

The conundrum of AI Art and copyright

In recent years, the development of sophisticated artificial intelligence systems is ushering in a new era of digital art. AI Art, the art that is produced by generative AI systems, has been met with equal amounts of excitement and uncertainty. These tools have the power to both enhance and threaten the creative industry as we know it. While it is creating new opportunities for artists, it is also raising legal and ethical questions from an intellectual property perspective.

Many in the creative industry are concerned about the so-called 'input' data that is used to develop these systems. They worry about how information is scraped from the internet, including art protected by copyright, and is then used to produce 'new' images without permission from the copyright holder. Further, there is confusion around whether the generated products, the 'output', attracts copyright protection. And if it does, who owns that copyright.

Consequently, a plethora of lawsuits involving generative AI systems are producing different answers to the question of whether AI-generated art attracts copyright and whether AI system training is lawful.

What is generative AI and system training?

Generative AI is a subset of AI that generates content strongly resembling the product of human creation. Users insert a text prompt, and the system produces outputs including texts and images in response. To do this, the generative AI system is trained on huge datasets of existing material 'scraped' from publicly available sources on the internet. The system can then be fine-tuned to a particular content domain, honing the data down to the most relevant. After the system has been trained, it draws on deep learning techniques and algorithms to 'learn' patterns and structures within existing content to generate 'new' content in response to a user's prompt.

Examples of generative AI platforms include the Chat-GPT, a text-to-text system and, more significantly for this article, DALL-E, Midjourney and Stable Diffusion, text-to-image systems.

Copyright issues: input

Text and data mining (TDM) is one technique used during the training stage of generative AI systems. The TDM process involves the extraction of vast amounts of data to analyse and identify patterns which can be used to improve the performance and output of AI systems. This data scraping often includes copyright-protected content reproduced without consent from the copyright holders.

This raises serious concerns for artists. Scraping data has the potential to impact artists' careers as AI systems learn how to create works in their artistic styles. This means users can easily recreate other artists' art. Digital artist Greg Rutkowski has reported that his name ranks among the most popular prompts – more so than Picasso. After scraping his images from the internet, users can prompt AI platforms to produce output that strongly resembles the classical painting style of fantasy landscapes for which he is known. Digital artists are being inadvertently punished for making their art accessible. TDM therefore raises difficult questions of ethics and legality, and the UK, EU and US are diverging on their approach to this practice.

Legal position

United Kingdom (UK)

In the UK, copyright is governed by the Copyright Designs and Patents Act 1988 (CDPA) under which there is one narrow TDM exception. Section 29A allows TDM of copyrighted works for non-commercial research provided that the user has lawful access. The UK recently abandoned a proposal to broaden the exception following criticism from creative industries who believed the proposal had not taken sufficient account of the potential harm it could cause them. However, not widening the exception could negatively impact the UK's

position as an AI leader as countries with broader copyright exceptions to TDM are attracting more AI developers.

While not extending the TDM exception, the UK Intellectual Property Office (UKIPO) has been working on establishing a voluntary AI copyright code. This code seeks to balance the need for generative AI platforms to have access to training data with the need for creators to be compensated for the use of their copyrighted works. Striking this balance presents a challenge and the UK government, with support from representatives in the technology and creative industries, does not want to rush its conclusion. However, with the increasing uncertainty around data training infringing copyright protection and the negative impact this is having on the creative industry, the need for guidance is becoming pressing. This non-legislative code was intended to be completed by autumn 2023, but it is now unlikely that we will see it until 2024.

European Union (EU)

In the EU, there is a broad TDM exception provided under the Directive on Copyright in the Digital Single Market.¹ Article 4 allows individuals such as commercial AI developers to make copies of works for the purpose of extracting information from text and data and retain them for as long as they are needed for AI training – as long as they had legitimate access to the content, and that the copyright owner had not expressly ‘opted out’ of the exception. If the copyright owner has not opted out, AI providers can rely on the exception to justify extracting or reproducing his/her works.

The forthcoming EU AI Act (the EU Parliament adopted amendments of 14 June 2023 combined with the EU Commission Proposal of 21 April 2021 is the latest text available at the time of writing) imposes specific requirements on generative AI systems. Providers will have to:

- train, design and develop the generative AI system in such a way that there are state-of-the-art safeguards against the generation of content in breach of EU laws;
- document and provide a publicly-available detailed summary of the use of training data protected by copyright; and
- comply with stronger transparency obligations.

The former two obligations aim to protect against the infringement of intellectual property rights (and in particular against copyright infringement).

The latter obligation aims to avoid, through transparency, the use of a generative AI systems to create manipulative content. Where a generative AI system has been used to create “deep fakes” (i.e. text, video or audio that appears to be authentic or truthful while it is not), the users that created such content must disclose that the content is AI generated or manipulated and (where possible) indicate the name of the legal or natural person that generated or manipulated the content.

