Artificial intelligence

In 2018, a painting created using artificial intelligence (AI), "Portrait of Edmond de Belamy," was sold at a Christie's auction for $432,500, while AI start-up JukeDeck composed music sung at a K-pop concert in Seoul. In 2016, Flow Machines – an AI system developed by SONY CSL Research Lab – composed new music based on everything from the Beatles to Bach. Veritone, a leader in enterprise AI software, recently partnered with the estate of Walter Cronkite to create a synthetic voice model of the iconic American broadcast journalist. Craiyon's AI text-to-image generator, which is publicly available, draws art based on word prompts.

■ "Portrait of Edmond de Belamy" by Paris-based arts-collective Obvious

Advances in technology, the development of the metaverse(s), and the expectations of today's consumers continue to propel the demand for next-level content. The considerable cost of producing high-quality, ultra-realistic artwork at a faster rate is a harsh reality for creators across many industries, including games, film, television, automotive, architecture and more. The finite amount of creators and time available to design adds another layer of challenges and causes an increasing number of industries to turn to AI assisted artistry to solve the problem of producing and scaling high-quality content.

Introduction

AI uses machine learning technologies to review, digest, and analyze vast quantities of data to create rules of application called algorithms. Once "trained," machine learning software can continually improve itself through the analysis of new data sources and through the observation of its own data output. In recent years, 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 in the legal world, from which two key questions have emerged:

1. Can AI digest massive databases that include works protected by copyright and use machine learning to "author" creative works without infringing on copyright?

2. Is the output generated by an AI system protectable under copyright laws?

Another area of increasing scrutiny in the sphere of machine learning and AI is that of ethical compliance of AI systems – as evidenced by the increasing number of academic papers and debates occurring in that space.

Training AI with data protected by copyright

Generating works using AI is a creative process that often differs from traditional computer-generation. With the latest types of AI, the computer program can make many of the decisions involved in the creative process without human intervention, thereby elevating it from the status of "tool" to that of "creator." At European policy level, considerable thought is currently being given to this particular question of AI-generated creations, as indicated in particular by the European Commission in its Communication of November 25, 2020.[1]

Separately, policy-makers continue to debate questions arising from the use of data that is protected by copyright for machine learning purposes, during the stage leading to the development of software capable of self-generating "creations."

Data and information used to train an AI system may or may not be subject to restrictions. Not all information is "protected" or "owned" – for example, protection is unlikely to extend to historical information about weather patterns, pollution levels, the shape of clouds, satellite imagery or birdsongs.

What about content protected by copyright? In any text and data mining ("TDM") process it is typically necessary to "clean" the text and data being mined (which in some cases takes up to 80 percent of the mining time), in order to remove inconsistent, unreliable or redundant data, and to "normalize" the data into a specific format adapted to the relevant application. These mining operations usually involve copyright issues because they involve upstream acts of reproduction of the works or databases concerned. In order to be "read" by an AI system, they must be stored, at least temporarily, and sometimes modified (e.g., by formatting, cutting, merging, compilation, etc.) to make them usable. Each of these copying operations is likely to engage the right of reproduction that is reserved to the relevant copyright owners, which requires the express authorization of

---

1   Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, Making the most of the EU's innovative potential – an intellectual property action plan to support the EU's recovery and resilience, 25 November 2020, available here.

those copyright owners for the exercise of those rights. In the same vein, the storage and, if necessary, the communication of copies of the initial data set to third parties without such authorization is likely to infringe the monopoly rights of those copyright owners, unless an applicable exception exists. One of the most frequently used exceptions, under U.S. law, is the doctrine of fair use. However, the U.S. law approach differs considerably in that respect from the approach adopted recently under EU law, at articles 4 and 5 of the Copyright Directive (2019-790).

