# Deloitte.
Legal

The legal implications of Generative AI

Considerations for enterprise intellectual property, data protection, and contracts

Deloitte AI Institute™

About the Deloitte AI Institute™

The Deloitte AI Institute helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With™".

The Deloitte AI Institute aims to promote a dialogue and development of artificial intelligence, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries, to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in

artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, deliver impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you're in; whether you're a board member or a C-Suite leader driving strategy for your organization, or a hands on data scientist, bringing an AI strategy to life, the Deloitte AI institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet ups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute

The current enthusiasm for AI adoption is being fueled in part by the advent of Generative AI. While definitions can vary, the EU AI Act defines Generative AI as "foundation models used in AI systems specifically intended to generate, with varying levels of autonomy, content such as complex text, images, audio, or video." (Art. 28b (4) AI Act).

As businesses explore how to use these new tools, there are potential concerns for enterprise stakeholders, particularly legal and compliance professionals.

Generally, a legal executive's role in examining the use of Generative AI is to knowledgeably advise stakeholders (i.e., business leaders, executive peers, the board, and others) on the risks associated with business applications of Generative AI. To this end, it is helpful to understand how Generative AI works and the risk implications the technology presents.

In our papers Generative AI is all the rage and Proactive risk management in Generative AI we explored the capabilities and technical function of Generative AI, how businesses can identify value-driving use cases, and some of the risks and considerations for trustworthy technology that emerge from its application. In this paper, we look at some common legal issues arising in the Generative AI space. Because regulatory frameworks applicable to Generative AI are emerging and quickly evolving, this article avoids a comprehensive discussion of existing or proposed regulations, except where a particular example might provide a better understanding of the relevant risks.

Intellectual property

Gear 7-38 1:2.75

s/n: 3941501

MT-450919

Teet

AI can process vast quantities of data, and without much noteworthy human intervention, transform it into an AI-generated output. The discussion on how to treat any intellectual property rights arising in both the materials used to train the AI (input) and the results created by the AI (output) is still in its early days.

To keep these issues comprehensible, we focus here on legal questions concerning copyright laws, but the same concepts are likely to be applicable to other sorts of protected IP.

Materials used to train the AI (input)

Depending on the law of the relevant jurisdiction, the materials used to train the AI could be copyright protected, and it is likely that reproductions of these materials are made during the training process. Unless certain exceptions to copyrights could be invoked, these kinds of reproductions may constitute an infringement to the copyrights of the author of these materials. These exceptions will vary from jurisdiction to jurisdiction. For example, in the United States, there is a concept of a "fair use" exception, whereas in the EU, the exceptions for transient or incidental copying and text and data mining may be relevant.

Therefore, it is difficult to identify which materials could be used to train an AI system without infringing any IP rights, including copyrights. The recent US Supreme Court ruling in the Warhol case on fair use, which focused more on the commercial purpose of new works than on the artistic expression, is likely to complicate the assessment of US-related copyright risks of AI training materials. However, the ruling's tangible repercussions are not clear yet and will likely be decided in the lower courts.

AI-generated output as a copyright-protected work

Broadly, current copyright law grants rights to the author of a protected work. The focus is on protecting the author's intellectual and personal relationship with their work and to ensure that authors maintain control over the exploitation of their works. However, when it comes to outputs from Generative AI, the question arises as to whether these outputs can have an author, as the composition of the output is not done by a human mind but by an AI system. Lawmakers in their particular jurisdiction will have a role in determining whether granting a copyright to the user complies with the purpose of copyright laws, not least because the user may not have made any free and creative choices that contribute in a meaningful way to the output.

