Trends and Developments

Contributed by:
George Ballas, Zoe Provatas and Nikolaos Papadopoulos
Ballas Pelecanos Law

Ballas Pelecanos Law is a full-service business law firm in Athens, Greece. Tracing its origins to the 1930s, it maintains a determinedly forward-looking and results-oriented mindset. Its specialised legal teams offer a wide range of expert legal services, including legal consultation and dispute resolution on a variety of matters, to its diverse client portfolio spanning most key business sectors. The firm operates on the basis of nine practice groups, each headed by a partner and comprising a dedicated team of associates. Some of its widely recognised core strengths include IP, technology, media and telecommunications, cybersecurity, data protection and privacy law. Maintaining an active interest in the interaction between new technologies and related rights, protection from discrimination and privacy, the team monitors new developments closely, ensuring it remains at the forefront of industry developments.

Authors

George Ballas is the founding, senior and managing partner at Ballas Pelecanos Law. Heading the firm's IP and Technology, Media & Telecommunications practice groups, he is a skilled litigator before the civil and administrative Supreme Courts of Greece, and a qualified European patent attorney. He is also a longstanding member of the Athens Bar Association, the International Bar Association, the International Trademark Association, the European Patent Institute and Lexing®, an international network of lawyers dedicated to digital and emerging technologies law. He is also often invited to author various articles and country chapters for a variety of IP, IT and TMT publications, as well as lecturing on patent issues.

Zoe Provatas is a senior associate at Ballas Pelecanos Law and a member of the IT/CT & Data Privacy practice group. She is an LL.B. and LL.M. graduate of the School of Law of the National and Kapodistrian University of Athens, a Certified Information Privacy Professional (CIPP/E) and a member of the Athens Bar Association, admitted before the Court of Appeals. She has experience in providing legal advice to top-tier clients on a broad range of data protection, new technologies, AI, cybersecurity and innovative projects combining law and IT. She provides consultation to clients in a variety of industry sectors including technology, energy, healthcare, retail and media.

Nikolaos Papadopoulos is an associate with Ballas Pelecanos Law and a member of the IP, IT & CT, the Company, Tax & Employment and the Consumer Rights practice groups of the law firm. Nikolaos graduated from Democritus University of Thrace Law School (LL.B.) and subsequently earned an Advanced Master of Laws degree (LL.M.) from Leiden Universiteit, Netherlands, in Law and Digital Technologies. He is a lawyer before the Court of First Instance and was admitted to the Athens Bar Association in 2017, and also a member of the pool of experts of the European Data Protection Board for new technologies and information security.

Ballas Pelecanos Law

10 Solonos Street,
106 73 Athens
Greece

Tel: +30 2103625943
Fax: +30 2103647925
Email: central@balpel.gr
Web: www.ballas-pelecanos.com

The Legal Implications of Generative AI

Generative artificial intelligence (AI) stands as a beacon of innovation, empowering machines to autonomously create original content across various mediums, from text to images and beyond. Leveraging technologies such as generative adversarial networks (GANs), reinforcement learning and deep neural networks, generative AI reshapes the landscape of creative expression and problem-solving. Its ability to produce diverse and high-quality content has found applications in fields as varied as art, medicine and finance, revolutionising how we approach complex tasks. As generative AI continues to evolve, understanding its multifaceted applications becomes increasingly crucial for researchers, developers and policymakers alike. From generating realistic images to crafting human-like dialogue, generative models exhibit remarkable versatility and potential. However, the widespread adoption of generative AI also prompts a nuanced examination of its legal and ethical implications.

In exploring the boundaries of AI creativity, questions arise regarding intellectual property rights, data privacy and algorithmic bias. As society grapples with these challenges, collaborative efforts are essential to ensure that generative AI realises its transformative potential responsibly. By fostering dialogue and developing comprehensive frameworks, we can navigate the evolving landscape of generative AI while upholding ethical principles and safeguarding individual rights.

Types of Generative AI

Exploring applications of generative AI

Generative AI constitutes a specialised branch within the field of artificial intelligence, leveraging sophisticated machine learning methodologies such as semi-supervised or unsupervised learning algorithms. These applications leverage AI models, such as GANs and Recurrent Neural Networks (RNNs), to create content that is often indistinguishable from human-generated content. Its primary function revolves around the creation of digital content spanning images, audio, videos, code and textual material. Generative AI is used, amongst other things, for text generation, such as content creation or coding, for image generation and manipulation, creating art or deepfakes or image-to-image translation, for audio generation, such as music composition and voice synthesis, for video generation, for game design, and even for medical and scientific research or simulation and training for autonomous vehicles, robots and AI systems.

