Considerations in procuring vendor AI

 Danielle Colliver

 Jessica R. Lumière

 Angela Jiao

AI is taking the world by storm, bringing with it the promise of endless possibilities for improving business efficiencies. While cautiously optimistic about the opportunities posed by AI, today's companies are struggling to strike the right balance between leveraging its benefits and effectively managing its associated risks. This article sets out key considerations for organizations that are contracting with third parties for AI technologies.

Knowledge first

From the outset of any contracting exercise, it is vital to have a good understanding of how AI works—both the third-party product itself and how it would interact with your organization's other tools. You don't need to be an expert, but you need to know enough to ask the right questions and assess the answers. Before engaging with an AI vendor, read up on AI, attend information or training sessions, and gain a basic level of understanding of what it is and how it works. Most vendors are happy to set up demos with their technical team to help you explore their product. In this instance, ask questions but be mindful of sales pitches that may not be backed by real performance. The goal is to be an informed user, not an expert.

Then due diligence

Before putting pen to paper on a contract for an AI product, it is important to understand key aspects of the specific technology, including:

- **What is the business use case?** Understand how your organization intends to use the technology, including the risks associated with its functionality and data exposure. If your organization is new to AI, the initial use case should be a discrete and manageable task. Consider:
  - limiting the variety of data points and the volume of records;
  - having a short-term, measurable ROI;
  - aligning your use case with your organization's overall AI strategy, data governance and risk management policies;
  - avoiding processes that underpin regulatory obligations or touch other critical or core processes, decisions or systems until you better understand the AI product and have tangible results on its validity and reliability; and
  - focusing on the desired business outcome rather than on the tool's superfluous features.

- **What are you purchasing and how does it work?** Although it seems like a simple question, the inner workings of AI tools are often multi-layered and complex, which can increase the risk of unintended and inexplicable consequences. Your technology team should be involved in reviewing the tool to ensure you understand its algorithms, training, and data processing methodologies to a level that is appropriate for the particular use case. Focus on obtaining as much information as possible about the third party's relevant policies, processes and developments, and avoid focusing on details that are not critical to your use case.

- **What are the limitations and risks?** Identify any limits of the technology, including its actual features and functionality, potential biases, training and data source limitations. Run through worst-case scenarios and ensure risks and limits are appropriate to your use case or can otherwise be mitigated by users or other operational mitigants.

- **What are the data security capabilities of the vendor and the AI technology?** Understand how the vendor manages and secures data, including cyber security and storage practices, locations, and measures.

- **Is there a generative AI component to the technology?** Understand the impact of hallucinations and consider whether results need to be validated by human users. Identify potential privacy, security and intellectual property risks, and investigate whether there are options to use private models with ring-fenced data sources or watermarking outputs for transparency—a useful practice whether or not emerging regulations requiring AI transparency and watermarking, such as the *California AI Transparency Act,* apply to you.

- **Who is the vendor?** Consider whether you have an established relationship with the proposed vendor, as well as your current level of insight into their relevant products and operations. If you don't have an established relationship with them, consider conducting diligence to confirm their expertise in this technology as it relates to your specific use case and ask for references from their existing customers. In addition, look out for agreements between different vendors for different segments of your purchase. For example, one entity may be providing or implementing the technology, but another entity may be servicing the technology over the life of the contract. Depending on the structure, this may present additional risks.

- **How have other customers experienced the technology?** Consider how long the technology has been on the market, whether bugs are still being ironed out, whether your data will be used to improve the service or product for other customers, and whether the tool has been used extensively by other customers. You may wish to further explore prior customers' experiences and speak with them if possible.

The diligence phase should conclude with an initial risk-based decision as to whether to proceed with contracting for the AI tool. The key consideration will be whether any risks identified through due diligence are acceptable for the given use case. This decision, and the rationale behind it, should be clearly documented and revisited throughout the contract and use case lifecycle as new risks are identified.