For example, in the EU, the European Parliament stated, in a resolution published on October 20 2020 that works created independently by an AI system are not currently eligible for copyright protection since intellectual property rights generally require an individual that is involved in the creation process. The AI Act does not deviate from this understanding. The US Copyright Office issued a statement in March 2023 that copyright protection does not extend to works generated by AI except to "the extent to which the human had creative control over the work's expression and "actually formed" the traditional elements of authorship" as demonstrated in the Zarya of the Dawn case. Additionally, in August 2023, the US District Court for the District of Columbia affirmed the US Copyright Office's position established in Zarya of the Dawn by granting summary judgment in the

Thaler v. Perlmutter case, where the US Copyright Office denied a copyright application for a work generated by a machine claiming human authorship is required for copyright protection.

Therefore, in broad terms, we may find lawmakers moving toward a position where modifying the output of an AI system and creating a new (derived) work allows the human author to obtain copyright; whereas, the more the output is created by the AI system itself, the less likely it is that such rights will attach. The implications of the Warhol case must also be taken into consideration.

There are uncertainties around the copyright protections for works generated by AI, and it is challenging to determine which data can be used for training without infringing copyright and other IP.

Personal data
and confidentiality

Generative AI systems both ingest and generate large amounts of data, including images, text, speech, video, code, business plans, and technical formulae. Training, testing, uploading, analyzing, consulting, or otherwise processing such input and output data requires various levels of protection.

Such levels of protection depend on the type of data, with a significant distinction between personal and non-personal data. When data qualifies as personally identifiable information (e.g., names, information on a person's life), data protection laws may apply, either locally (e.g., CCPA in California) or regionally (e.g., GDPR in Europe).

Business data, such as financial and technical information, strategic know-how, and trade secrets, may also be classified as confidential information under local laws or by contract, providing for both civil and criminal penalties in cases of mishandling. In this context, when using Generative AI systems, organizations must carefully consider proper categorization of data inputted into these systems and take steps to ensure data is processed lawfully, securely, and confidentially.

To this end, we turn to some of the main challenges organizations face when using personal and confidential data in Generative AI systems, as well as the measures they could adopt to mitigate the relevant legal risks.

Personal data roles and responsibilities

From an EU perspective, a starting point for a personal data protection assessment when using Generative AI is to consider the roles of the parties involved (i.e., data controller, data processor/service provider etc.). This helps define which entity bears the primary responsibility for compliance and what specific actions are to be taken.

In this respect, a Generative AI system provider—as per principle and in a simplified business model—would operate as the data controller for the first layers of training and testing data. Moreover, the provider would most likely operate as an independent data controller for all data, while offering off-the-shelf, data-embedded products. The provider may also act as a data processor on behalf of a customer organization for input and output data, especially where the provider simply licenses the AI "engine" to enterprise customers without any embedded data.

In both cases, the customer organization will likely operate as a data controller for any additional layers of training and testing, for input or output data, depending on the applicable business model. Mixed roles or even joint controllership are also possible and should be considered on a case-by-case basis, in the context of the required data protection and algorithmic impact assessments.

It must be noted that the aforementioned scenarios have not been ruled by any court or supervisory authority yet.

Data protection principles

Across jurisdictions, there are several common personal data principles and protections that are highly impacted by Generative AI systems. When using Generative AI, organizations should pay specific attention to the following aspects of the solutions they use.

Transparency

In their privacy policies and statements, organizations should consider describing (in straightforward language) the use and purpose of AI systems, explain the logic behind AI-powered automated decisions, and highlight risks for the individuals.

Data minimization

While vast amounts of data are required to train Generative AI systems, organizations should consider whether they must limit or exclude personal data from the training set. This could be achieved by using tactics such as filtering personal data from training data, using synthetic data as training data, or preventing end users from inputting personal data into the system's search function.

Lawfulness

In certain jurisdictions, there are specific legal grounds for processing personal data, even when such data were "publicly available" at collection (e.g., consent, contractual necessity, legitimate interests, and legal obligation). Some Generative AI systems seem to be invoking their legitimate interests for processing personal data for system training purposes and contractual necessity for providing the "service". It is plausible that, following a respective careful internal assessment (e.g., proving that their interests in processing the said data outweigh the risks to individuals), organizations may also be able to invoke the same legal grounds, for their own business purposes.

