


Generative AI

Navigating intellectual property



Generative artificial intelligence (AI) tools are rapidly being adopted by many businesses and organizations for the purpose of content generation. Such tools represent both a substantial opportunity to assist business operations and a significant legal risk due to current uncertainties, including intellectual property (IP) questions.

Many organizations are seeking to put guidance in place to help their employees mitigate these risks. While each business situation and legal context will be unique, the following Guiding Principles and Checklist are intended to assist organizations in understanding the IP risks, asking the right questions, and considering potential safeguards.

 **Generative AI introduces numerous risks and questions. Businesses and organizations should contemplate implementing suitable policies and providing training to employees regarding the technology's opportunities and limitations. This proactive approach is crucial in navigating the challenges associated with the use of generative AI.**

What is generative AI?

Generative AI tools can create new content, such as text, computer code, images, audio, sound, and video, in response to a user's prompt, such as a short, written description of the desired output. Current examples of generative AI tools include ChatGPT, Midjourney, Copilot and Firefly.

Generative AI is based on machine learning and generative AI tools are trained using enormous amounts of data, often including billions of pages of text or images. Depending on the approach of the AI tool developer, training data sets may consist of freely available, unencumbered information (pure data), protected data (such as copyright protected works) or a mixture of both.

The trained AI tool is then prompted by human input which triggers a complex series of often billions of calculations that determine the output. It is generally not possible to predict the output or determine whether and to what extent certain parts of the training data influence the output produced.



Overview of issues

Developing generative AI can be extremely costly, reaching tens of millions of US dollars, and most businesses and organizations are opting to adopt third-party generative AI tools or fine-tuning such models using their own data. General issues and business risks include:

Determining use cases

Generative AI can perform many tasks and optimal use cases are still evolving and they will vary across businesses and organizations.

Differences in contractual terms

Generative AI tools are new and best practices and norms for commercial contract terms are still developing. There can be considerable differences in the terms on which developers are licensing their AI tools, including the approach to trade secrets and other confidential information, the ownership of outputs, the availability of indemnities, and obligations on users to mitigate risks through the implementation of staff monitoring and training.

Training data issues

Some generative AI tools have been trained using materials scraped from the internet, including copyright works, personal information, biometric data, and harmful and illegal content. There is ongoing litigation over whether the scraping, downloading, and processing of materials, the trained AI models, and their outputs involve breaches of IP, privacy, and contract. Debates are ongoing about the balance of interests between IP owners and AI developers.

Output issues

Generative AI may produce inappropriate or illegal outputs, including incorrect information, IP infringements, deepfakes, personal information, defamatory allegations, and discriminatory, biased, and harmful content. Technical safeguards are being developed, but given the complexity of the calculations involved, predicting AI behavior in all circumstances is challenging. Additionally, the IP laws of most countries were written before the advent of AI, leading to uncertainties in the ownership of rights in AI outputs.



Changing regulatory landscape

Governments and regulators are considering new laws, regulations, policies, and guidelines for generative AI. These may impose requirements on businesses and organizations using generative AI. Specific regulations are already in force in China, and the European Union aims to implement regulations soon.

This list of issues is not exhaustive and there are potentially many other challenges, including the energy-intensive nature of training and using generative AI.

Many international organizations, such as UNESCO, the OECD, and the Global Partnership on AI, have published guidance on the general principles for the responsible use of AI. Businesses and organizations should consider implementing a staff policy and training for generative AI to encourage responsible experimentation and use.

Generative AI and IP

Generative AI has many IP touch points and uncertainties. While complete mitigation of these IP risks is impossible, the following considerations may be useful for businesses and organizations navigating IP considerations in this evolving technical field.

Confidential information

Confidential information is information that is not publicly available, may or may not have commercial value, is communicated in confidence, and is reasonably protected. It includes trade secrets, which are a type of confidential information that has (potential) economic value or provides a competitive advantage due to its secret nature.