Further, generative AI systems qualify as a type of foundations model² (or general-purpose AI (GPAI) model, the new terminology apparently voted on when the Act (not yet published) was adopted on 8 December 2024), accordingly providers of generative AI systems must comply with the obligations imposed by the Act on providers of foundation models/GPIA models. They include:

- only using datasets that are subject to an appropriate data governance ensuring that the datasets are suitable and unbiased;
- designing, developing and testing the foundation model to ensure performance, predictability, interpretability, corrigibility, safety and cybersecurity throughout its lifecycle;
- developing technical documentation and intelligible instructions for the foundation model. The provider must keep this technical documentation available for the competent authorities for a period of ten years from the date of market introduction;

- establishing a quality management system to ensure and document compliance with the AI Act; and
- registering the generative AI model in a public EU database that the EU Commission has been tasked to introduce.

Finally, generative AI systems must comply with the obligations that apply to AI systems depending on their risk categorisation. The Recitals of the AI Act clarify that the development of a generative AI system or foundation model, as such, does not lead to a high-risk classification. For each specific generative AI system, one must assess the risk classification of such AI system – and comply with the corresponding obligations.

The final text of the EU AI Act is not available yet. It is possible that when published, the terminology and the obligations imposed on providers of generative AI systems will have been slightly modified. However, we do not expect significant changes.

The EU AI Act will come into force progressively, most likely between summer 2024 and summer 2026. Compliance obligations affecting generative AI systems are expected come into force in mid-2025.

United States (US)

In the US, in place of a TDM exception, the law relies on ‘fair use’ under section 107 of the US Copyright Act 1976. The fair use doctrine permits the unlicensed use of copyrighted material in certain circumstances. Section 107 provides the legal framework for determining what uses may qualify. This involves balancing, on a case-by-case basis, the purpose and character of the use (transformative use), the nature of the copyrighted work, the portion used and the effect of use on the potential market for, or value of, the protected work.

AI developers have long taken the view that fair use would protect them from copyright infringement claims. In *Authors Guild v Google*, the court ruled in favour of Google that the actions of scanning and digitalising printed copyright-protected books for an online searchable database was protected under the fair use doctrine.³ The court deemed Google’s actions sufficiently transformative. However, considering the recent *Andy Warhol Foundation for the Visual Arts, Inc v Goldsmith* case, fair use may not be the strong foundation that AI companies thought it was.⁴ Here the courts shifted the fair use analysis away from transformative use and considered how the secondary work was being used.

AI is inherently transformative, taking existing data to generate something new. However, transformation should be more than just an aesthetic alteration. It should significantly alter the original work’s function and purpose. Though not a matter of AI Art, the court in *Warhol*, by applying the doctrine of fair use as it did, could undermine the strength of generative AI developers’ case in future copyright infringement claims.

Currently, there are ongoing claims on TDM issues being brought in the US. Earlier in the year, Getty Images filed a lawsuit accusing Stability AI of infringing its copyright after more than 12 million images were used for training data.⁵ A similar claim was also filed in the UK on the same basis. No court in either country has issued a ruling in the disputes but given the differing TDM laws, it will be interesting to compare the approach taken by each court.

Google also recently asked a US court to dismiss a lawsuit against them claiming that data scraping to train a generative AI system violated millions of people’s intellectual property rights. Google argues that the lawsuit would ‘take a sledgehammer not just to Google’s services but to the very idea of generative AI’.⁶ If every jurisdiction took a restrictive approach to TDM, then AI systems would struggle to gather enough data to adequately train.

Potential solutions

Since raising multiple TDM lawsuits, Getty Images released its own AI photo-generation platform trained solely on their stock image library. So confident that their platform is free of any third party copyright-protected material, they are offering indemnities to their users. This is an increasing trend. Other system providers are

offering similar indemnities to cover any copyright infringement claims brought by third parties against users of their platforms. This is an attempt to combat the wariness some are feeling about using generative AI tools for fear of the intellectual property risks.

Since it is not possible for all generative AI platforms to have their own copyright free libraries to train on, and with copyright lawsuits already underway, alternative solutions will need to be discussed for safeguarding AI training data. Ideas including data-sharing agreements or royalty-based compensation models have been suggested.⁷ Data-sharing agreements could regulate the access, use and administration of protected content used in the AI training process. A royalty-based scheme could be implemented that provides creators with a set fee or portion of revenue for the use of their copyright-protected material. Such solutions could also act as an incentive for users to contribute their works to the future improvement of generative AI systems.

