews and safeguards.

Beyond the AI-specific laws and legal frameworks, the principles remain that: (1) responsibility for compliance with laws should be allocated to the party that is in the best position to control the relevant area and defend the claim and (2) the use of AI or dependence on a third-party AI provider to satisfy legal requirements does not change these underlying legal requirements.

But for a variety of reasons — key among them uncertainty and relative leverage of the parties — contractual solutions to the allocation of responsibility for compliance with laws in the AI space vary significantly and must be considered on a case-by-case basis.

It would not be surprising if, in the near future, contractual responsibility for compliance with AI laws were carved out and addressed separately from the general compliance with laws warranty — similar to the approach to data protection laws — in part based on the need to address regulatory requirements with specificity in the contract, and to supplement them with ongoing operational reviews that may be more extensive than what has been "operationalized" in connection with data protection laws.

Conclusion

From business seminars to boardrooms, deal makers are evaluating the ways in which AI can streamline mission-critical processes and generate value for companies. In parallel, businesses are racing to develop internal buffer policies governing the use of AI.

As companies leverage AI tools, technology lawyers and their clients must be aware of the impact of AI

on technology deals, and develop plans for its use that include both contractual and operational safeguards.

Ultimately, clients will benefit from changes to traditional deal terms and novel contractual clauses drafted for new, AI-specific considerations.

Marina Aronchik is a partner and Samuel Hartman is an associate at Mayer Brown LLP.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

[1] Kevin Roose, The Brilliance and Weirdness of ChatGPT, N.Y. Times, December 5, 2022, at <https://www.nytimes.com/2022/12/05/technology/chatgpt-ai-twitter.html>.

[2] Nico Grant and Cade Metz, Google Releases Bard, Its Competitor in the Race to Create A.I. Chatbots, N.Y. Times, March 21, 2023, at <https://www.nytimes.com/2023/03/21/technology/google-bard-chatbot.html>.

[3] Jordan Novet, Microsoft's \$13 Billion Bet on OpenAI Carries Huge Potential Along with Plenty of Uncertainty, CNBC.com (April 9, 2023 at 10:40 PM EDT), <https://www.cnbc.com/2023/04/08/microsofts-complex-bet-on-openai-brings-potential-and-uncertainty.html>.

[4] Jordan Novet, AWS Is Investing \$100 Million in Generative A.I. Center in Race to Keep Up with Microsoft and Google, CNBC.com (June 22, 2023 at 6:33 PM EDT), <https://www.cnbc.com/2023/06/22/aws-invests-100-million-in-generative-ai-as-it-sees-a-long-race-ahead.html>.

[5] Richard Assmus and Emily Nash, Generative Artificial Intelligence and Intellectual Property, (2023).

[6] Ayesha Gulley, Why We Need to Care About Responsible AI in the Age of the Algorithm, World Economic Forum (March 22, 2023), <https://www.weforum.org/agenda/2023/03/why-businesses-should-commit-to-responsible-ai/>.

[7] Draft of the European Commission's proposed AI Act (May 11, 2023) at <https://www.europarl.europa.eu/news/en/press-room/20230505IPR84904/ai-act-a-step-closer-to-the-first-rules-on-artificial-intelligence>.

[8] Ana Bruder, Oliver Yaros, and Ondrej Hajda, Developments from Europe: New Laws Regulating Artificial Intelligence, Data and Digital Operational Resilience (2023).

[9] Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People, The White House (October 2022) at <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

[10] Artificial Intelligence Risk Management Framework, US Department of Commerce, National Institute of Standards and Technology (January 2023)

at <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

[11] Guidance on AI and Data Protection, Information Commissioner's Office (March 15, 2023) at <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>.

[12] Measures for the Management of Generative Artificial Intelligence Services (Draft for Comment), Cyberspace Administration of China (April 12, 2023) at <https://digichina.stanford.edu/work/translation-measures-for-the-management-of-generative-artificial-intelligence-services-draft-for-comment-april-2023/>.

MAYER | BROWN

AMERICAS | ASIA | EMEA

MAYERBROWN.COM

Mayer Brown is a leading international law firm positioned to represent the world's major corporations, funds, and financial institutions in their most important and complex transactions and disputes. Please visit www.mayerbrown.com for comprehensive contact information for all our offices. This Mayer Brown publication provides information and comments on legal issues and developments of interest to our clients and friends. The foregoing is not a comprehensive treatment of the subject matter covered and is not intended to provide legal advice. Readers should seek legal advice before taking any action with respect to the matters discussed herein. Mayer Brown is a global services provider comprising associated legal practices that are separate entities, including Mayer Brown LLP (Illinois, USA), Mayer Brown International LLP (England & Wales), Mayer Brown (a Hong Kong partnership) and Tauil & Chequer Advogados (a Brazilian law partnership) and non-legal service providers, which provide consultancy services (collectively, the "Mayer Brown Practices"). The Mayer Brown Practices are established in various jurisdictions and may be a legal person or a partnership. PK Wong & Nair LLC ("PKWN") is the constituent Singapore law practice of our licensed joint law venture in Singapore, Mayer Brown PK Wong & Nair Pte. Ltd. Details of the individual Mayer Brown Practices and PKWN can be found in the Legal Notices section of our website. "Mayer Brown" and the Mayer Brown logo are the trademarks of Mayer Brown.