Businesses and organizations using generative AI tools may inadvertently give away trade secrets or waive confidentiality in commercially sensitive information if such information is used for training or prompting AI tools. They should consider putting in place a combination of technical, legal, and practical safeguards to prevent this.

Risks

Generative AI tools may save and train on users' prompts. If users include confidential information in prompts, confidentiality may be lost because the AI supplier has a copy of the information and, further, the information may become part of the model and the output shared publicly with other users.

When businesses and organizations train generative AI tools from scratch or fine tune existing tools using their confidential information, there is a risk of the information becoming available to the public.

Hackers may be able to extract training data, including confidential information, using techniques such as "prompt injection".

Providers of private generative AI tools may monitor and store prompts to check for inappropriate use. In some cases, prompts may be reviewed by the provider's staff.

Mitigations

Check the settings on generative AI tools to minimize the risk that the provider stores or trains using your prompts.

Consider using generative AI tools that operate and are stored on a private cloud.

Check if the providers of an AI tool will store, monitor, and review your prompts. Seek suitable protections and assurances from the provider concerning any confidential information.

Limit access to generative AI tools that use confidential information to staff with authorized access to that information.


Implement a staff policy and provide training on the risks of including confidential information in prompts.

Consider having information security specialists vet and monitor generative AI tools.



IP infringement

Many generative AI tools are trained on enormous quantities (sometimes billions) of items protected by IP. There are several ongoing legal disputes alleging that the scraping and use of these works to train AI, the trained AI models, and their outputs are IP infringements. These cases are largely focused on copyright and trademarks but, in theory, other IP rights could be involved, such as industrial designs, database rights, and patented inventions.

 There is significant legal uncertainty whether AI tools, their training, use, and outputs represent IP infringements. The answer may vary by jurisdiction. Businesses and organizations should consider mitigating the risk by using IP compliant tools, seeking indemnities where possible, vetting datasets, and implementing technical and practical measures to reduce the likelihood of infringement.

Risks

There is pending litigation worldwide to determine whether the training of AI using IP protected items, the use of such trained AI models, and the outputs generated by them amount to IP infringements.

The risk is not limited to AI developers but potentially extends to users of generative AI tools. In many countries, liability for various forms of IP infringement, such as making a copy of a copyright work, does not depend on the intention or knowledge of the alleged infringer.

The courts are yet to resolve whether generative AI developers, providers, customers, and users can be liable for IP infringement, payment of compensation and the destruction of infringing


Mitigations

Consider using generative AI tools that have trained solely on licensed, public domain, or a user's own training data.

When choosing an AI tool, consider if there are providers willing to offer indemnities against IP infringement, specifically copyright infringement. Assess the extent and suitability of the indemnity. For example, the protection might be limited to third-party compensation and conditional on compliance with contractual restrictions and implementation of risk mitigations.

Thoroughly vet datasets when training or fine-tuning generative AI. Verify IP ownership, license coverage for AI training, and compliance with Creative Commons licenses





models or outputs. It is unclear whether courts would consider it proportionate to make orders preventing the use of an AI model trained on IP-protected items.

Regarding potential copyright infringement, some countries' IP laws include exceptions that might apply to generative AI, such as fair use, text and data mining, and temporary copying. However, a lack of harmonization between countries and the yet unknown application of these exceptions for generative AI introduces uncertainty.

Even where courts have issued judgments these may depend on the specific circumstances of the case as well as the provisions of the national law.

or public domain status. Ensure comfort with applicable copyright exceptions in the intended jurisdiction.

Be aware that regulators are considering putting in place obligations to disclose details of IP-protected items used to train models. Consider keeping records documenting how an AI model was trained.

Implement staff policies and training to minimize the risk of producing infringing outputs. Advise against prompts referencing third-party business names, trademarks, copyright works, or specific authors/artists.


Consider implementing measures to check for infringements before using outputs. These may include plagiarism checkers, image searches, and freedom-to-operate reviews.

Evaluate mitigation measures, related costs and the business risk based on context.



