nature machine intelligence

PERSPECTIVE

https://doi.org/10.1038/s42256-020-0161-x

Check for updates

Human ownership of artificial creativity

Jason K. Eshraghian

Advances in generative algorithms have enhanced the quality and accessibility of artificial intelligence (AI) as a tool in building synthetic datasets. By generating photorealistic images and videos, these networks can pose a major technological disruption to a broad range of industries from medical imaging to virtual reality. However, as artwork developed by generative algorithms and cognitive robotics enters the arena, the notion of human-driven creativity has been thoroughly tested. When creativity is automated by the programmer, in a style determined by the trainer, using features from information available in public and private datasets, who is the proprietary owner of the rights in AI-generated artworks and designs? This Perspective seeks to provide an answer by systematically exploring the key issues in copyright law that arise at each phase of artificial creativity, from programming to deployment. Ultimately, four guiding actions are established for artists, programmers and end users that utilize AI as a tool such that they may be appropriately awarded the necessary proprietary rights.

Artistic creativity has traditionally been viewed as a product of the human mind. This association is increasingly cast into doubt as artificial intelligence (AI) becomes more pervasive in everyday life and, more recently, the world of art.

As generative algorithms and robots continue to evolve, they will soon be capable of creating more than just art, namely inventions and brands. Researchers have already used machine learning for backend design flow in electronic design automation tools[1]. McKinsey predicts that AI will create US$3.5–5.8 trillion annual value in the global economy[2] in the coming years. Entrusting the protection of artificially generated works to legal frameworks that predate current generative techniques presents two key problems: (1) in the absence of clear guidance, corporations may be disincentivized in using AI-based tools to develop works where protection is not guaranteed; and (2) end users of generative AI may unknowingly infringe on the rights of other artists, leaving themselves vulnerable to liability.

Ambiguity in law may lead to innovation in AI moving to more permissive jurisdictions. A leading example is the portrait *Edmond de Belamy* (Fig. 1), which looks like a painting from the seventeenth century, but in actuality is the creation of a machine learning algorithm trained on a dataset of 15,000 portraits from the 1300s to the 1900s. The algorithm stemmed from a concept published in 2014 titled 'Generative adversarial nets' (GANs)[3], which has made the generation of synthesized data, images and audio significantly more accessible (Fig. 2). Multiple contributors built on each other's code, until it was ultimately used by French-based art collective Obvious to generate the portrait[4–6].

On 24 October 2018, the portrait was auctioned at Christie's in New York and sold for US$432,500. In this new world where brushstrokes have been replaced by lines of code, who owns the rights to AI-generated creations? And how can programmers and contributors protect their own proprietary interests? While there are legal analyses that seek to answer similar questions, none have yet considered the specific practices that machine learning engineers and data scientists undertake in developing GANs and other generative algorithms[7,8]. These practices will be essential in assisting courts with determination of proprietary rights.

This Perspective will seek to navigate and apply the present legal frameworks to AI-generated works, in a manner accessible to engineers, programmers and artists. This is achieved by considering the human tasks that enable the automation of artwork, including programming a neural network, dataset curation, training, and execution or inference[9]. How various tasks in AI-based art generation give rise to proprietary rights will be explored across various jurisdictions. For end users to confidently utilize AI-generated works, a set of four guiding principles, which consider the programmer, trainer, user and the output, will be provided to assist AI artists with being appropriately awarded the necessary proprietary rights.

Fundamentals of copyright

An AI-generated creation must first satisfy some basic requirements to be afforded protection by copyright law. The overarching principles of copyright[10] in common law systems, including the United States, United Kingdom, Canada and Australia, indicate that if the artwork is an original work of authorship fixed in a tangible medium, then it will be afforded protection. Civil law jurisdictions, including most European and Asian nations, generally do not have a fixation requirement[11]. Additionally, there is some variation in the definition of 'original' across jurisdictions. In the United States, a modicum of creativity must be present for originality to subsist in the artwork[12]. In the United Kingdom, there used to be a lower threshold requiring the exercise of skill or labour[13], but in 2009, along with the rest of Europe, they adopted the view that the work must be the author's own intellectual creation[14]. How these thresholds are harmoniously applied in practice is arguably not wholly settled in UK jurisprudence. If a spectrum were to exist between creativity and labour, then Canada[15] and Australia[16] would be placed somewhere in between.

Is it possible for AI to exercise creativity, skill or any other indicator of originality? At this point in time, the capacity for AI to generate abstract or inventive thought is severely limited. Neural networks fundamentally transform a set of discrete, limited-domain input parameters into another set of discrete, limited-domain output parameters, using a set of pre-defined functions. The US Copyright Office relies on long-standing Supreme Court precedent that "copyright law only protects 'the fruits of intellectual labour' that 'are founded in the creative powers of the mind'"[17]. This does not include works generated by a machine[12]. Of course, this position may change as AI improves at solving ill-posed problems without human intervention. As Lord Briggs observed in the recent patent case of Warner-Lambert Co Ltd versus Generics, "the court

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA. e-mail: jasonesh@umich.edu

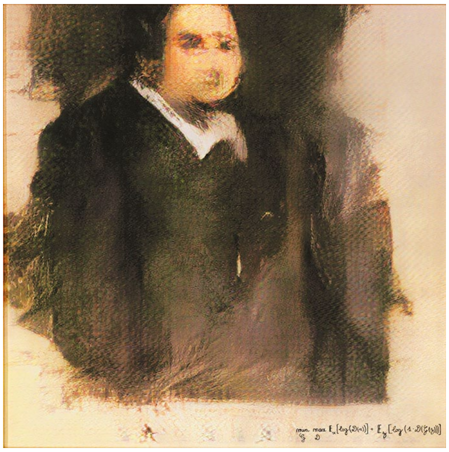

Fig. 1 | *Edmond de Belamy*. Obvious, using a modified implementation of art-DCGAN by Robbie Barrat[5], and DCGAN initially developed by Soumith Chintala[6].

is well versed in identifying the governing mind of a corporation and, when the need arises, will no doubt be able to do the same for robots"[18]. But until that time, it is unlikely that a court will find a machine to be the propagator of originality.