This process hinges on a training regimen wherein algorithms are exposed to extensive datasets containing pairs of input and output examples. Through iterative learning, these algorithms discern intricate patterns within the input data, enabling them to generate outputs that align with the desired specifications. This training paradigm facilitates the development of AI systems capable of autonomously producing content that exhibits remarkable fidelity and complexity, mirroring human-generated counterparts in various domains.

Advances and innovations

Today's AI systems can autonomously generate creative content across written, visual and auditory realms with minimal human input, creating works virtually indistinguishable from human creations. For instance, advanced text-to-image generators such as DALL·E 2 swiftly produce images based on textual prompts. Trained on a massive dataset of over 650 million image-text pairs, DALL·E 2 goes beyond simple imitation, grasping contextual understanding. OpenAI, the developer behind ChatGPT and DALL·E 2, pioneered the use of Aesthetic Quality Compari-

son, training a model to predict human aesthetic judgements using video data. This approach allows DALL·E 2 to craft art consistent with human perception, though it operates differently from human perception itself.

Listed below are a few interesting examples that gained traction during the past months:

- ChatGPT, from OpenAI and backed by Microsoft, leads in generative AI. Its latest version, GPT-4, is known for human-like responses and boasts improved AI capabilities. Microsoft's substantial investment in it and its integration into Bing highlight its growing importance.
- DALL.E, from OpenAI, leverages GPT implementation to connect words to visual elements, generating images from user prompts and showcasing its multimodal capabilities.
- Google Bard, inspired by Microsoft's use of GPT in Bing, offers a chatbot interface. Initially flawed, it improved with PaLM 2, providing better visual responses.
- Midjourney, a top generative AI, swiftly generates high-quality images from text prompts, making it preferred alongside DALL.E for its accessibility.

While chatbots like the above have quickly risen in popularity, generative AI is also used in different areas and sectors.

Image generation and manipulation

Generative AI commonly generates images from text prompts, allowing users to describe their desired image. The AI interprets these prompts to produce realistic images that are customisable in subjects, styles and settings. This interaction creates diverse visual content, aiding creative expression and design across domains. These types of generative AI include, amongst others, functionalities such as semantic image-to-image translation, image completion, image super-resolution and image manipulation.

Software and coding

Generative AI is already affecting software development and coding, boosting productivity and code quality. This rapidly evolving field holds vast potential to inspire new avenues of software innovation while improving efficiency.

A key application in software development is code generation, which extends to code completion, automated testing and enabling natural language interfaces. This enables developers to interact with software systems using human language instead of programming languages.

Video creation

Generative AI streamlines video production with novel features, automating tasks such as video compositions, animations and special effects. These tools create high-quality video content from scratch, enhancing resolution, manipulation and overall completion.

Such functionalities entail video style transfers and video predictions.

Audio generation

Generative AI is also used to create audio. Audio generation can be categorised as follows:

- *Text-to-speech generators:* AI creating realistic speech audio from a user's textual prompts.
- *Creating music:* AI generating complete audio for novel pieces of music by learning the styles and patterns of the music inputs.
- *Speech-to-speech conversions:* AI creating new speech or voices via existing audio files.

Text generation and summarisation

Certain AI models are trained on large datasets to generate up-to-date and authentic content. Some of the most common use cases of generative AI applications used for text generation and summarisation are listed below:

- *Content creation:* Generative AI models assist in various writing activities. Applications like these expedite writing by generating ideas, quotes and content outlines.
- *Language translation:* Text analysis and translation from one language to another.
- *Virtual assistants and chatbots:* Generative AI fuels chatbots and virtual assistants, generating natural responses in real-time conversations. ChatGPT exemplifies this, enhancing user engagement and offering personalised assistance for businesses.
- *Content aggregation:* Generative AI tools can automatically summarise bulk texts to help users understand them better.
- *Automatic report generation:* In business environments, the summarisation of large datasets into understandable reports could provide users with an advantage. Generative AI analysis may lead to identifying patterns and highlighting insights.

Legal Implications

Data privacy

Generative AI systems are fed with training data and learn to generate statistically probable outputs that have similar characteristics. It is therefore clear that AI enables the collection and use of large amounts of data, both personal, ie, information relating to an identified or identifiable individual, and non-personal; data feeds AI systems and AI systems generate more data.

The fact that the effectiveness and fairness of AI tools depend on the quality and quantity of data puts individuals' fundamental right to protection of personal data and private life in jeopardy. Inclusion of personal data in training sets poses privacy and other risks to individuals, including, inter alia, that information in training data could foreseeably be produced as part of a generative AI system's output. In the absence of a more specific regulatory framework at European level, the General Data Protection Regulation (GDPR) is called upon to address and regulate key aspects of the functioning of AI as far as the processing of personal data is concerned.