Open-source obligations

Code generated by AI might be subject to open-source obligations. When a software application or code is open source, it means that the source code is made available to the public, and users are often granted certain rights and freedoms to use, modify, and distribute the software. However, these rights and freedoms come with obligations that users must adhere to, such as attribution, and these obligations vary depending on the specific open-source license governing the software.

 **Businesses and organizations should consider whether this risk is appropriate for their code, investigate potential indemnities, and implement technical and practical measures to reduce the likelihood of open-source obligations arising.**

Risks

Generative AI could be trained on code subject to open-source requirements, potentially breaching obligations like restrictions on commercial use or attribution. There is an ongoing legal dispute in the US concerning this.

Some open-source licenses specify that any code incorporating the open-source code becomes subject to the requirements of the same open-source license. Users integrating AI-generated code might therefore inadvertently introduce open-source obligations into their projects.

Mitigations

Consider obtaining generative AI tools from providers training exclusively on licensed examples or implementing technical safeguards, such as detecting relevant open-source licenses.


Consider procuring generative AI tools from providers offering indemnities against open-source infringements. Check the extent and suitability of the protection and conditions that apply.

When training or fine-tuning generative AI tools, thoroughly vet training data for sufficiently permissive licenses.

Adopt a risk-benefit approach to generative AI use in coding. If ensuring code is free from open-source obligations is vital, consider prohibiting suppliers and staff from using generative AI on those projects.

Deepfakes: rights in likeness and voice

Likeness and voice are protected in many countries, though such protection is not harmonized. Forms of protection include some IP rights (such as passing off in common law countries), unfair competition laws, human rights, constitutional rights, and publicity rights.

 **Generative AI has the potential to mimic the likeness or voice of specific people, with some tools explicitly designed for this purpose. Businesses and organizations should consider the risks associated with such capabilities.**

Risks

Unauthorized use or imitation of someone's voice or likeness may result in infringement of IP or other rights, with challenges arising from non-harmonized legal frameworks across jurisdictions.

Mimicking likeness and voice may also risk reputational harm or legal actions, such as fraud or defamation. Many countries are considering specific laws and regulations for deepfakes. For example, China has already passed regulations applying to "Deep Synthesis".

Mitigations


Establish a staff policy and provide training explicitly restricting the use of "deepfake" generative AI tools. For approved generative AI tools, enforce policies that prohibit references to specific individuals in prompts.

In cases where there is a legitimate business reason to synthesize someone's voice or likeness, obtain necessary consent and licensing from the subject.



IP rights in and ownership of AI outputs

It is unclear whether new content generated by AI tools such as text, images, or other creative works, can be protected by IP rights, and if so, who owns those rights. Even if AI output is not IP protected, there may be contractual provisions that govern its use.

 **The existence and ownership of IP rights in generative AI outputs is unclear. Businesses and organizations should seek contractual clarity over ownership and consider using generative AI only in cases where IP ownership in the outputs is not crucial for their business model.**

Risks

The IP laws of most countries were written without considering generative AI, leading to uncertainty over whether there can be IP in AI outputs and who would own any such rights. This may not be an issue for some IP rights, like trademarks, but there is widespread concern for copyright.

Recent patent applications, naming an AI system, “DABUS”, as an inventor, have consistently been rejected in countries that have issued judgments because no human inventor has been identified. It is not yet clear whether generative AI can make inventions without human inventors or whether such inventions are patentable.

The US Copyright Office has issued guidance on registering works containing material generated by AI, indicating that a creative contribution from a human is required. Decisions by the Office suggest that a user’s

Mitigations

Review the terms and conditions of generative AI tools to understand who owns the IP (if any) in outputs.

Explore ways to enhance control or rights over outputs by incorporating IP elements like brand names and logos or involving human creativity in modifying or creating new versions of the outputs.

Document the role of humans in the invention or creation process.


Where possible, establish an agreement about who owns the copyright in computer-generated works. The legal tests vary between counties and may be difficult to apply, so an agreement improves certainty.

