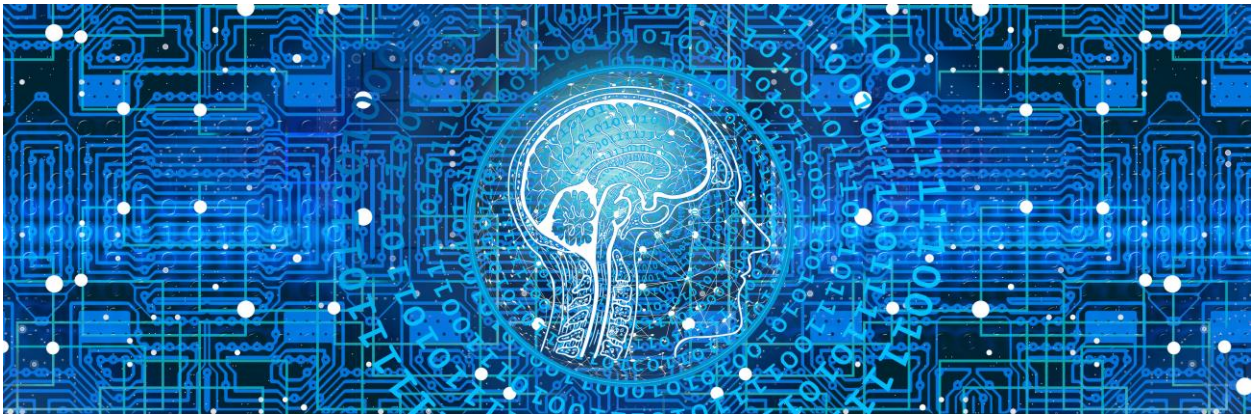

Practice Note: Issues to Consider When Contracting for AI Systems

By: Jeff Monassebian[1]

Introduction:

Artificial Intelligence Systems, or "AI Systems," employ algorithms and software programs with the goal of enabling computer systems to emulate "human-like" analysis to generate Output and recommend actions. AI Systems are rapidly being deployed within the technology infrastructure of organizations across all industries. As such, counsel must:

- understand how AI Systems operate, ingest and analyze data, and create Output;
- advise clients with respect to the business and legal issues unique to licensing and using AI Systems; and
- ensure terms in the AI license agreement have been informed by careful consideration of such issues.

Understanding Terminology:[2]

- *Text and Data Mining ("TDM")*: Text and Data Mining is the process of extracting relevant information from vast amounts of data for purposes of identifying patterns, relationships, trends and anomalies.

- *Training Data*: Training Data includes data extracted during TDM. Training Data may be (i) tagged or labeled (typically by human annotators) to generate targeted results, (ii) not tagged or labeled, in which case the AI System autonomously groups and/or classifies data by common attributes, but with no intended target, or (iii) a hybrid of both approaches.

- *Machine Learning*: Machine Learning falls under the umbrella, and is a subset, of Artificial Intelligence. Machine Learning utilizes algorithms to learn from patterns, relationships, trends and anomalies in data.

- *Production Data*: Databases and other information repositories against which an AI System is executed to generate Output.

[1] Jeff Monassebian is the Managing Attorney of Technology Practice Group LLC (www.technologypracticegroup.com), a law firm that regularly advises its Fortune 100 clients in connection with their technology and IP transactions. LinkedIn: https://www.linkedin.com/in/jeff-monassebian-0a0b67/). Gary S. Greenstein, Esq., Kevin Davis, Esq. and Kelly Monassebian, Esq. of Technology Practice Group assisted with this Practice Note.

[2] Terminology is listed in "waterfall" order as opposed to alphabetical order.

- *Models*: Similar to Machine Learning, Models are also a subset, of Artificial Intelligence. A Model is a mathematical or algorithmic formulation that has been trained with Training Data and informed through Machine Learning to generate Output and/or recommend actions in response to user requests, instructions and queries.

- *Prompt*: Before a Model can provide responsive Output, the user must pose a request, instruction or query. In its simplest form, a Prompt is the request, instruction or query to which the Model will deliver responsive Output. Additionally, a Prompt can include 'context', such as a repository of documents or library of systems/code, which can be referenced in a Prompt to improve a Model's Output. The adage "garbage-in-garbage out" applies when crafting a Prompt. The more precise a Prompt is crafted, the higher the likelihood that responsive Output will satisfy the user's request, instruction or query.