Until a satisfactory arrangement can be reached, AI developers need to be aware that the creative industry is coming up with their own solutions and protection. A new data poisoning tool, Nightshade, was announced at the end of October 2023. This tool can 'poison' AI training data and damage the outputs of image-generating systems to the point of rendering them useless. The poisoned pixels are invisible to the human eye – an artist's image of a dog will still look like a dog. However, an AI system will read the poisoned pixels of the dog's image and generate the image of a cat.

It is still too early to tell what the significance of Nightshade will be, but the software developers behind it have created it as an open source. This means that other software developers can take the programme and build their own versions thus broadening the data poisoning tools available. This poses a real issue for the future of AI systems. Once a system has been exposed to poisoned data, it is difficult to fix. Exposure to enough poisoned images may permanently damage an AI system to the point of being unusable.

Copyright issues: output

Further copyright issues arise in the output produced by generative AI. With generative AI systems, it is easy for anybody to be an artist. However, whether the art produced by generative AI systems attracts copyright protection is uncertain. Copyright protects the author from others copying their creative output. Where there is no direct human involvement in the creative process, there is typically no author.

UK

The UK is unique in that it already protects computer-generated works which do not have a human creator/author (s 178 of the Copyright Designs and Patents Act 1988 (CDPA)). Section 9(3) CDPA provides that the author of a computer-generated work is 'the person by whom the arrangements necessary for the creation of the work are undertaken'. There is uncertainty around who this definition of an author or 'arranger' relates to. It is likely to be the user of the generative AI platform, the one who inserts the prompts. An AI system cannot be an author as it is not human, and it is argued that the AI developer is too far removed. This provision was established in 1988 and it remains to be seen whether it is adequate for today's advanced technology. The UK Government has announced that it has no plans to amend the CDPA to cater for AI generated works at this time.

EU

There is no statutory provision granting copyright protection to computer-generated works in the EU. Rather, a work must satisfy the originality criterion to attract copyright. The prevailing approach for determining originality is through evaluating whether the work is of an author's 'own intellectual creation'. The EU has not ruled out the possibility that this criterion could extend to AI Art. In an EU Commission paper, it was recommended that AI Art qualify for copyright protection if a human made sufficient creative choices which are expressed in the final AI-assisted output.⁸ AI-assisted work would require that a human edit, refine and

improve the generated content to have some input in the final work. It appears that the key is human creativity though there is no case law or parameters to guide this yet.

US

The US has taken a strict approach to 'human authorship'. Like in most countries, the US Copyright Act 1976 does not extend protection to computer-generated works. There have been multiple AI Art copyright claims and so far, the US is taking a hard and fast approach of dismissal.

The US Copyright Office (USCO) amended their decision to grant copyright for the graphic novel, *Zarya of the Dawn* in February 2023 after discovering the author used the AI platform Midjourney to generate images. In their updated decision, they limited the copyright registration to only the text and the section and arrangement of the images, but not the individual images themselves.

In August 2023, a U.S. court in Washington D.C. upheld the refusal of USCO to grant copyright to the AI generated image *A Recent Entrance to Paradise*. The USCO Review Board stated that they would not register work 'produced by a machine or mere mechanical process' that operates 'without any creative input or intervention from a human author'.⁹ In court, Judge Howell expanded on this ruling that 'copyright has never stretched so [...] as to protect works generated by new forms of technology operating absent any guiding human hand [...]. Human authorship is a bedrock requirement of copyright'.¹⁰

The above indicates that, in some circumstances, the US may grant copyright provided there is creative input from a human. However, in September 2023, USCO rejected copyright for the AI generated image, *Théâtre D'opéra Spatial* for a second time despite the author arguing that beyond entering the prompt, they adjusted the scene and dictated the tone of the image. USCO requires that a 'meaningful' amount of human effort in the final work must be proven. There must be 'sufficient' variation to distinguish it from the original generated image.¹¹ However, with USCO consistently refusing to grant copyright in any AI Art, what constitutes a 'meaningful' or 'sufficient' amount of human input is uncertain.

Generative AI enthusiasts have argued that the outputs of generative AI systems are no different to works produced by a camera. The act of a user inserting a prompt is akin to a photographer pressing a button, thus they argue that the output should attract copyright protection in the same way. However, Judge Howell distinguished photographs from AI Art in her ruling upholding that photographs are the production of a photographer's, the author's invention and creativity. The ability to copyright photographs rests in the fact that the human creator, not the camera, conceived of and designed the image and then used the camera to capture it.

For now, there are minimal guiding principles when it comes to copyright eligibility in AI Art. Human involvement and creative input continue to be the reference point for copyright protection, but how far these concepts will be stretched is unclear. The UK and EU courts have yet to form a view. However, it is inevitable that AI Art copyright claims will arise, and courts throughout Europe will have to take a position. Currently, signs are that they will offer copyright protection to AI Art more readily than in the US. The US's strict human authorship requirement may evolve, in time, to accept AI Art where there is a degree of human involvement, however, current case law does not suggest this will happen any time soon.

