




Artificial intelligence

Artificial intelligence (AI) applications are rapidly gaining the ability to behave as intelligent entities and to generate music, art, and other creative works. Three years ago, an AI-created work of art, *Portrait of Edmond de Belamy*, was sold at a Christie's auction for \$432,500. Also notably, the SONY CSL Research Lab has developed an AI system called Flow Machines that composes new music based on everything from the Beatles to Bach.

Whatever the metaverse is – whether an augmentation of the real world, any number of artificial virtual worlds, or both – it is certain that it will be characterized by an overlay of unfathomably vast amounts of information or “data.” A feature of that information is that it will be created and distributed from within the metaverse itself, that is, from within an environment created and imagined by a person and controlled by a particular entity (for example, the developer of a game, and increasingly any other business wanting to be present in the metaverse). But the metaverse, unlike the real world, is entirely manufactured. There will be no digital tree or cloud in the metaverse that doesn't “belong” to its creator. From the look of our avatars, to the clothes we wear and the cars we drive in the metaverse, we can expect that almost everything will be somebody's intellectual property.

AI uses machine learning technologies to review, digest, and analyze vast quantities of data to create rules of application called algorithms. Once “educated,” machine learning software can continually improve itself through the analysis of new data sources and through the observation of its own data output. More recently, AI has expanded to include computing systems that aim to replicate the function of the human brain in analyzing and processing information, called artificial neural networks, as well as pairing computer networks in generative adversarial networks where the computers learn from each other.

The massive ingestion of data by AI machines, and the works they create, have generated considerable debate. Can AI digest massive databases that include copyrighted works and use machine learning to “author” creative works without infringing on copyright? In addition, is the output generated by AI protectable under copyright?

Machine learning and fair use

As AI search engines crawl through the worldwide web endlessly seeking, digesting, and aggregating content, they inevitably digest copyrighted works such as music videos, songs, novels, and news stories. Since this digestion is frequently performed without the consent of the copyright holder, its legality depends on whether it is a permitted exception to, or outside the framework of, copyright law. Under U.S. copyright law, the exception that is most frequently relied on is “fair use.”

Under section 107 of the Copyright Act, “fair use” is a four-factor test: (1) the purpose of and character of the use; (2) the nature of the copyrighted work; (3) the amount and substantiality of the portion used in relation to the whole; and (4) the effect of the use on the potential market for, or value of, the copyrighted work. Fair use of a copyrighted work for such things as teaching, scholarship, and research is specifically permitted by section 107. A key consideration that courts have used in deciding whether fair use exists is whether the use is “transformative.”

Whether machine learning of copyrighted material constitutes fair use is a hotly debated topic that will affect the future of AI. For example, Thomson Reuters and West Publishing Corp. recently sued Ross Intelligence, Inc. over, among other things, its alleged use of machine learning to create a legal research platform for Ross from the Westlaw database. Will fair use protect machine learning?

The Second Circuit found Google Books' scanning of more than 20 million books, many of which were subject to copyright, to be a "non-expressive" and transformative fair use of the texts because Google Books enabled users to find information about copyrighted books, as opposed to the expressions contained in the books themselves. If the use of the copyrighted materials is "non-expressive" fair use, protection is likely available. As long as the AI used in machine learning is not "too smart," the mechanical digestion of copyrighted works may be permitted.

Of course, AI has evolved far beyond Google Books. AI now has the ability to learn from the way authors express ideas and to generate its own creative output. This expressive machine learning may in turn harm the market for works by human authors. The fact that AI can create outputs that mimic human expression and personalization means that AI's use of copyrighted works for purposes of machine learning may result in copyright infringement if permission has not been obtained from the owners of those works.

Training AI with metaverse content

This "intellectual property everywhere" scenario is likely to affect how we may access and re-use the data created within the metaverse.