- *Input*: Prompts and, if provided by licensee, Production Data and Training Data.

- *Output*: Data, images, information and other results generated from execution of a Model or other AI System in response to a Prompt or other instructions.

- *Fine Tuning*: Fine Tuning trains a Model with Training Data specific to the intended use, user's industry and/or applicable industry vernacular such that responsive Output is more relevant, accurate and reliable. Models are usually trained/Fine Tuned on an iterative basis to improve their performance. One example of Fine-Tuning is *Harvey.ai*, which allows Fine-Tuning to a law firm's unique requirements.[3]

- *Fine Tuned Models*: A Fine Tuned Model is a Model that is Fine Tuned by Training Data and/or feedback.

License Agreement – Key Terms:

The current legal framework that governs AI Systems, including their creation, training, and Output, applies existing laws and principles to technology not contemplated when such laws and principles were developed. As such, existing laws may not adequately address, let alone answer, the novel legal issues that will invariably arise in the context of AI System licensing agreements. While legislative activity is underway to specifically address AI, the process is still nascent.[4] Therefore, counsel need to ensure that key contractual terms memorialize the parties' understanding and intent, because it will take some time for the law to "catch up" to the technology.

- *Ownership*: Simply said, "it's complicated." There are many moving parts when negotiating the respective ownership interests arising out of Artificial Intelligence.[5]

[3] Harvey.ai uses the GPT-3 technology (*not* ChatGPT) to enable lawyers to create legal documents or perform legal research by providing simple instructions using natural language.
[4] See, for example, the proposed EU AI Act (https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence) and "*A Pro-innovation approach to AI regulation*" (https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper).
[5] Discussion and sample agreement provisions with respect to ownership rights are limited to Models.

- *Models*: Development of Models, and in particular, large language Models utilizing deep learning,[6] is a major undertaking, and costs can run into the hundreds of millions of dollars. Beyond capital requirements, expertise to train a Model requires highly-skilled engineers and computer scientists. Therefore, licensors will zealously protect and preserve their IP rights in a Model. This paradigm is similar to the majority of software agreements.

- *Fine Tuned Models*: No matter the extent to which a Model is Fine Tuned, it is unlikely that a licensor will pass ownership in a Fine Tuned Model to licensee. Licensors defend this outcome because the significant investment in development is exponentially higher than contributions made by any one licensee. Moreover, licensors will further rationalize their position by suggesting all licensees benefit from use of Fine Tuned Models because the Output is more accurate and precise. Therefore, it is important for counsel to recognize and understand how to protect licensee's rights in its Training Data that is shared with licensor, while at the same time preserving the benefit to licensee in using a Model that is Fine Tuned with Training Data licensor receives from others.

    - *Licensee Provided Production and Training Data*: In order for a licensee to benefit from Models Fine Tuned with Training Data provided by others, licensor will request a license to use licensee's Training Data to continue Fine Tuning the Model, i.e., "*give-to-get.*"[7] Licensee's ownership and commercial interests in its Training Data should be weighed against the benefit licensee receives from using a Model Fine Tuned with Training Data provided by others. Balance between these competing factors may be achieved by requiring licensor to anonymize licensee's Training Data.[8] A sample provision follows:

        - *"Licensor may use Licensee's Training Data to create Anonymized Data and use such Anonymized Data solely for purposes of improving the accuracy of Licensor's Model. As used in this Clause, "Anonymized Data" means aggregated information prepared or produced by Licensor from Training Data provided to Licensor by Licensee and Licensor's other licensees, provided, however, that in all cases it is not possible to identify Licensee or any of its personnel or other users or any of its or their respective behavior. Without limiting the foregoing, to constitute Anonymized Data, Licensee Training Data must: (i) be aggregated with Training Data of at least [**] other similarly situated licensees of Licensor; (ii) not comprise more than [**] percent of the aggregated information; and (iii) not include any Personal Information."*

- *Prompts*: Prompts are usually crafted in an iterative manner until Output is responsive to the user's request, instruction and/or query. Crafting a Prompt can be time consuming and labor intensive. Therefore, as between licensor and license, Prompts should be owned by licensee. Defensible positions for licensee's ownership can be put forward in: (i) copyright if the Prompt is: (a) original and created