The overriding issue that emerges is the applicability of some principles governing the processing of personal data, which is highly impacted by generative AI systems. In particular, the following issues arise with regard to such principles:

Transparency

The principle of transparency is fundamental to the protection of personal data. Transparency is primarily achieved by providing individuals with the necessary information regarding the processing of their personal data, whether collected directly by the individuals themselves or by third parties. In their privacy notices, organisations must inform individuals that their personal data may be used to train and test a generative AI system as well as about the purpose for which their personal data are processed, explain the logic behind AI-powered automated decisions, and highlight risks for the individuals.

Purpose limitation

Personal data must be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes or beyond the affected individuals' reasonable expectations. During the development and deployment life-cycle of an AI system, organisations should carefully evaluate

the compatibility with the purpose for which the personal data used in its development were collected.

Data minimisation

Personal data must be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed. While vast amounts of data are required to train generative AI systems to achieve their full potential, developers, providers and deployers of generative AI systems should limit the collection, use and further processing of personal data only to what is necessary to fulfil the legitimate identified purposes. Therefore, personal data must only be used as training data if required to achieve the legitimate identified purposes of the generative AI system, while the use of anonymisation or pseudonymisation techniques should be taken into consideration.

Accuracy

Pursuant to Article 5 of the GDPR, personal data must be accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay. It is obvious that the accuracy of the output of generative AI systems highly depends on the accuracy of training data. Therefore, generative AI systems must rely on accurate, reliable and representative data. False or inaccurate personal data must be excluded from training data. But even when trained with representative high-quality data, the output generated by generative AI systems may contain inaccurate or false information including personal data leading to hallucinations, in which a tool confidently asserts that a falsehood is real. To mitigate the risks posed by the potential lack of accuracy of generative AI systems, it is important that it must be indicated when there is uncertainty regarding generative AI responses so individuals have the chance to validate the output, eg, by citing the sources on which the output is based and using technical safeguards.

Privacy by design and by default

Rapid technological change poses new risks to data protection. Some of the unique characteristics of AI render compliance with data protection laws more challenging in comparison with more "traditional" IT systems. In line with the privacy by design and by default principle, organisations should conduct a data protection impact assessment to identify, assess and address the risks posed by generative AI systems at every stage of their life-cycle. The state-of-the-art security measures designed to implement data protection principles must be implemented in an effective manner, and security safeguards should be integrated into the processing in order to meet the requirements of the GDPR and protect the rights of individuals.

Biased or inaccurate information

Bias can occur in various stages of an AI system life-cycle. AI systems are based on machine learning data-driven techniques, so the primary source of bias is data collection. If generative AI systems are trained with data which may not be diverse or representative and/or reflect discrimination, they may generate outputs which have discriminatory effects on individuals based on their gender, race, age, health, religion, disability, sexual orientation or other characteristics. Furthermore, if training data is not balanced or the system architecture is not designed to handle diverse inputs, the generative AI system may produce biased outputs. In addition, bias may be introduced if the generative AI system is not tested with diverse inputs or monitored for bias after deployment.

Apart from the above, inadequate or biased training data and incorrect model assumptions may lead generative AI systems to generate responses which contain false or misleading information presented as facts, the so-called "hallucinations". AI hallucinations may have various significant consequences/liabilities for organisations and/or persons using incorrect output from generative AI systems. For example, such an organisation or user could potentially suffer reputational damage or even charges of libel as well as negligence claims if they have used generative AI to provide advice.

Mitigation of AI hallucinations

The best way to mitigate the impact of AI hallucinations is to stop them before they happen. It must be ensured that generative AI systems are trained on diverse, balanced, well-structured and high-quality data from reliable sources. Datasets must be transparent to make AI outcomes understandable and traceable. Also, organisations should establish the AI system's responsibilities and limitations; this will help the system complete tasks more effectively and minimise irrelevant, hallucinatory outputs. Finally, the generative AI system must be tested and evaluated thoroughly before use and on an ongoing basis as it evolves and improves.

In any case, human oversight must be established to mitigate the risks of AI hallucinations. As a final measure for the prevention of hallucination, a human being should review, filter, correct and validate generative AI outputs. Outputs generated by generative AI systems must be verified against a credible source. Human reviewers may also provide relevant expertise and increase the accuracy and safety of AI systems, upholding human values.