Contractual risk mitigation strategies (when available)

The ability to negotiate a contract for AI technology may only exist for custom solutions or high-value agreements. In many cases, particularly with commercially available products, there may be no ability to negotiate, or negotiations may be limited to material concerns. However, when the opportunity to negotiate presents itself the goal should be to negotiate a contract that (a) is appropriately tailored to your use case; and (b) mitigates risks identified during the due diligence phase. If there is no opportunity to negotiate, the use case should be tailored based on the contractual limitations and the risks identified during the diligence phase.

In addition to terms and conditions that would be applicable to any technology contract, the following are some of the key terms and conditions that should be explicitly considered in the context of a contract for AI technology:

- ***Ownership and use rights related to outputs and inputs***
  - *Outputs* – Ensure your organization is granted the rights it needs to use the outputs of the AI technology for each intended use case. Consider what pre-existing ownership rights may exist in, or overlap with, the outputs. This is important depending on how you intend to use the outputs and whether you have the necessary rights to do so.
  - *Inputs* – Ensure that the vendor's ability to use data and information your organization provides as inputs into the AI model is appropriately limited based on the sensitivity of the data. Consider the purpose for which the vendor is using your data and the method of their use (such as in anonymized and aggregated form). For example, AI technology vendors often push to have broad rights to use AI input data to further train their models. Whether this is appropriate will depend on the sensitivity of the data in question, whether your organization has obtained appropriate rights for such use of the data (e.g., consents), and whether your organization obtains any benefit from the vendor's use of your data (e.g., through improved outputs or some form of compensation).
    - If the vendor will be using third party data as an input, ensure that they have the appropriate rights to use this data and that your organization is protected from any claims as a result of using the AI tool that is leveraging this data (see indemnities below).
- ***Warranties (performance and quality)*** – Focus on obtaining strong outcome-based warranties where the vendor guarantees that the AI technology performs in accordance with certain criteria, such as the production of expected outputs or results.
- ***Indemnities*** – Given the current uncertainty in case law around intellectual property rights relating to AI inputs and outputs, the contract should include intellectual property indemnities that shift the risk of third-party intellectual property claims to the vendor. It is also important to consider whether the risk of incorrect outputs or outcomes should be shared with the vendor, particularly if your organization is an early adopter, if the outputs are based on your ring-fenced data, or if there is additional financial or regulatory risk to your organization if the AI doesn't work or works incorrectly (e.g., for automated decision-making tools). Note that indemnities remain rare in this space and, where they may be offered by larger providers (e.g., Microsoft and Google), carefully consider any exclusions and the organization's obligation to comply with any conditions.
- ***Disclaimers*** – It is common for AI technology vendors to disclaim any liability because of the client organization's reliance on the validity of outputs or subsequent decisions based on the model's output. Whilst determining whether such a disclaimer is acceptable will be based on your use case, consider factors pertaining to the AI's operation, such as whether there is an element of human review, the level of custom data used to train the model, or the custom functioning of the system. Where such disclaimers are considered acceptable, they should be narrowly drafted to avoid situations where they are broad enough to extend to disclaim liability for outputs generated in violation of the agreed functioning of the AI technology. For example, if an AI decision-making tool was given a rule that X data results in Y decision but the tool produced Z decision instead of Y (despite being given X data), then this is an issue resulting from a technological malfunction, not an issue stemming from improper reliance on the tool.
- ***Transparency, audit and governance*** – The contract terms should, wherever possible, allow your organization to monitor the AI technology in order to ensure it performs as intended and adheres to the agreed-upon terms and industry practice. Where your organization has significant leverage, the AI tool is particularly critical, or the vendor is providing AI development services, these terms might include the right to review the modeling approach and techniques used by the AI tool, the training data, the testing done by the vendor, the explainability and transparency analysis, and the vendor's policies and procedures (including around making and testing changes to the tool, security, and tracking incidents and user complaints).