[6] Deep learning Models have been trained to recognize complex patterns in pictures, text, sounds and other datasets to produce responsive Output.
[7] In cases where licensee is installing a Model on-premises behind its own firewall or in a partitioned segment of the cloud-hosted environment, the "give-to-get" paradigm may not provide sufficient benefit to licensee. In this case, licensee will be training the Model with its Training Data. Training Data provided by other licensees may not be meaningful to licensee's use of the Model.
[8] Any decision to allow licensor to use licensee's Training Data must be carefully considered. For example, if licensee's Training Data is confidential or provides licensee with competitive advantage, it is unlikely that any amount of anonymization will provide sufficient protection to licensee.


by a human, (b) creative, and (c) fixed in a tangible medium of expression; and (ii) trade secret if: (y) the Prompt can be regarded as a formula, compilation, method, or process; and (z) the Prompt's secrecy is maintained so that it derives independent economic value. Regardless of the approach used to protect licensee's intellectual property right interest, the license agreement should expressly state that Prompts are, as between the parties, owned by licensee and constitute licensee's confidential information. A sample provision follows:

- *"Licensor acknowledges that Prompts are prepared through the application of methods and standards of judgment used and developed through the expenditure of considerable work, time and money by Licensee. Licensor also acknowledges that, as between the parties, Prompts are the exclusive property of Licensee and constitute confidential information and trade secrets of Licensee."*

- *Input and Output*:

Input:[9] Input includes Training Data provided by licensee and Prompts. Ownership with respect to Prompts was discussed earlier in this Practice Note. Ownership by licensee of its Training Data can be supported in: (i) copyright if the Training Data is: (a) original and created by a human, (b) creative, and (c) fixed in a tangible medium of expression; and (ii) trade secret if: (y) the Training Data is a compilation, method, or process; and (z) the Training Data's secrecy is maintained so that it derives independent economic value.

Output: Copyright and trade secret law[10] may support the conclusion that Output is owned by licensee.

- *Copyright*: Copyright protection for Output will depend on whether the Output (i) was created entirely from the application of Prompts or autonomously by the Model, or (ii) resulted from some creative input by the human user (e.g., Output selected and arranged by a human). If Output was generated entirely by application of Prompts or autonomously, the current policy of the US Copyright Office is that copyright protection is not available.[11] The extent of creative human input required to afford copyright protection will be factually specific in each instance.

- *Trade Secret*: Output may constitute a formula, pattern, compilation, method, or process. To the extent such Output derives independent economic value and its secrecy is maintained, trade secret protection may apply. Where a Model is executed on cloud-hosted servers, particular care should be given to ensure Output is stored in a partitioned segment of the server or graphic processing unit with access limited to the licensee in order to preserve secrecy.

[9] While ownership of Training Data in this paragraph discusses licensee rights, the same analysis and outcome is relevant for the protection of a licensor when the licensor is providing Training Data.
[10] Output is not protectable under current patent law in the United States because the "inventor" is not a human. This result may change in the future as the United States Patent and Trademark Office considers patentability in cases where there the "invention" results from the combination of AI and human ingenuity.
[11] See: "*Statement of Policy*", issued by the Library of Congress dated March 16, 2023: https://www.federalregister.gov/documents/2023/03/16/2023-05321/copyright-registration-guidance-works-containing-material-generated-by-artificial-intelligence#:~:text=IV.-,Guidance%20for%20Copyright%20Applicants,author's%20contributions%20to%20the%20work.


Because existing laws may not definitively resolve ownership rights, the license agreement should memorialize the ownership paradigm as it relates to Input and Output. A sample provision follows:

- *"As between the parties, Licensee has exclusive title and ownership rights, including all Intellectual Property Rights, throughout the world, in all Licensee Data.[12] To the extent that such rights may not originally vest in Licensee, Licensor hereby irrevocably assigns, and shall cause all relevant Licensor personnel irrevocably to assign, to Licensee (or its designee) all such rights in the Licensee Data. Licensor shall not: (i) dispose of, distribute or otherwise use or exploit the Licensee Data, for any purpose other than for the sole and exclusive benefit of Licensee or to create Anonymized Data; and (ii) assert any lien or other right over any Licensee Data."*