Sensitive data

Several jurisdictions impose increased diligence obligations to organizations when it comes to personal data concerning minors or other sensitive information, such as criminal convictions, medical health, or biometric data. These obligations can include things like age verification, stricter legal grounds for processing (e.g., consent), or even banning processing.

Individual rights

In several jurisdictions, individuals have direct data protection rights. These may include the right to: access and request a copy of any personal information an organization may hold about the individual; ask for the rectification of inaccurate data, such as in case of untrue representations; request human intervention in AI-automated decisions that have significant impacts; opt-out from "legitimate interests" processing; and permanent data deletion. However, considering the underlying technical principles of Generative AI technology, implementing processes that allow compliance with individual rights may be a challenge on its own.

Confidentiality

A breach of confidentiality, imposed either by law or by contract, is a risk to the rights and freedoms of both people and organizations. As such, ensuring the ongoing confidentiality of data across the entire AI lifecycle is an essential factor.

Generative AI models can inadvertently learn and reproduce sensitive information present in the training data. This can result in the generation of outputs that contain confidential information, which, if shared or made public, may compromise confidentiality.

Businesses also need to be aware of their own confidentiality obligations. If a business's use case requires confidential information that has been shared by customers, suppliers, or other third parties, the business will need to first consider any duties of confidentiality and other contract terms under which the information was shared and whether they are permitted to use of that data within a Generative AI system.

Measures to consider adopting

As the use of Generative AI continues to increase, organizations need to carefully assess the existing legal, financial, and reputational risks connected with personal data and confidentiality. Organizations may want to think about the following non-exhaustive list of considerations in addition to legal and regulatory requirements as they come into force:

- Should data access be limited to authorized personnel? What role should physical and logical access control mechanisms, such as authentication systems play?

- What specific policies and procedures for the use of Generative AI tools will be adopted and how will they be maintained, and compliance audited?

- Will policies and procedures be adapted, ensuring the exercise of individual rights (e.g., data deletion)?

- What training and awareness sessions for employees on the ethical, lawful, and secure use of this technology are appropriate?

- How do supply chain audits and controls impact organizations whether they are a supplier or recipient of AI Generative services?

- What technical and organizational measures (e.g., AI governance, privacy-by-design and by-default, pseudonymization, anonymization, encryption, and secure storage) should be put in place to ensure organizations and the personal or confidential data they ingest or retrieve are protected against unauthorized disclosure, alteration, or loss of availability?

- Will legal specialists and technologists be involved in the designing of controls to protect personal data and confidentiality from the early stages of any AI project and will the expertise be in-house or external?

Contractual terms

Given the legal risks associated with the use of Generative AI in a business context, when licensing or otherwise entering into a contract that relates to a Generative AI solution, careful consideration to the terms of the contract under which the solution is procured is important. There are a number of key points that will likely be required to be addressed and understood:

Liability

Organizations may seek indemnities from the Generative AI solution provider for potential IP infringements, data privacy breaches, or confidentiality breaches that arise, and providers will have to consider their own risk appetite in this regard.

Insurance

Especially when dealing with smaller AI solution providers, organizations will consider whether the provider would be able to pay any claims or whether relevant insurance is available.

Business continuity

Since Generative AI solutions may become essential to day-to-day business operations, due consideration is likely to be given to the impact that unavailability may have on the business.

Privacy and confidentiality

As discussed above, provisions regarding confidentiality and data privacy are likely to be a key focus of any contractual framework for the provision of Generative AI services.

AI regulations

Many jurisdictions are developing or about to enact new AI laws and regulations, many of which could override any conflicting contract provision or need to be addressed contractually. This dynamic is likely to be reflected in contractual terms.

The path ahead with Generative AI

Given the evolving legal and regulatory position, legal executives are increasingly likely to be undertaking legal assessments to determine their approach to many of the issues highlighted in this paper.