AI and machine learning are great examples of technology whose ability to operate – given their reliance on ingesting vast amounts of data – may be hampered in an "intellectual property everywhere" scenario. Today, data and information used to train a machine learning model may or may not be subject to restrictions. Not all information is "protected" or "owned" – for example, protection is unlikely to extend to historical weather information, pollution levels, the shape of clouds, or birdsongs. In the metaverse, every birdsong is likely to be the product of a machine, coded by a human, and may thereby become protectable (for instance, the code used to write the song may be protected, or the song itself if written by a human).

This may give rise to new and fascinating legal disputes. In an "intellectual property everywhere" scenario, the use of almost any type of information in a machine learning system could likely constitute a restricted act for which authorization is required. If we consider copyright, for example, simply "reading" information should not constitute a restricted act, but acts of copying or reproduction – which are likely to take place in the real-world functioning of a machine learning system – almost certainly are unless a relevant copyright exception is shown to apply, such as the doctrine of fair use in the United States, specific machine learning exceptions in jurisdictions such as Japan, or the more limited (and highly compromised, as far as commercial operators are concerned) text/data mining exceptions in European law.

The last point raises another certainty of the metaverse. The application of fragmented and variegated national intellectual property frameworks to "international" machine learning and output distribution will be at least as complicated as they have proven to be in the context of Internet distribution of traditional content. It is certain that the jurisdictional arbitrage that has characterized the development of the Internet will be repeated in the metaverse.

Is AI-created output infringing?

Even if the creation of the AI machine learning model in and of itself is not infringing, if output generated by an AI system that has been trained on a particular type of data is substantially similar, it may be an unauthorized "derivative work" that infringes copyright in the preexisting works. For example, companies like Jukedeck, which was purchased by ByteDance and taken off the market, have used machine learning on recorded music to create algorithms that in turn create new music. Because of the potential for companies like Jukedeck to generate automated music that would hurt the market for music composed by humans (such as production music typically used in film or television), these creative outputs will almost certainly receive heightened scrutiny.

“In the future, making the metaverse a safe place for all is likely to require that every AI-generated three-dimensional gaming environment is devoid of biases, bullying, and other man-made expression of violence.”

Is AI-created content copyrightable?

AI creations are certain to constitute large parts of the landscape of the metaverse’s virtual worlds – sometimes literally, as in the case of the Azure-driven location models and maps generated in Microsoft Flight Simulator. The questions of rights and ownership in the outputs of AI systems raise their own problems.

International law espouses the human-centric concepts of personal expression, authorship, and originality as prerequisites for the existence of copyright in a creative work (and therefore for its protection and “ownership”). Those concepts break down when the link between a human author and the creative work is interrupted – most infamously in the “monkey selfie” case, where a photograph taken by a monkey was found not to enjoy copyright protection. Outputs generated purely by AI systems (which are, depending on the facts, distinguishable from works created by humans with AI assistance) challenge the norms that only contemplate human creation of copyright works. Even the UK’s unique provision governing “computer generated works,” where the person “by whom the arrangements necessary for the creation of the work are undertaken” is deemed the author, confirms the need to identify a human rather than a system as the author of a “creation.”

Likewise, traditional justifications for copyright protection, such as incentivizing creation of works or protecting the natural rights of creators, break down when the creator is a machine requiring no incentivization and having no personality.

In short, the UK legal system does not appear to welcome or accommodate creations by robots, which (currently) seem destined to fall into the category of information that is free and free-flowing. Could an AI-generated metaverse reset our world by providing a great space for the public domain and “commons” to thrive? Will an AI-generated metaverse compete with human-generated worlds in a great clash of intellectual property battles? The android’s doodle of an electric sheep may have no author and no copyright protection, but the programmer of the android may still want to license it to you.

In the United States, the primary purpose of copyright law is to promote the production of creative works by providing an economic incentive to authors through the protection of their works. This economic incentive is provided to authors for the public good, because enabling authors to be rewarded monetarily for their works will lead to the production of more creative content. As AI companies continue to invest in the technologies necessary for the machine-based production of creative works, will they be able to enjoy the economic protections of copyright?