- *Confidentiality*: As discussed earlier, intellectual property law may afford ownership rights to licensee with respect to licensee data (including Prompts and Output). However, as an additional layer[13] of protection for licensee, the license agreement should contain an obligation for licensor to treat and maintain licensee data as confidential. A sample provision follows:

  - *"As used herein "Confidential Information" means all confidential or proprietary information related to the business of the other party to which a party has access, acquires or otherwise processes, whether in oral, written or other form, in the course of or in connection with this Agreement, together with all copies of, and all materials incorporating, any such information. Confidential Information of Licensee includes Licensee Data (including Prompts and Output) and Fine Tuned Models.[14] Except for Personal Information (which is always deemed Confidential Information), Confidential Information does not include information that: (i) was in the possession of or demonstrably known by the receiving party prior to its receipt by the receiving party without restriction on its use or disclosure; (ii) is independently developed by the receiving party without use of, reference to or reliance on the other party's Confidential Information; or (iii) becomes known by the receiving party from a source apart from the other party without breach of this Agreement and is not subject to an obligation of confidentiality."*

  - *"Without limiting or modifying Section [*] of this Agreement ('Ownership'), the receiving party shall keep all Confidential Information secure and strictly confidential using procedures no less rigorous than those used to protect and preserve the confidentiality of its own similar confidential and/or proprietary information but in no event less than a reasonable standard of care given the nature of the Confidential Information. Notwithstanding the foregoing, Licensee may disclose Confidential Information of Licensor to the extent required pursuant to applicable*

[12] "*Licensee Data*" should be defined *as collectively: (i) Inputs; (ii) Usage Data; and (iii) Output. For the avoidance of doubt, Licensee Data shall not include any of Licensor's commercially available models and their respective model weights, or model architecture. As used herein, "Usage Data" means all usage data (i.e., telemetry data, metadata, and similar data) collected, generated or derived by Licensor from use of model(s) by or on behalf of Licensee. Usage Data does not include any such data that is aggregated and anonymized by Licensor such that it is not possible to identify Licensee or its personnel or any of its or their respective behavior.*

[13] Confidentiality clauses usually contain exceptions to the confidentiality obligation. Therefore the clause should be drafted in a manner such that the confidentiality obligation is supplemental to and does not diminish or modify the ownership provisions of the license agreement.

[14] If the Fine Tuned Model is licensed for the exclusive benefit of licensee.

law, court order, to satisfy a request by any regulator or in connection with any review of its model risk management activities."[15]

- *Personal Information*: Personal Information will likely be included in datasets against which Text and Data Mining is performed[16] and therefore may conflict with a data subject's right to: (i) have notice of who uses their personal information, how it is used, and the purpose for such use; and (ii) give consent to such use.[17] The practical difficulties of giving notice and obtaining consent from each data subject whose personal information may be processed, or establishing a lawful basis to process the data, exposes both licensor and licensee to claims by data subjects that their privacy rights have been violated. Moreover, anonymization of personal information may not provide a work-around because the rules for anonymizing personal information are unsettled.[18] In the event a data subject brings a claim against a licensee that development or use of the AI System violates such data subject's privacy rights, the license agreement should include a defense and indemnity obligation requiring licensor to defend licensee and hold licensee harmless from any liability, unless there is a factual basis for an exemption to apply.[19] A sample indemnity provision is provided in the Indemnities section later in this Practice Note.

- *Indemnities*: Given the current legal framework governing AI Systems applies existing laws and principles to technology not contemplated when such laws and principles were developed, claims can arise beyond conventional allegations that an AI System infringes, thereby resulting in unanticipated liability. For example, Text and Data Mining may result in infringement of copyrighted material[20] and, as noted earlier, violate the privacy rights of data subjects. Claims may also arise if licensee does not have adequate rights to provide Training and/or Production Data, or when Prompts crafted by licensee return infringing Output. The license agreement should equitably balance each party's defense and indemnity obligations taking into account the following considerations:

  - ❖ Which party is best positioned to defend the claim? AI Systems are developed by licensor and therefore licensor possesses the factual basis to defend allegations of infringement. Conversely, if licensee provides the Training or Production Data, licensee possesses the factual basis to defend claims alleging provision or use of such data infringes third party rights.

---

[15] With respect to model risk management, see Risk Management and Disclosure section later in this Practice Note.