The risk of infringement on IP rights and/or risk to the award of IP protections; the applicability of personal data protection or confidentiality obligations, as well as the implementation of respective safeguarding measures; and the suitability and enforceability of contractual terms governing the acquisition and implementation of Generative AI tools will all be in focus.

Going forward, legal executives can take a leading role in strategic decision-making related to any use of Generative AI within the enterprise. They are likely to develop responsibilities and accountabilities

in respect of developing ethical and legal frameworks, curating the organization's own risk appetite, in addition to ensuring compliance with law and regulation. Specifically, legal executives should consider staying closely engaged with the evolution of the technology itself, as well as changing laws and regulations. Taking a whole-of-enterprise approach, important stakeholders will include the C-suite, the lines of business, internal expertise and external advisors and consultants who may have the technical expertise to help identify risks, opportunities, and changes to business strategy and processes.

Training people and transforming their approach to understanding the ethical and legal implications of using Generative AI may also fall into the domain of the legal executive.

While the competitive advantage of Generative AI is enticing, adoption of this powerful, differentiating technology demands attention to the risks that could imperil an enterprise's brand, reputation, stakeholder trust, or critically, its compliance with legal and regulatory obligations.

Reach out for a conversation

Key authors

Donna Bartlett
Partner
Deloitte Legal Australia
dbartlett@deloitte.com.au

Willem-Jan Cosemans
Managing Associate
Deloitte Legal Belgium
wcosemans@deloitte.com

Matt Saunders
Partner
Deloitte Digital Canada
masaunders@deloittelegal.ca

Dr. Till Contzen
Partner
Deloitte Legal Germany
tcontzen@deloitte.de

Klaus Gresbrand
Director
Deloitte Legal Germany
kgresbrand@deloitte.de

Maria-Alexandra Papoutsi
Managing Associate
Deloitte Legal Greece
mapapoutsi@kbvl.gr

Pietro Boccaccini
Director
Deloitte Legal Italy
pboccaccini@deloitte.it

Peggy Anstett
Associate Director
Deloitte Legal
United Kingdom
panstett@deloitte.co.uk

Bruce Braude
Chief Technology Officer
Deloitte Legal
United Kingdom
bbraude@deloitte.co.uk

Richard Reeve-Young
Director
Deloitte Legal
United Kingdom
rreeveyoung@deloitte.co.uk

Louis Wihl
Director
Deloitte Legal
United Kingdom
lwihl@deloitte.co.uk

Don Fancher
Principal
Deloitte Financial Advisory
Services LLP
United States
dfancher@deloitte.com

Lori Lorenzo
Managing Director
Deloitte Transactions and
Business Analytics LLP
United States
lorilorenzo@deloitte.com

Our global Deloitte AI Institute leadership

Beena Ammanath
Global Deloitte AI Institute,
Lead
Deloitte Consulting, LLP
bammanath@deloitte.com

Wessel Oosthuizen
Deloitte AI Institute,
Africa, Lead
weoosthuizen@deloitte.co.za

Dr. Kellie Nuttall
Deloitte AI Institute,
Australia, Lead
knuttall@deloitte.com.au

Audrey Ancion
Deloitte AI Institute,
Canada, Lead
aancion@deloitte.ca

Jan Hejtmanek
Deloitte AI Institute,
Central Europe, Lead
jhejtmanek@deloittece.com

Roman Wei Fan
Deloitte AI Institute,
China, Lead
rfan@deloitte.com.cn

Dr. Björn Bringmann
Deloitte AI Institute,
Germany, Lead
bbringmann@deloitte.de

Masaya Mori
Deloitte AI Institute,
Japan, Lead
masayamori@tohmatsu.co.jp

Nicolas Griedlich
Deloitte AI Institute,
Luxembourg, Lead
ngriedlich@deloitte.lu

Tiago Pereira Durao
Deloitte AI Institute,
Portugal, Lead
tdurao@deloitte.pt

Sulabh Soral
Deloitte AI Institute,
United Kingdom, Lead
ssoral@deloitte.co.uk

Deloitte.
Legal

About Deloitte