- ***Compliance with laws*** – The contract should not only include an obligation for the vendor and the AI technology to comply with laws and regulations applicable to AI technologies (for example to avoid prohibited uses under the Regulation (EU) 2024/1689 (the EU's *Artificial Intelligence Act*)), but also require the vendor to facilitate your organization's compliance obligations, as applicable. For example, NYC Local Law 144 applies to an employer's use of decision-making tools for employment practices in New York; however, the vendor will need to supply specific information to ensure the employer-organization can comply. While AI regulation continues to develop, consider including mechanisms for ensuring the contract can be updated to reflect changes in laws, regulations, and policies.

- ***Remedies*** – Carefully consider the circumstances under which your organization should be able to exit the contract, get its money back, be compensated for losses it may suffer, or rely on some combination of these remedies. Be cautious when considering liability limitations and disclaimers that may carve away at your remedies.

Other risk mitigation strategies

In addition to contractual risk mitigants noted above that may be available to you, the following mitigants can help manage other risks related to the procurement of AI technologies:

- ***Organizational alignment –*** Ensure that:

  - the standard contractual terms are well understood by each of the relevant stakeholders in your organization (or those potential future use cases);

  - if there are terms or conditions that exceed your organization's risk tolerance for the current use case, ensure the use case is narrowed accordingly; and

  - if there are certain terms and conditions that exceed your organization's risk tolerance when certain criteria are met or for certain types of future use cases, ensure there is an internal mechanism for ensuring the AI tool is not used in the event those criteria or use cases materialize.

- ***Vendor management*** – Use the contractual rights at your disposal, including governance and audit rights and the vendor's performance obligations, to oversee the vendor and their use and development of the AI tool, and to evaluate the tool's performance on a regular basis. Also, ensure the organization is aware of any ways to manage the risks of using AI tools to operate and perform services where the AI component may not be as obvious (e.g., Microsoft Copilot contains AI components that users are often not aware of).

- ***Project management*** – Oversee the project internally to ensure compliance with the contract terms and conditions and prepare a project plan to avoid scope creep that might expand the risk of using the AI tool beyond your organization's risk tolerance. Where appropriate, consider human oversight of any AI tools, particularly those used for automated decision-making or critical functions.

- ***Knowledge management and change management*** – Ensure that AI and its associated benefits and risks continue to be a primary focus for employee training, including ethical use, data privacy, security, and compliance with regulatory obligations and your organization's policies. Carefully manage any changes to the tool's functionality or use, as well as to the regulatory landscape.

- ***Internal governance and controls –*** Employ technical controls on your side of the demarcation (e.g., access controls, limiting data that can be inputted, etc.). In addition, ensure that your organization has comprehensive and clear data governance policies and practices, as well as an AI policy or other organizational policies or guidelines governing how and when AI tools can be used. Ensure there are controls in place to monitor compliance, and that that the organization's board is well versed in AI-related risks and policies.

Conclusion

Whether your goal is to lessen administrative burdens, automate processes, hasten decisions, or simply avoid having to take notes in a meeting, AI can offer many solutions. Evaluating your planned use of this technology in a systematic and informed manner and ensuring the right controls—be they contractual or otherwise—are in place will allow you to take advantage of these efficiencies while managing their associated risks.

To discuss these issues, please contact the author(s).

More from our AI resource hub

AI RESOURCE

What should be included in my organization's AI policy?: A data governance checklist

Organizations will need to adhere to a clear data governance framework in order to wield AI responsibly and ethically.

AI RESOURCE

Rules for AI tools: how can legal teams source suitable tech?

Legal professional can avoid wasting resources on AI by focusing on specific business use cases before implementation. We share how.

AI RESOURCE

AI for employers: balancing risk and reward

AI has a role to play in streamlining the employee lifecycle. We provide insights into the evolving legal landscape surrounding the use of AI for employee recruitment and retention.