[16] To the extent licensor (and not licensee) is providing the data sets against which TDM is performed, it will likely be deemed the "controller" with respect to any personal information included in such data and is therefore responsible for compliance with applicable data privacy laws.

[17] In the United States, notice to and consent from the data subject is required. EU and UK data protection laws may also allow processing of personal data if there is a lawful basis to process the data.

[18] The following is from IAPP Publication dated June 27, 2023: "*The definition of anonymization is changing in the EU: Here's what it means*" (https://iapp.org/news/a/the-definition-of-anonymization-is-changing-in-the-eu-heres-what-that-means/).

> *At a high level, a risk-based approach to anonymization allows for the residual risk that the data still could theoretically be identified in the future — the lower the risk, the stronger claims to anonymization can be. This* ***risk-based approach*** *is commonly applied in several jurisdictions and has been a central tenet of anonymization standards in the U.S. The Federal Trade Commission, for example,* ***promoted this standard in 2012, shaping state-level privacy laws around the U.S. ever since. A risk-based approach typically entails maintaining tight control over the way the data is reused, which is why closed data environments, with monitoring and auditing capabilities, are so important****. … Meanwhile, the European Data Protection Supervisor* ***focused on the criterion of irreversibility*** *— meaning techniques used to anonymize* ***can never be reversed*** *— in a way that also appears to be at odds with the Article 29 Working Party's initial risk-based approach. To make matters even more confusing, those in the technical community who have tried to formalize the EU concept of anonymity* ***have adopted the most restrictive approach*** *to anonymization, often oversimplifying legal tests for anonymization and driving the adoption of the most conservative methods possible.*

[19] From a practical standpoint, licensor should defend the claim because licensor created the AI System and has the records to show what and how data was used, but see the exemption in the sample provision provided later in this Practice Note.

[20] To the extent a "copy" needs to be made and/or stored during the course of TDM, such activity may constitute copyright infringement of the target data-sets.

- Standard exceptions to a licensor's indemnity obligation are not necessarily applicable or appropriate in connection with AI System infringement. For example, an exception that excuses licensor from defending an infringement claim to the extent the claim arises from a combination of the AI System with products and services provided by third parties may not always apply. AI Systems will access third party systems and ingest and analyze data from such systems – arguably a "combination", unless the AI System is deployed and executed on-premises, trained with Training Data and executed against Production Data provided only by licensee.

- License agreements usually impose broad defense and indemnity obligations on licensees to defend and indemnify licensor from claims that arise from licensee's use of the AI System. Such an obligation is too broad and should be avoided. While licensees submit Prompts or instructions, Output is largely a function of how the AI System has been designed, developed and trained by licensor. Therefore any indemnity given by licensee relating to use should be limited to an allegation that licensee has used the AI System to craft a Prompt or instruction with the knowledge or intention to generate infringing Output.

- Licensors should defend and indemnify licensees against claims by a data subject that its privacy rights were violated in connection with development, training or use of an AI System, except to the extent violation occurs from licensee provided Training and/or Production Data.

A sample provision follows:

> *"Licensor's Defense and Indemnification. Licensor shall defend and hold harmless Licensee from and against all third party claims, actions and demands, and shall indemnify Licensee against all Losses[21] suffered or incurred, in each case arising out of or relating to any third party allegation or determination that the AI System or its documentation, or any portion of them, or Licensee's use thereof, infringes or misappropriates any Intellectual Property[22] or other proprietary right of any third party. Licensor shall have no obligation to defend Licensee under this Section with respect to any third party claim that the AI System or its documentation, or Licensee's use thereof, constitutes an infringement or misappropriation of any Intellectual Property Rights of any third party, solely to the extent such infringement or misappropriation arises from: (a) modifications to the program code underlying the AI System made by, or by a third party on behalf of, Licensee (other than modifications made by or on behalf of Licensor), except where such modifications are agreed to by Licensor in writing or are necessary for the reasonable use of the AI System; or (b) Training Data or Production Data in the form provided by Licensee."*

[21] "Losses" should be defined as: *collectively, all losses, liabilities, damages, awards, costs, expenses (including lawyers' and other professional advisers' and experts' fees) and amounts paid in settlement or compromise.*