The differing, patchwork approaches of different jurisdictions to TDM exceptions creates opportunities for arbitrage of national copyright laws when it comes to carrying out TDM, particularly for commercial purposes. The absence of an untrammeled TDM exception within the EU clearly has potential to encourage AI users to train their AI systems on data placed on servers in jurisdictions with clear copyright exceptions, and to create consequential effects in areas such as business structuring, investment decisions and talent retention.

Text and data mining in the United States

As AI search engines crawl through the World Wide Web endlessly seeking, digesting, and aggregating content, they inevitably digest copyrighted works such as music videos, songs, novels, and news stories. Since this digestion – which generally requires the making of a copy – is frequently performed without the express consent of the copyright holder, its legality often depends on whether it is permitted under an exception to, or outside the framework of, copyright law. Under U.S. copyright law, the exception that is most frequently relied upon 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 copying of copyrighted material for the purpose of machine learning constitutes fair use is a hotly debated topic that will affect the future of AI in the United States. For example, Thomson Reuters and West Publishing Corp. have 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. The outcome of this case is still pending, although Ross' motion to dismiss was denied.[2]

Will fair use protect machine learning?

In a seminal case from 2015, 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.[3] A key learning from the case was the distinction made between "expressive" and "non-expressive" use of copyrighted materials, the latter being deemed fair use by the court. Applied to AI, could the solution mean that so long as the original text does not "express" in the final work product, the act of machine reading is fair use?

We are not aware of U.S. courts applying fair use in the context of TDM, in part because cases considering AI functionality have often involved the express use of copyrighted material that qualified as traditional copyright infringement. For example, the Second Circuit found in a 2018 case, that although TVEyes' "search feature" for Fox News content in and of itself might have been

2   Thomson Reuters Enter. Ctr. GmbH v. ROSS Intelligence Inc., 529 F. Supp. 3d 303 (D. Del., Mar. 29, 2021).
3   Authors Guild, Inc. v. Google Inc., 804 F.3d 202 (2d Cir. 2015).

sufficiently transformative to be fair use, the fact that TVEyes also had a "watch feature" that redistributed copyrighted Fox News content to TVEyes users for a monthly fee did not permit a fair use defense (*Fox News Network, LLC v. TVEyes, Inc.*, No. 15-3885 (Feb. 27, 2018)).

In practice, major TDM search projects are generally dealt with under contract, which has resulted in low instances of litigation. Academic and commercial arguments have also been raised against over-reliance on "fair use" for TDM. As a practical matter, a key factor that U.S. courts will look at is whether TDM deprives the copyright owner of the value of their copyrighted material.

Text and data mining in the European Union (Directive 2019/790)

In Europe, the recent Copyright Directive adopted in 2019 created two TDM-specific exceptions.

1. TDM for research that focuses on TDM by research organizations and cultural heritage institutions, limited to the purposes of scientific research (art 4).

2. TDM for any purpose that applies for everyone else, but with a significant caveat: the ability for copyright holders to opt out of that exception (art 5).

The caveat allowing rights owners to opt out is significant, and could potentially place a considerable burden on the shoulders of businesses that would arguably need to verify, each time a training set needs to be copied, whether owners of the underlying copyright-protected material have opted out or not. Otherwise, businesses could inadvertently be infringing copyright.

Given that there is no incentive for rights owners not to reserve their rights, we suspect that a great number of (traditional) copyright owners will want to reserve their rights and "opt out." With regard to the manner in which rights owners could exercise their opt out, the Directive is somewhat unclear. It explains that a rights owner may only reserve those rights by the use of machine-readable means, and should be able to apply measures (e.g., technical measures) to ensure that their reservations in

this regard are respected. This raises significant questions such as: (1) the exact manner in which the opt-out must be expressed; (2) at what point the TDM user needs to check whether the opt-out has been exercised (e.g., at the time when it first accesses the data, or on a continual basis?); (3) who bears the burden of proof as between the rights owner and the user (bearing in mind the difficulty a user will have in "proving a negative," i.e., that the opt-out right has not been exercised); or (4) how to determine the period of permitted retention.