The Draft EU AI Act

The European Commission, proposing the first-ever legal framework on AI ("Draft EU AI Act"), which is expected to be adopted soon, attempts to address the challenges posed by AI following a proportionate risk-based approach. Pursuant to the Draft EU AI Act, AI systems will have to meet data transparency obligations. AI system providers will have to ensure that AI systems intended to interact with individuals are designed and developed in such a way that individuals are informed that they are interacting with an AI system unless this is obvious from the circumstances and the context of use. In addition, users of an AI system that generates or manipulates image, audio or video content that appreciably resembles existing persons, objects, places or other entities or events and would falsely appear to a person to be authentic or truthful ("deepfake") will have to disclose that the content has been artificially generated or manipulated. As far as high-risk AI systems are concerned, they are faced with much stricter transparency obligations as well as the requirement for appropriate human oversight.

Intellectual property

Generative AI applications pose a new challenge to current intellectual property laws due to their capability to independently produce original content. Different levels of AI involvement in content creation are described in academia. For instance, "AI-assisted work" implies that a natural person, not AI, is functionally considered the author of the work, while "AI-generated" indicates that no natural person qualifies as the author.

A significant concern revolves around attributing ownership to AI-generated works. Traditional copyright laws designate human creators as owners, but determining authorship becomes unclear with generative AI. This ambiguity may

trigger disputes over intellectual property rights, as multiple parties could claim ownership of AI-generated content. High-profile lawsuits by content creators, such as the New York Times and Getty Images, against generative AI developers in the US and EU have escalated these concerns.

The discourse on this issue appears twofold, as legal theorists focus on both stages: the algorithm's training and the generation of outcomes.

Training of the algorithm

The primary copyright concern regarding AI training revolves around the possibility that training datasets might contain copyrighted text or materials. Lawfully reproducing or using these materials in the training process requires permission from rights holders or specific legal provisions allowing their use in training language models.

Training generative AI algorithms, such as language models (LLMs), encompasses large-scale datasets and numerous potential rights holders, making it highly challenging to seek all the rights holders and obtain explicit licences.

On the one hand, it has been argued the use of training datasets could be lawful by applying the text and data mining (TDM) exception provided by Directive 2019/790 to train language models (LLMs). Based on the definition of TDM as "any automated analytical technique aimed at analysing text and data in digital form to generate information, including patterns, trends, and correlations", such activities could fall within this definition.

On the other hand, Article 4(2) of this Directive dictates that reproductions and extractions of content, such as described above, may only be retained for as long as is necessary for the purposes of text and data mining. This could result in the obligation of the trainers of LLMs to delete copyrighted content as soon as the training of the algorithm is concluded.

Output generation

The discussion on output generation includes not only the potential infringement of materials used during the training of LLMs by the outputs produced but also the possibility of legally protecting such outputs via copyright or patent legislation.

In a nutshell, the output generated by a generative AI application could lead to two main outcomes:

- infringing intellectual property rights of works used for the training of the algorithm;
- autonomous creations, legally separate from the pre-existing materials.

A generative AI output could potentially infringe on legal rights in two primary ways. Firstly, if the output closely resembles legally protected elements of pre-existing materials, and secondly, if the output incorporates protected aspects of pre-existing materials through unauthorised adaptations or modifications, it would likely be considered a derivative creation. Another important aspect of this intriguing problem is that the CJEU has recently determined in YouTube v Cyando (Joined Cases C-682/18 and C-683/18) that if platforms fail to comply with any of three distinct duties of care, they will be directly accountable for violations of the right to communicate a work publicly.

Regarding the outcome's protection, the European Parliament emphasises that existing intellectual property legislation still applies when the creative outcome primarily stems from human

intellectual activity, albeit with assistance from an AI system. The CJEU, as demonstrated in the Painer case (Case C-145/10), confirms that copyright-protected works can indeed be created with the aid of a machine or device. According to CJEU case law, predominant human intellectual activity is evident when a human creator utilising generative AI exercises free and creative choices during the conception, execution and/or editing phases of the work.

Similarly, a broader interpretation of the inventive step requirement may be warranted when considering patent protection for inventive outcomes produced with generative AI support. This interpretation would focus on non-obviousness to a person skilled in the art, assisted by AI, ie, an AI-aided human expert, as many scholars have pinpointed.

Conclusion

Generative AI holds immense promise for innovation and creativity, yet its proliferation underscores the pressing need for robust legal and ethical frameworks. Intellectual property rights face new complexities as AI-generated content blurs the lines of ownership and authorship. Moreover, concerns regarding data privacy and algorithmic biases demand careful consideration in the development and deployment of generative models. As society grapples with these challenges, collaborative efforts between policymakers, technologists and legal experts are paramount. By fostering dialogue and cultivating responsible practices, we can harness the transformative potential of generative AI while mitigating its risks. Ultimately, navigating the evolving landscape of generative AI requires a balanced approach that fosters innovation, safeguards ethical principles, and upholds the rights and dignity of individuals.