Generative AI

WIPO
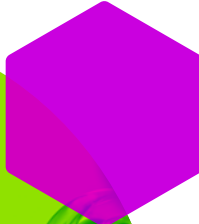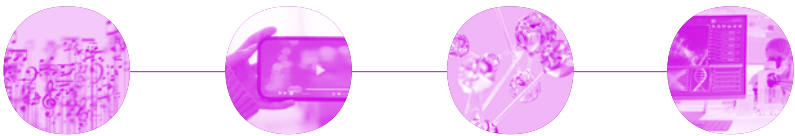
Generative AI

"Once in a while, a new technology emerges that captures public imagination around the world, dominating public headlines and private conversations, dividing detractors and supporters and seemingly putting our world at a crossroads." WIPO Director General Daren Tang

What is generative AI?

Traditional AI systems are primarily used to analyze data and make predictions. Generative AI, however, expands on these capabilities by creating new data similar to its training data. This advancement enables the generation of various forms of new content, including audio, code, images, text, simulations, and videos.

Generative AI refers to any machine learning model capable of dynamically creating output after it has been trained. Importantly, the scope of generative AI extends far beyond content generation. The potential applications of generative AI models are vast, ranging from music composition and video generation to molecular modeling in drug discovery and medical diagnostics. As the technology continues to evolve, its applications are poised to extend into new domains.

While current generative AI models are particularly well suited for correlating language and generating content, the potential of generative AI has wide applications both in the creation and innovation space.

The launch of ChatGPT in November 2022 brought generative AI to the forefront, garnering mainstream attention, and transforming how we work and create. ChatGPT utilizes natural language processing and deep machine learning to generate written content. While it was the first of its kind, similar models including Google's Bard, Microsoft's Bing Chat and Meta's Llama 2 soon followed. Other AI models including Midjourney, DALL-E, and Stable Diffusion generate images from text prompts.

💡 **ChatGPT, a first-of-its-kind large language model, utilizes natural language processing and deep machine learning to generate written content.**

Although some marvel at the accuracy and speed with which generative AI can produce content, critics are concerned about the potential lack of reliability and mistakes in the output generated. Basic-level generative AI models, given a context like a text snippet, merely predict the most probable next word.

💡 **However convincing the output may appear, current generative AI models still exhibit serious limitations, operating by combining words in statistically likely orders without genuine understanding.**

Definitions

Deep learning, a subset of machine learning, draws inspiration from the structure and operation of the human brain to acquire patterns and representations from data. Deep learning employs neural network models with interconnected nodes, or neurons, arranged into multiple layers. The network processes input data through these layers and produces a non-linear output. Each neuron has several adaptable parameters that learn from the data, rather than being pre-programmed.

A recent advancement in generative AI, commonly referred to as language models (LMs) or large language models (LLMs), is due to new models well suited for correlating language. Importantly, LMs and LLMs are not confined to the processing and generation of human languages. In machine learning terms, language can refer to any symbols (such as words) that within context (grammar, relationship between the words) can convey meaning. In language, words create sequences, and meaning relies not only on word choice but also on the relationships between these words. To effectively capture these relationships, LLMs are suitable for processing sequences and maintaining sufficient memory and are designed to handle large datasets. LLMs can perform various natural language processing tasks, including text generation, language translation, text summarization, question-answering, and more.
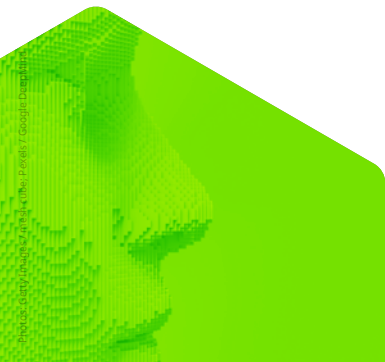
Regulating generative AI

Generative AI represents both substantial opportunities and significant risks. On the side of risks, it raises many concerns beyond its technical limitations, touching on issues of reliability, accuracy, and ethics. At the forefront of much of the public debate are ethical concerns, such as the possibility of societal risks of disinformation, market manipulation, cybercrimes, threats to privacy and democracy, and unintended consequences on content creation and cultural diversity. Furthermore, concerns about job displacement, language bias, lack of transparency, the influence of large companies on regulations, economic inequalities, and the monetization of data collected from countries without commensurate benefits are significant.

Many are concerned about the disruptive potential of generative AI, necessitating regulatory measures. They are advocating for international collaboration and global governance of AI. The need for fast, short-term solutions is pushing policymakers to explore legislative patches and alternative avenues.