When commissioning works, consider seeking a warranty that generative AI has not been used.

text prompt alone may not establish copyright, as the prompt merely “influences” the output. However, the Beijing Internet Court has recently decided that a user owns the copyright in an AI-generated image because he adjusted the prompts and parameters so that the image reflected his aesthetic choice and judgment. These varying interpretations of copyright for AI-generated works introduce legal uncertainty regarding global recognition of copyright in generative AI outputs.

A few countries (such as India, Ireland, New Zealand, South Africa, and the United Kingdom) provide copyright protection for “computer-generated works” without human authors. Ukraine has introduced rights for “non-original objects” generated by computer programs.

Consider using generative AI only where IP rights are not essential, such as for internal use, idea generation, and for ephemeral uses such as (personal) social media posts



Checklist

There are many measures that businesses and organizations can use to foster responsible and legally compliant use of generative AI. The following check list may be useful for businesses and organizations looking to put in place responsible practices and to navigate this fast-evolving field.

Staff policies and training

- ❑ Implement a staff policy and training to guide appropriate usage and to encourage responsible experimentation and use of generative AI, including:
 - ❑ Understand the opportunities, risks and limitations associated with generative AI.
 - ❑ Avoid using confidential information in prompts.
 - ❑ Limit access to generative AI trained on trade secrets to staff with authorized access to that information.
 - ❑ Avoid using third-party IP in prompts, to minimize infringement outputs.
 - ❑ Avoid using “deepfake” generative AI tools.

Risk Monitoring and risk profile management

- ❑ Monitor case law and regulations for changes.
- ❑ Regularly assess and update policies based on evolving risks and court decisions.
- ❑ Communicate legal risks clearly to the business to adopt practices according to the business risk appetite.
- ❑ Maintain a list of AI tools, categorizing them based on risk profiles, for example whitelists for tools that can be used by all staff, restricted tools that use confidential information, and prohibited tools.

Record-keeping

- ❑ Consider documenting how AI tools were trained.
- ❑ Ask staff to label AI-generated output and to keep records of prompts used.
- ❑ Document the role of humans in the creation process.



AI tool assessment

- ❑ Review the terms and conditions and settings on externally procured tools (including those trained on internal data) to
 - ❑ Understand if the provider stores your prompts.
 - ❑ Understand what data the tools have been trained on.
 - ❑ Seek tools that use properly licensed or public domain training data or have technical safeguards against using protected data.
 - ❑ Determine if the provider is offering indemnities against IP infringement and what the conditions are.
- ❑ Vet and monitor generative AI tools by information security specialists.
- ❑ Explore private generative AI tools stored on-premises or in private clouds to enhance control and assurance.
 - ❑ Seek suitable protections and assurances from the provider concerning confidential information.

Data assessment

- ❑ Vet datasets when training AI and consider IP ownership and license coverage.

AI outputs

- ❑ Check generative AI providers' terms on IP rights and ownership in outputs.
- ❑ Check for IP infringements before using outputs.
- ❑ Integrate human input and creativity with AI outputs to maintain control over ownership of outputs.
- ❑ Establish agreements on ownership of outputs.
- ❑ Document the role of humans in the creation process.
- ❑ Obtain necessary consent and licensing to synthesize someone's voice or likeness.



Further reading

The WIPO Conversation on IP and Frontier Technologies is a leading global forum to facilitate discussion and share knowledge among all stakeholders on the impact of frontier technologies, including AI, on IP.

The discussion in the eighth session of the WIPO Conversation focused on generative AI and IP to help policymakers understand potential policy choices. More information about the [eighth session of the WIPO Conversation](#), including the program, presentations, and webcast, can be found on the meeting page.

More information about IP and frontier technologies is available on the WIPO website: www.wipo.int/ai.

Next steps

To keep informed about the next session of the WIPO Conversation, sign up for the IP and Frontier Technologies Division's newsletter by emailing frontier.tech@wipo.int.



The document was prepared by WIPO's IP and Frontier Technologies Division, drawing on commissioned work by Matt Hervey (Growling WLG, United Kingdom).



