Bristows

Generative AI

Key issues when integrating AI into your products and services.

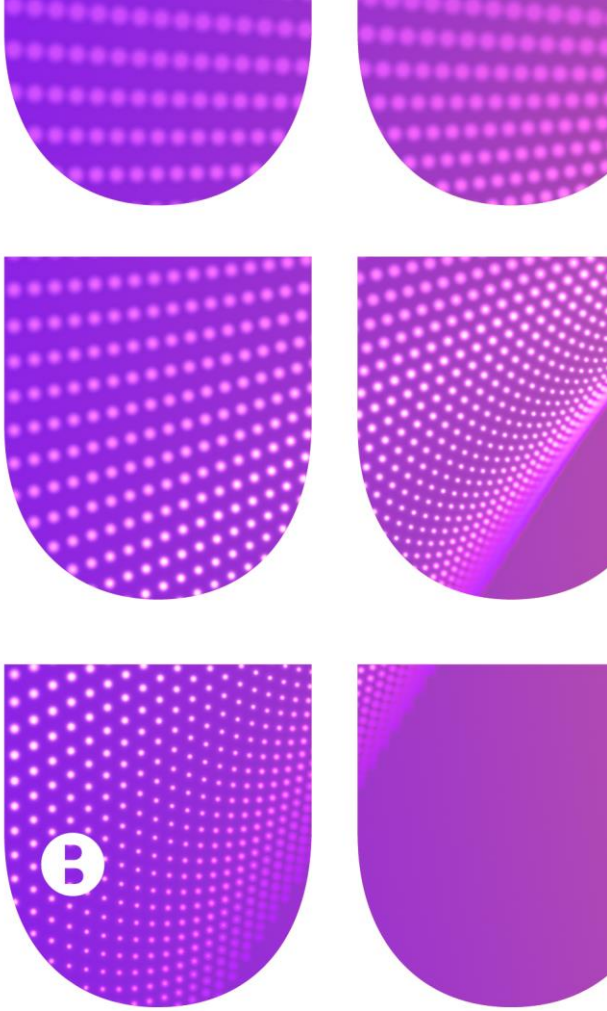

One way organisations can leverage the generative AI wave is to incorporate it into their own (non-AI) products and services.

Imagine an advertising agency connecting its content creation program to a large language model (LLM), so creatives can produce tailored campaigns at speed and scale. Or a medical imaging company enhancing its outputs using a model trained on a huge database of ultrasound and x-ray images. Many businesses will be looking to buy paid subscriptions to generative models and integrating them (through APIs) into their own products and services, so that they become "powered by AI".

This can be achieved with one of the major AI companies (like OpenAI) or a sector-specific developer (like Huma.AI in life sciences). In any case, legal teams should consider the issues involved when contracting for generative AI, which in most cases will be on standard terms and conditions. Below, we set out some key issues to look out for and guidance on the **due diligence** and **internal mitigation** that can ensure the AI project meets its objectives.

1. How will your confidential information be used?

Generative AI tools are often fine-tuned on input data to make their outputs more reliable and accurate. While some AI contracts are clear that ownership of all prompt and input data will remain owned by its customer, many vendors ask for a licence to use a customer's input prompts and images to improve their services. It's therefore important to consider the type of prompt and other input data you will upload. The risks include: the upload could be considered a disclosure that may put you in breach of confidentiality obligations to your third parties; the protection offered by confidentiality or trade secrets could be lost or weakened; and your information could find its way into the public domain.

Due diligence:

- Engage with your business/product teams to understand which datasets will be exposed to the AI tool (e.g. types of prompts).
- Check whether confidential information or personal data is included in those datasets (if so, see the internal mitigation below).

Internal mitigation:

- Investigate whether your vendor offers an option to opt out your data from being used to train and improve the model (more common when the model is used via an API).
- Circulate acceptable use policies that set appropriate guardrails for users (see our accompanying article on this).
- If confidential information or personal data will be used, explore whether it is possible to have a dedicated instance of the model running in an environment you control to help ensure particularly sensitive data is segregated (note this may increase costs compared to public or multi-tenant instances).

2. Will you own the generated outputs?

The main output of generative AI, as the term suggests, is the content it generates. There is a broad spectrum of IP ownership and licensing positions across AI vendor contracts. Legal teams therefore need to consider who owns the generated output, who has a right to use it and for what purposes.