Generative AI raises many regulatory concerns, especially ethical concerns. Regulators are actively seeking fast, short-term solutions to match the rapid pace of technological change.

Key IP considerations

Overview

Many are embracing generative AI tools for writing, visual content, and coding. Due to its perceived ability to independently create content, generative AI has the potential to compete economically with human creators and to boost human creativity and productivity across various industries. One significant example illustrating its potential disruption in the creative industries is the song "Heart on My Sleeve," which used unauthorized vocals generated by AI to mimic musicians Drake and the Weeknd. This highlights the ongoing debate around copyright issues: on the one hand copyright works are being used to train generative AI models and on the other hand the models themselves are able to generate realistic content. Content generated by these AI systems may challenge traditional notions of authorship and originality.

However, the IP touch points of generative AI are much wider, and it is worthwhile bearing in mind the whole spectrum of IP rights.

> Generative AI has sparked intense discussions regarding potential copyright infringement by including protected works in training data and copyright protection for AI generated output. However, it is essential to recognize that IP questions related to generative AI are not limited to copyright.

Generative AI models

Patents and trade secrets play crucial roles, particularly when it comes to protecting AI models, alongside the use of open-source regimes.

Patent protection is available for novel and inventive AI models, algorithms, and learning methods. However, the patentability of AI systems varies across jurisdictions, posing the challenge of distinguishing between patentable AI systems and unpatentable abstract concepts implemented with standard computing technology. Some companies also choose to openly share untrained models, even if they hold patents for fundamental technologies.

In some jurisdictions, copyright protection for the source code of AI models and safeguarding key method components through trade secrets may offer immediate and reliable protection compared to patents.

💡 **The best approach for protecting generative AI models with IP depends on several factors, including the AI model's nature, the jurisdiction, and the specific elements intended for protection. Patents, trade secrets, and copyright all play a part.**

Inventing with generative AI

Generative AI has the potential to combine sets of knowledge in ways that humans may not find obvious and hence assist innovative breakthroughs. It can explore vast data sets and identify patterns and solutions, leading to more efficient problem-solving and the discovery of novel approaches.

💡 **Using an AI tool in the invention process can complement human innovation by generating fresh ideas and solutions. Humans still play a critical role in the invention process by defining problems, setting objectives, and determining how AI-generated insights are applied.**

However, employing generative AI in assisting invention may raise the bar for inventive step, potentially making it hard to obtain patents for such inventions. In turn, inventors might resort to safeguarding their innovations through trade secrets, potentially stifling follow-on innovation.

Generative AI inputs

Generative AI models require extensive training data. For example, the original GPT-3 model was trained on 570 gigabytes of text. In some cases, portions of these training data sets were scraped from the internet and contain copyright-protected text and images. This has raised issues of potential copyright infringement.

The developers of generative AI models argue that the models do not keep a permanent copy of any copyright-protected work but simply generate weightings in their neural networks. In this way, the use of training data is similar to human learning by reading books or viewing artworks. They contend that training models falls under existing copyright exceptions and limitations as they do not copy training data. Large training data sets also ensure that AI models are reflective of humanity and there is less risk of bias. Finally, they say that generative AI is not replacing human creativity but complementing human creators and serving as an additional tool that fosters new forms of creativity.

Copyright owners, on the other hand, assert that the use of their works as part of training data sets, if unauthorized, represents copyright infringement. Several lawsuits have been initiated by copyright owners against AI developers, however clear answers to these complex legal questions are unlikely to emerge in the immediate future. Similarly, there is no easy way to assess how much a single work contributed to training an AI model and how to potentially compensate the copyright owner. Furthermore, copyright registration is not compulsory in many jurisdictions, making the identification of copyright works difficult at times.

> Conflicting viewpoints exist regarding the balance between the interests of copyright owners and the need for generative AI developers to use copyright works as part of their training data sets.

Generative AI outputs

Generative AI can produce an extensive array of creative outputs and there is significant discussion whether these AI-generated outputs should benefit from copyright protection. The prolific nature of generative AI may lead to an influx of content varying in quality and creativity, challenging established legal concepts of work originality, authorship, and ownership.

> Copyright was designed as a human-centric right and in many countries only humans can be creators and hence copyright owners. However, that concept is currently being challenged in some jurisdictions.

For example, some applicants in the United States have requested protection for outputs generated by AI alone, which was declined. More recently, the Beijing Internet Court said images generated by Stable Diffusion counted as an original work, because many human prompts amounted to sufficient human creativity.