Final thoughts

The rise of generative AI is regularly cited as 'unprecedented' in the media. However, generative AI systems have been around as early as the 1960s when the first chatbot, ELIZA, was released (though the term was coined later). In the past year, with the emergence of systems like Chat-GPT and DALL-E, generative AI has been forced to the forefront of our minds. With governments looking to regulate AI, proposing various legislation, frameworks, and codes, as well as hosting safety summits, it is understandable that some are feeling a deep sense of distrust towards AI.

However, with the current criticisms it can be easy to forget that between 2016 and 2020, AI Art was regarded with interest and welcomed in the art market. In 2016, *The Next Rembrandt*, a 3D printed painting was created using AI that imitated Rembrandt van Rijn's style and brushstrokes. Five years later, AI was used again to reconstruct the missing pieces of another Rembrandt work housed in the Rijksmuseum in The Netherlands.

In 2018, Christie's auctioned the AI Art *Portrait of Edmond Belamy* for \$432,500. The following year, Sotheby's also engaged in the sale of AI generated art. Exhibitions and art fairs showcased various AI artists and the world seemed genuinely interested in this new artistic movement.

This fusion of art and technology and its rapid and currently under-regulated development has raised concerns. However, AI has many applications capable of adding value to the work of scholars and experts. Already AI systems are assisting in the deciphering of ancient, damaged scrolls. AI applications also are being developed to assist the authentication process. For now, AI cannot and will not replace the human eye and expertise. But AI is a tool that can create opportunities and improve current practices and accessibility.

Ellie Pritts is a good example. Pritts is an LA-based artist who developed a degenerative neurological condition which impacted her ability to continue to create art. To overcome the physical limitations caused by her condition, she uses generative AI to make new art. Trained on a dataset of 15 years of her own work, generative AI has allowed Pritts to continue her art practice. Such systems have also helped re-establish the work of game illustrator, Sean Aaberg, who lost his ability to sketch after a stroke left him partially paralysed. A self-proclaimed Luddite, Aaberg now believes that generative AI has saved his artistic life.

Unfortunately, it is unclear at this stage whether AI art produced by those artists qualifies for copyright in their respective jurisdiction. While not all AI Art will or should be protected by copyright, these are two examples that demonstrate how AI Art can be infused with emotion and humanity. Despite what cynics and some current legal systems say, creative output such as these should be protected.

Generative AI is in its infancy, yet the law must catch up and adapt to new forms of creative output. Current copyright laws are no longer fit for purpose given the progress of technology and the availability of data in the digital world. The legislator and the courts will, eventually, take a view, however, until then, the uncertainty is damaging both the creative industry and users of creative materials. The other issue is that in a global world, the national, piecemeal, approach to copyright protection adds risk for all concerned. The regulation of copyright in AI art and other areas where art and technology intersect, should be addressed internationally.

Eloise Calder and Pierre Valentin

¹ Directive (EU) 2019/790 of the European Parliament and of the Council 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directive 96/9/EC and 2001/29/EC

² Article 1(c) of the EU Parliament Proposal defines 'foundation model' as "an AI system model that is trained on broad data at scale, is designed for generality of output, and can be adapted to a wide range of distinctive tasks"

³ *Authors Guild v. Google, Inc.*, 804 F.3d 202 (2d. Cir. 2015)

⁴ *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith*, 143 S. Ct. 1258 (2023)

⁵ *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135 (D. Del. Mar. 29, 2023)

⁶ *J.L. v. Alphabet Inc.*, U.S. District Court for the Northern District of California, No. 3:23-cv-03440 at *2

⁷ Nicola Lucchi, 'ChatGPT: A Case Study on Copyright Challenges for Generative AI Systems' (12 June 2023) *European Journal of Risk Regulation* (2023) 1-23

⁸ European Union Commission, 'Trends and Developments in Artificial Intelligence – Challenges to the Intellectual Property Rights Framework' (2020)

⁹ US Copyright Office, 'Second Request for Reconsideration for Refusal to Register *A Recent Entrance to Paradise*' (Correspondence ID 1-3ZPC6C3; SR # 1-7100387071), U.S. Copyright Off. Rev. BD. (14 February 2022), pp. 3

¹⁰ *Thaler v Perlmutter*, No. 1:22-cv-1564 (D. D. C. Aug. 18, 2023) at *4

¹¹ Chapter 300 of the *Compendium of U.S. Copyright Office Practices, Third Edition* (2021) referenced by the U.S. Copyright Office Policy Statement 'Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence' (16.03.2023)