Assuming that certain types of rights owners will largely seek to exercise their opt-out rights, these new TDM exceptions are likely to provide a contrasting level of protection to businesses, depending on the type of data they use. If the data being used is likely to belong to the most traditional areas of the entertainment industry, then these exceptions may provide little support for use in commercial AI applications. The geopolitical context thereby created is one in which other jurisdictions have positioned themselves favorably in the race to become global centers for TDM and AI development, through their more developed, fit for purpose copyright exceptions.

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.[4] Outputs generated purely by AI systems (which are, depending on the facts, distinguishable from works created by humans with AI

---

4   Naruto v. Slater, 888 F.3d 418, 426 (9th Cir. 2018).

PROCESSING

"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."

BETA : 7648 34 55

ALPHA : 973 86 90 56

TD : 3765 5355 5

SCORE : 3 7 8 4 0 6 6

481 802 2829 88
39 98239 8238
JKVTT UVSU

CONTROL ANALYSIS

1239 98239 8238
JKVTT UVSU

SP   VM   BA   DA   PM

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, both the EU and the UK legal systems do 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 works to be copyrightable, they must be "original works 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.[5]

In 2018, the Copyright Office rejected Stephen Thaler's application to copyright "A Recent Entrance to Paradise," a work generated by his AI system and listed author, the Creativity Machine, on the grounds that it "lacks the human authorship necessary to support a copyright claim." The Copyright Office also rejected Thaler's claim that AI can be an author under the work-for-hire doctrine.[6]

The view of the Copyright Office is that a work generally needs to be of human authorship in order to be copyrightable, with the computer merely being an assisting instrument, and where the traditional elements of authorship (such as literary, artistic or musical expression) were conceived and executed by a human.[7] This means that AI-created works in the United States will likely 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.

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.[8]

5   "[T]he Office will refuse to register a claim if it determines that a human being did not create the work." U.S. Copyright Office, Compendium Of U.S. Copyright Office Practices § 306 (3d ed. 2021).
6   Letter from Shira Perlmutter, U.S. Copyright Office Review Board, to Ryan Abbott, Esq. (Feb. 14, 2022) (on file with the U.S. Copyright Office).
7   U.S. Copyright Office, Compendium Of U.S. Copyright Office Practices § 313 (3d ed. 2021).
8   See Wenqing Zhao, AI Art, Machine Authorship, and Copyright Laws, 12 Am. U. Intell. Prop. Brief 1 (December 2020).

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

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.[9]

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.

Is AI-created output infringing?

The fact that AI can create output that mimics human expression and personalization means that AI's use of copyrighted works for the purposes of machine learning may harm the market for works by human authors and thus come under increased scrutiny by (human) rightsholders. Even if the creation of the AI systems 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 to the data in the dataset, it may be an unauthorized "derivative work" that infringes copyright in the preexisting works, which is a scenario far more likely to unfold with small and very small datasets.

Should AI copyright be based on creativity?

Some countries, such as the United Kingdom, have moved toward protecting computer-generated works (steered by humans) 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 is deemed 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 expressions 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.

9   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

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 harmonized rules on AI, Artificial Intelligence Act).

The EU Commission's 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.

Various EU member states have already reacted to the proposed AI Act. A decision on the proposal is intended for November 2022. However, it is not yet clear whether this timeframe can be met as there are still too many topics being heavily discussed. Moreover, it also seems that there are still some gaps in data protection law, which could be a major barrier to the Artificial Intelligence Act.

To whom does the proposal apply?

The newly proposed regulation would apply to (1) 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; (2) users of AI systems in the EU; and (3) 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, and these 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 clarification. Once the law is adopted and passed, the regulation would be directly applicable in all member states in the EU.

Authors

Sophie Goossens
Partner
London
sgoossens@reedsmith.com

Jess Drabkin
Partner
New York
jdrabkin@reedsmith.com

Thomas Fischl
Partner
Munich
tfischl@reedsmith.com