Section 102 of the Copyright Act requires that for a work to be copyrightable, it must be “an original work of authorship fixed in any tangible medium of expression now known or later developed...” While neither the Copyright Act nor the U.S. Constitution addresses the requirement of human authorship, the courts and the Copyright Office have operated on that basis. The Copyright Office has rejected attempted registrations of works produced solely by mechanical processes, and has included the requirement of human authorship in its Compendium of Copyright Office Practices. Three years ago, the U.S. Court of Appeals for the Ninth Circuit dismissed a claim for copyright infringement based on the publication of selfies taken by a crested macaque monkey in a wildlife book on the basis that an author that was not human had no standing to sue under the Copyright Act.

This means that AI-created works will become part of the public domain when created and can be freely distributed. As it stands, this has profound implications for the development of AI-created works because the companies and investors behind the machines that produce them at present are not afforded protection under U.S. copyright law. There has been a lot of discussion as to whether U.S. copyright will evolve to afford this protection.

“Governments have used versions of the technology in criminal justice and the allocation of public services like income support.”

One argument for extending copyright protection to non-human authors is that other non-natural persons have been extended legal rights. Corporations in the United States have long been afforded the right to enter into contracts and enforce contracts to the same extent as human beings, as well as the obligation to pay taxes.

Some commentators have argued that the end user of an AI program generating creative content should be the owner of that content, using a concept of a machine-based work-for-hire doctrine, with the AI program being deemed the equivalent of a contractor who is hired by an employer to produce content owned by that employer.¹ Others have cited the creative contributions that the end user makes in directing the AI program to produce a creative work as a justification for the end user being deemed an author of the AI-produced content, viewing the AI program as a tool of the end user.²

AI as an enforcement mechanism to protect copyright

Beyond having the ability to produce creative works, machine learning also provides human authors with the ability to enforce their rights and to better monetize their rights. Companies like Audible Magic, as well as Google and YouTube, have developed AI software that recognizes content and helps detect potential copyright violations. Their technologies should yield significant economic benefits for human authors.

Should AI copyright be based on creativity?

Some countries, such as the United Kingdom, have moved toward protecting computer-generated works based on the elements of creativity contained in the work in order to encourage investment in AI systems. As AI continues to develop and generate more “creative” works, the debate over the ability to copyright these works, and who can own them, will undoubtedly grow.

Ethics

The other area of considerable interest in the sphere of machine learning and AI is that of ethical compliance of AI systems – witness the increasing number of papers and debates happening in that space.

Today, the ethical ramifications and pitfalls of AI are considered to be highly application-specific. The potential for in-built biases of the AI system to create serious consequences for human subjects are deemed very much more obvious in the context of, for example, criminal justice applications than that of an AI generator of artwork. This underlies the identification by the European Commission in its recent draft AI Regulation of “high risk” AI applications, which are to be subject to statutory standards.

In the future, making the metaverse a safe place for all is likely to require that every AI-generated three-dimensional gaming environment is devoid of biases, bullying, and other man-made expression of violence all too often experienced in our real-world environment.

When the day comes, it seems very likely to us that all AI operators – to a greater or lesser extent depending on the nature of their applications, and whether as a matter of legal compliance or commercial best practice (for example, in adhering to voluntary sector standards and benchmarks) – will need to consider their internal processes and governance with respect to the high level of safety and security that will be required to enter the building site of the metaverse.

The scope for bias in systems and outputs; the quality and nature of training data; systems resilience and accuracy; human oversight and intervention – to name but a few factors – are likely to be necessary to ensure that humans feel comfortable, safe, and at ease in the metaverse.

1. See Wenqing Zhao, AI Art, Machine Authorship, and Copyright Laws, 12 Am. U. Intell. Prop. Brief 1 (December 2020).

2. See Nina Brown, Artificial Authors: A Case for Copyright in Computer-Generated Works, 20 Sci. & Tech. L. Rev. 1 (Fall 2019).

Europe's approach to AI and the metaverse

To date, no specific EU legal framework to regulate AI and the metaverse exists. The development, deployment, and use of AI are subject to a range of horizontal laws and principles, such as on data protection and privacy, consumer protection, product safety, and liability.