[22] "Intellectual Property" should be defined as: *all: (i) patents, patent applications, patent disclosures and inventions (whether patentable or not); (ii) trademarks, service marks, trade dress, trade names, logos, corporate names, Internet domain names, and registrations and applications for the registration thereof, together with all associated goodwill; (iii) copyrights and copyrightable works (including computer programs and mask works) and registrations and applications thereof; (iv) design rights and database rights; (v) trade secrets, know-how and other confidential information; (vi) waivable or assignable rights of publicity, waivable or assignable moral rights; and (vii) all other forms of intellectual property.*

- *"Licensee's Defense and Indemnification. Licensee shall defend and hold harmless Licensor from and against all third party claims, actions and demands, and shall indemnify Licensor against all Losses suffered or incurred, in each case arising out of or relating to any third party allegation or determination that: (a) Training or Production Data, in the form provided by Licensee, infringes the Intellectual Property Rights of a third party; or (b) Licensee has submitted a Prompt or other instruction (in the form submitted by Licensee) with the knowledge or intent that the Output generated from such Prompt or instruction will infringe the Intellectual Property Rights of a third party. Licensee shall have no obligation to Licensor under this Section for claims that Training or Production Data (in the form provided by Licensee) constitute an infringement or misappropriation of Intellectual Property, when such data was used by Licensor in a manner not permitted under the License Agreement.*

- *"Indemnity for Violation of Privacy Rights. Licensor shall indemnify Licensee against all Losses suffered or incurred in connection with all claims, actions, investigations and demands, in each case to the extent arising out of or relating to an allegation: (a) by a data subject that data subject's privacy rights have been violated in connection with the development or permitted use of the AI System; or (b) by a governmental or regulatory entity that the AI System, its development or its permitted use violates applicable privacy and data protection laws or regulations. Licensor shall have no obligation to defend Licensee under this Section solely to the extent such violation arises from the use by Licensor of Licensee provided Training and/or Production Data (in the form provided by Licensee) in a manner permitted by this License Agreement."*

- *Security*: The license agreement should require licensor to implement and maintain security procedures to buttress the ownership rights of licensee and confidentiality obligations of licensor with respect to licensee data.[23]

  Security procedures should include:

  - minimum security requirements (including encryption of licensee data in transit and at rest). Each licensee organization should have established minimum security requirements for cloud-hosted services. These requirements should be shared with licensor, and licensee should request licensor's confirmation that licensor's security procedures satisfy licensee's minimum security requirements. The purpose of this process is not to require licensor to implement bespoke security procedures, but for security gaps to be identified, thereby enabling licensee to assess the applicable risk;
  - configuration settings that enable licensee authorized personnel to determine if, and for how long, licensee data is stored on the cloud-hosted environment;
  - access restrictions so that configuration and security settings can only be modified by named individuals designated by licensee; and

[23] Security is most relevant where the AI System is provided in a cloud-hosted environment or in a hybrid environment. In the hybrid environment: (i) Prompts and instructions are created on licensee's internal systems and then transmitted via API to the AI System for execution in the cloud-hosted environment and (ii) Output is generated in the cloud-hosted environment. Only licensee should have the ability to set system parameters defining whether (a) Prompts/instructions remain in the cloud-hosted environment; and (b) if Output should be transmitted back licensee's internal systems, deleted or maintained on the cloud-hosted environment.


- requirements as to whether the AI System must only be executed and/or licensee data processed/stored in a partitioned segment of the server or graphic processing unit with access limited to licensee (a "Partitioned Segment").

A sample provision follows:

- *"Licensor shall, throughout the term of this Agreement, maintain and comply with a written information security program of administrative, technical and physical safeguards that are appropriate for Licensor's size and complexity, the nature and scope of Licensor's activities, and the sensitivity of the data that Licensor is handling. Without limiting the foregoing, Licensor shall ensure that such security program includes: (i) the minimum security controls (set out in Exhibit [*]); (ii) controls to secure and protect Licensee Data; (iii) controls that limit permission to implement and manage configuration and security settings to the named individuals identified in Exhibit [*]; and (iv) practices to detect, report and resolve security vulnerabilities and threats as quickly as possible."*

- *"The AI System shall process Licensee Data only in the Partitioned Segment. Prompts, Instructions and Output may only be retained in temporary cache of the Partitioned Segment and only for the minimum period of time necessary to process such Prompts and Instructions, and to generate Output. Licensor may temporarily retain Training Data provided by Licensee, but only for the minimum period of time necessary to train the AI System, and only if such retained Training Data remains only in the Partitioned Segment."*