Due diligence:

- Consider the use case and whether IP ownership is key or not:
  - If it involves incorporating AI into your product/service, it is likely you will want to own the full "stack" including all outputs (e.g., the enhanced medical image, or the advertising copy). Indeed, you may need to own outputs in order to stay in line with contractual commitments to your own customers (e.g. IP ownership warranties and onward transfers to customers)…
  - …But if the tool is being used in a way that is ancillary to your product/service (e.g. to facilitate background R&D, or customer communications) you may be more relaxed around IP ownership and licensing.

Internal mitigation:

- Check the vendor terms to ensure the position you need is reflected (e.g. full ownership or an adequate licence to use outputs as you intend).
- As an emerging technology, generative AI is challenging the boundaries of copyright protection, so organisations should tread carefully if relying on having exclusivity over generated content (for example, if creating a new brand logo).

3. What if the output infringes a third party's IP rights?

As generative AI models are trained on large volumes of (often publicly available) data, there is a risk that IP infringement claims could be brought by third party rights owners. The risk is higher if the model relies on a specific dataset (think industry-specific journals, or photos of niche categories of subjects) or reproduces its training data in whole or in part (imagine a third party's logo appearing in an AI output). Unless the vendor has excluded protected material from its training datasets, the original rights owner may seek to assert its IP rights over their work (which may even include outputs of the tool).

Due diligence:

- Ask how the AI vendor obtains and uses model training data. Some state the sources of their training data, for example by saying that they use stock images licensed from a rights holder, openly licensed work or public domain content in which copyright has expired. Such an approach to training will reduce the risk of the model's output infringing a third party's IP rights.

Internal mitigation:

- Before using a specific output, the internal team could run general searches to assess if it may infringe a third party's IP.
- As any infringing outputs of a model may be co-caused by the customer's choice of input prompts (not just the underlying model data), consider internal training for your users on best practice for prompts (or engage a prompt engineer) to help reduce the risk of infringing outputs.
- Have an internal policy to respond to taking down any outputs that (potentially) infringe a third party's IP.

4. How do you know that the AI tool will deliver what you need?

To produce high quality outputs, the generative AI tool will need to be trained on a large quantity of data relevant to the problem it is solving. The current state of the AI market indicates that vendors will not typically offer strong contractual performance commitments, as they may argue that these tools are used for a wide variety of purposes and not any one customer in mind, and that the customer is better placed to assess whether the outputs meet their needs. The risk to a customer is the outputs are inaccurate or otherwise do not deliver the business objective and there is little contractual recourse against the AI vendor.

Due diligence:

- Consider the use case and whether the risk of inaccuracy is likely to be too great, e.g. it may not be appropriate to use generative AI to diagnose patients.
- Consider whether to commission an AI vendor to develop a proof-of-concept using your own data so you can test whether it will perform as you expect.
- Check what information the vendor gives about its products – several describe the way in which their models are trained, the nature (volume, quality, source) of data used and whether they verify or curate this.

Internal mitigation:

- Ask internal business/product teams to assure themselves that the tool performs as they expect through user testing and internal demonstrations (and low-risk trials with select customers willing to share risk).
- Emphasise the need for the internal team managing the product/service into which the AI tool has been embedded to check and assure the outputs deliver what their product needs.

5. How will your use comply with current and future regulation?

The rapid advance of AI has led to calls for specific regulation to control AI risks and even for a moratorium on further development. This is in addition to various long-running legislative proposals to regulate AI, most notably the EU AI Act. The use of generative AI may need to take various steps in order to comply with incoming laws. Then there is the matter of how sector-specific regulation applies to generative AI use, for example how an advertising agency ensures its AI-powered content platform complies with the CAP Code,

or how a medical imaging company's AI-infused tool complies with medical device rules. How should an organisation integrating generative AI think about compliance with AI law as it evolves?

Due diligence:

- Taking steps to verify the AI vendor's own compliance (with legal requirements or sector-specific standards) can help demonstrate to your regulator that the decision to use the tool was appropriate.
- Understand the use case to determine whether AI-specific regulation may apply, e.g. if the use of the outputs would constitute a "high risk" activity under the EU AI Act or would make impactful decisions about individuals.
- Check if the vendor offers a level of output filtering to reduce the risk of unlawful content being outputted, in the way that some AI chatbots have implemented recently.

Internal mitigation:

- If AI-specific regulation may apply to your use, implement a best practice compliance approach (including impact assessments, internal governance and quality assurance).
- Have an internal team filter out any outputs that pose a clear legal risk, e.g. obscene or defamatory material or outputs that could constitute a malicious communication.
- Have internal teams review and take down potentially infringing outputs on a regular basis.

If you would like any assistance in contracting for generative AI models, or advice on other implications of integrating AI into your products and services, please do get in touch with us.