Very recently, however, on April 21, 2021, the European Commission published their long-awaited proposal for a regulation on AI, aiming to turn Europe into the global hub for trustworthy AI (Proposal for a Regulation laying down harmonised rules on AI (Artificial Intelligence Act)). The proposal is the result of several years of preparatory work by the Commission, including the publication of a "White Paper on Artificial Intelligence." The vision of the Commission is to protect and strengthen fundamental rights of people and businesses while at the same time encouraging AI innovation across the EU.

Whom does the proposal apply to?

The newly proposed regulation would apply to (i) providers that place on the market or put into service AI systems, irrespective of whether those providers are established in the European Union or in a third country; (ii) users of AI systems in the EU; and (iii) providers and users of AI systems that are located in a third country where the output produced by the system is used in the EU.

What is in this proposal?

The Commission takes a risk-based but overall cautious approach to AI and recognizes the potential of AI and the many benefits it presents, but at the same time is extremely aware of the threats these new technologies pose to the European values and fundamental rights and principles.

They follow a risk-based approach that is essentially

divided into four parts:

1. Unacceptable risk: AI systems that are considered as a clear threat to the safety, livelihoods, and rights of people are generally prohibited. An unacceptable risk exists especially when systems or applications manipulate human behavior to influence the user's free will and that could lead to psychological or physical harm. For example, toys using voice assistance to encourage minors to engage in dangerous behavior would fall in this category.
2. High risk: AI systems identified as high risk are permitted, but subject to special requirements and conformity assessments. Such systems include AI technologies used in various areas that need higher protection, such as education, critical infrastructure, employment management, security components of products, law enforcement in cases of interference with people's fundamental rights, or asylum and border control management.

Just to name a few special obligations: The systems must go through adequate risk assessment and mitigation systems before being placed on the market. In addition, they have to provide a high quality of data sets, a detailed documentation about all information necessary on the system, and its intended purpose so that authorities can assess compliance. The systems must meet the requirements of transparency and information for the user and must be overseen by humans to minimize risks.

In particular, all remote biometric identification systems

are placed in this category and are subject to these strict requirements. Their live use in publicly accessible spaces for law enforcement purposes is generally prohibited. Very few strict exceptions are allowed, which must be authorized by a judicial body (for instance, when absolutely necessary to search for a missing child).

3. Limited risk: AI systems with limited risks are generally permitted but also have to fulfill specific transparency obligations. AI systems such as chatbots shall make users aware of the fact that they are interacting with a machine so that they can make an informed decision to either continue or stop.
4. Minimal risk: The vast majority of AI systems, such as video games or spam filters, fall into this category and are legally allowed as there is minimal risk or no risk at all for users' rights or safety.

What's next?

The European Commission's 108-page proposal is an attempt to regulate an emerging technology before it becomes mainstream. As the European Union has been the world's most aggressive watchdog of the technology industry, it may serve as a blueprint for similar measures around the globe.

The rules have far-reaching implications for major technology companies that have poured resources into developing AI, but also for scores of other companies that use the software to develop medicine or judge creditworthiness. Governments have used versions of the technology in criminal justice and the allocation of public services like income support. The broad definition of AI systems ensures that the regulation would have a significant impact in all industry sectors, in particular in those sectors that want to have success with the metaverse.

The proposal now goes to the European Parliament and

the Member States in the ordinary legislative procedure. Given the controversial nature of AI and the large number of stakeholders and interests involved, it seems likely that this will not be a straightforward process. There will likely be many amendments and, hopefully, also some further clarifications. Once the law is adopted and passed, the regulation would be directly applicable in all Member States in the EU.



Sophie Goossens

Partner
London
sgoossens@reedsmith.com



Jess Drabkin

Partner
New York
jdrabkin@reedsmith.com



Tomos Jones

Associate
London
tomos.jones@reedsmith.com



Thomas Fischl

Partner
Munich
tfischl@reedsmith.com