- *Liability*: License agreements typically (i) limit liability to an amount equal to fees paid or payable for the preceding 12 month period, and (ii) disclaim all liability for consequential, special and indirect damages. Such limitations and exclusions should be avoided when it comes to AI Systems. They essentially render confidentiality, security, privacy, compliance with law and indemnity obligations illusory as there would be no meaningful remedy in damages for their violation. Damages sustained from the violation of confidentiality, security or privacy obligations will substantially be in the form of consequential, special and indirect damages (e.g., loss of reputation, business opportunity, clients and profits). Moreover, fines for violation of law and the costs and liability relating to IP defense and indemnity obligations can amount to tens of millions of dollars. Therefore, liability clauses should include applicable carve-outs from both the limits and exclusions of liability. A sample provision follows:

  - *"NO LIMITATION OR EXCLUSION OF LIABILITY SHALL APPLY WITH RESPECT TO ANY CLAIM ARISING OUT OF, RELATING TO OR ARISING UNDER: (i) CONFIDENTIALITY OR PRIVACY OBLIGATIONS, COMPLIANCE WITH SECURITY PROCEDURES, OR DEFENSE AND INDEMNIFICATION OBLIGATIONS; (ii) WILLFUL MISCONDUCT OR GROSS NEGLIGENCE;[24] OR (iii) ANY FINE OR PENALTY IMPOSED ON A PARTY FOR FAILURE TO COMPLY WITH APPLICABLE LAW."*

[24] In most jurisdictions, a party cannot limit its damages for its willful misconduct or gross negligence.

- *Risk Management and Disclosure*: As the use of AI Systems become more pervasive, an organization's actions will frequently be informed by Output and other information generated by the AI System. When health, safety, financial, privacy or other rights are affected by such actions, regulatory requirements and/or licensee's own internal compliance procedures may require transparency as to the AI System's structure, algorithms, and processes. Therefore, the license agreement should require licensor's cooperation in providing such information. Moreover, the license agreement should require licensor to monitor the AI System to verify it is performing as intended and that Output and similar information is not affected by bias. A sample provision follows:

    - "*Licensee may request detailed information in relation to the AI System (including its structure, algorithms, and processes) and Licensor shall provide reasonable co-operation and assistance to Licensee in respect of the provision of information referred to herein.*"

    - "*On at least a semi-annual basis, Licensor shall perform and make available to Licensee the results of performance monitoring to demonstrate that the AI System is performing as intended, including that Output is generated without bias.*"

- *Terms of Use*: Licensor will typically include "Terms of Use" to proscribe inappropriate or unsafe use of an AI System. Examples of proscribed use include Prompts or instructions that generate Output that is: (i) illegal, (ii) child sexual abuse material, (iii) hateful, harassing, or violent, (iv) fraudulent or deceptive, or (v) a privacy violation. As discussed earlier, Output generated is largely a function of how the AI System has been designed, developed and trained by licensor. Therefore any such proscription should be limited to the case where licensee has acted with the knowledge or intention to create inappropriate or unsafe Output.

Conclusion:

AI Systems are at the "bleeding edge" of technology, and new capabilities are being developed rapidly.[25] Existing legal frameworks struggle to provide clear and definitive answers to the legal and business issues arising from the use of AI Systems, and will likely do so for the foreseeable future. Eventually, AI Systems may autonomously identify and assess each parties' respective legal and business obligations, risks and requirements, and then generate the appropriate license agreement.[26] However, an attorney or business person who solely relies (or overly relies) on recommendations or conclusions generated by artificial intelligence does so at their own risk. AI Systems will continue to advance, but it remains to be seen if the day will come when AI can truly replace human judgement, discretion and experience. Until that day, diligence and careful consideration of the issues, including the issues explored in this Practice Note, must be an on-going exercise.

[25] For example, a new feature in some large language Models trigger action, e.g., an instruction within a Prompt to send a request or query to a public or private application. Use of this feature may raise privacy and related "data controller" issues.
[26] While various vendors currently provide AI tools for contract generation, it is foreseeable that the entire negotiation and contracting process, from inception to signature, can be performed using an AI System.