At the same time generative AI can be used as a tool for human creators to open new fields of creativity and works. If a human creator is using a generative AI as a tool but is making a significant human contribution the resultant work could potentially be awarded copyright protection. The question remains as to how to measure such contribution and establish a threshold.

> There is an ongoing, passionate debate about the meaning of originality in a world where AI plays a substantial role in the creative process.

Case studies

Generative AI is flourishing and is employed for many purposes. The diversity was showcased at the eighth session of the WIPO Conversation. The cases that were discussed ranged from choreographic expression to writing news articles. Another example showed how generative AI serves marginalized communities, such as deaf and hard-of-hearing individuals, with sign language recognition. These case studies also demonstrate the potentially wide range of relevant IP concerns. These examples illustrate the potential tensions between the interests of generative AI model developers, human creators, and copyright owners. Tech companies are asking for more access to data, while rights holders believe that their works should not be used without consent and fair remuneration.
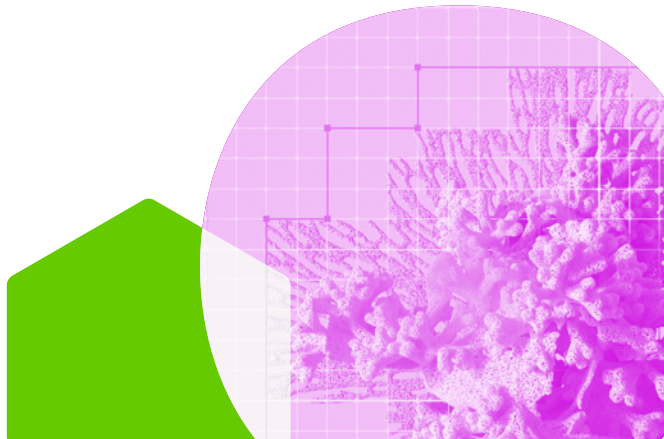
Case study example	IP challenges/questions
Open generative AI models	Access to quality training data is crucial for enhancing AI safety and reducing biases. Navigating jurisdictional variations related to copyright limitations and exceptions, and fair use, is challenging.
AI-augmented fashion and TV production	Determining ownership of AI-generated content is difficult. Lack of diversity in training data may lead to misappropriation and misrepresentation of traditional cultural expressions.
Automatic text processing	Obtaining training data and generating trained models is costly and time-intensive. There is a risk of competitors or customers reverse engineering competing systems if models cannot be protected by IP rights. Compliance with IP rules regarding training data is a concern for small firms, as large players can take a much higher-risk approach.
Music production	AI-generated creations are ineligible for copyright protection in most jurisdictions that have considered the question, posing issues for radio and streaming services. Unauthorized use of songs, music, and lyrics for machine learning is a concern. Authors and artists should be able to opt out of their content being used to train AI.
Sign language recognition and translation software	Protecting AI models might prevent further research and access
Personalized television advertising	Copyright in and ownership of outputs, especially in scenarios with multiple contributors, remains uncertain.
Stock image supplier	Visual data is often protected by copyright, necessitating authorization for machine learning applications.

Risk mitigation strategies

Mitigating risks associated with the utilization of generative AI is crucial for companies entering this dynamic field.

💡 **Businesses can implement several strategic measures to harness the potential of AI while minimizing adverse consequences.**

For AI developers

Exercising due diligence is essential. Companies should consider establishing comprehensive contracts with both service and data providers as crucial to protect against potential legal liabilities that may arise during AI development. For example, when training AI models, using licensing agreements for data obtained from third parties can help ensure IP compliance. Companies may also consider adopting contracts and technologies that help safeguard their AI systems and sensitive data from potential threats and reverse engineering.

For users of AI models

Businesses should have a clear understanding of the terms of use of the generative AI tools they employ. This can help companies navigate potential licensing restrictions and limitations and understand their legal liabilities.

For copyright owners

It may be worthwhile to explicitly state that the work is not available for training AI models or the conditions under which copyright works are available. Additionally, implementing technical protection measures can further safeguard copyrighted works from unauthorized use.
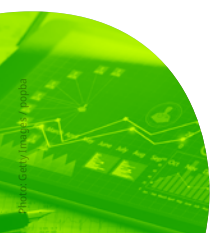
Further reading

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 sending an email to frontier.tech@wipo.int.

WIPO

WORLD
INTELLECTUAL PROPERTY
ORGANIZATION

wipo.int