This does not preclude the human behind the AI from claiming authorship in the resulting output. A program cannot generate artwork without a human's initial input, and programmers are likely to exercise a degree of skill or creativity in either developing the code or procuring the training set that forms the basis of the output. In fact, the wording of legislation from countries including the United Kingdom[19], Ireland[20] and New Zealand[21], recognize this by expressly stating that computer-generated works may be copyrightable, and would belong to "the person by whom the arrangements necessary for the creation of the work are undertaken". Though in the United Kingdom's case, this comes with the added complication of the harmonizing effect of European case law from the Court of Justice, at least for the time they have been part of the European Union[14,22].

The resolution becomes less apparent where the programmer and user is not the same person, or on the extreme end of the spectrum, where a neural network is passed into a for loop to generate an endless stream of new work. The human contributions of programming, training and executing must all be factored in when determining whether the output of the AI is an original work.

The programmer

In jurisdictions where computer-generated works are explicitly recognized, it is likely that the programmer will have a strong claim to the output artwork, especially where they both trained and executed the network[23]. However, jurisdictions including the United States[24] and Australia[25] appear to take a narrower stance by providing no reference to computer-generated works, indicating that only human-generated works are subject to copyright protection. The question becomes: can AI-generated work be extrapolated to be human-generated?

Given the present state of narrow AI, the short answer is yes. If a computer can be established as a tool that enables a human author to produce an original expression, akin to a digital camera or word processor, then a neural network also acts as a tool, and the output is a result of human authorship. The use of AI, as a tool, may not be any different from other tools like a word processor or a digital camera. As such, machine-aided tasks do not require special treatment by law (in contrast to 'computer-generated' works, which are specifically legislated for), and do not require further statutory categorization[26].

On this basis, the classical black-box problem of neural networks becomes irrelevant, as understanding the operation of a tool is not a prerequisite to authorship. A photographer only needs to understand the tunable parameters (for example, shutter speed, aperture, composition and lighting) to exercise creative control over the aesthetic of their photograph. They are not required to fathom the internal mechanics of the camera. Likewise, a programmer need not understand why a neural network 'learned' a certain set of weights, or the mathematics behind a cost function. Rather, the programmer will be able to demonstrate the minimal amount of creativity required by United States copyright law by applying a basic knowledge of the parametric decisions involved in developing a GAN (training set, hyperparameters, learning rate and so on), or demonstrating that the act of coding stems from the author's own intellectual creation as required in Europe. Even with the variation in determining originality, it is possible for a programmer to demonstrate they used a machine as a tool in attaining a copyrightable result. This is achievable with version control to prove the programmer's contribution.

The trainer

Conflicting claims are likely to arise where multiple people are involved in the creation of a computer-generated work. A person who retrains a network is likely to be making some active contribution to influence the generated output, an extremely common scenario with the advent of online repositories, such as GitHub. Should a generated image belong to the programmer, the trainer or the user/curator? And at what stage does the task of retraining a network override the proprietary interests of the original programmer?

To train a network, a human must first gather data to form the training set. A training set may be randomized using a web crawler or selectively curated; common practice involves a mix of the two. The role of the trainer can be likened to that of an art curator. While a curator exercises a great degree of creative selection in setting up an exhibition, they typically cannot claim copyright in their exhibition, in common law jurisdictions. This is particularly the case where an exhibit is artefactual in nature, and relies on the individual pieces of artwork rather than the collation. It is because artefactual exhibitions are often seen as temporary or fleeting and lacking in tangible form, a prerequisite of copyrightable work. In contrast, AI-generated works are tangible in a digitized format. This suggests that selective curation by a trainer will be sufficient for satisfying the fixation requirement.

The use of a web crawler is more doubtful. Unless it can be found that programming a web crawler satisfies the jurisdiction-dependent requirements of originality that led to the creation of the output, then it may not be enough. For example, the United States demands a modicum of creativity[12], whereas under Australian law it is suggested that skill and effort are sufficient[27]. The outcome may once again be different where the trainer has deployed narrow and selective criteria, such as by filtering the images through a convolutional neural network to only include images with specific features. By narrowing the selection criteria, the trainer asserts a degree of creative control over the output[28]. Jurisdictions that use labour or effort, rather than creativity, as an indicator of originality typically require a lower threshold to satisfy, and programming a web crawler indeed requires a degree of labour. Implementing selective criteria should be catalogued by the trainer in order to satisfy the higher thresholds of creativity where necessary, be it through modification of the dataset or in the curation of the randomized output data.

The user

Once a network has been trained, it is ready to be deployed. If a user merely executes another's program, they will not be afforded proprietary rights due to the lack of creativity, labour or intellectual effort taken to click 'run'. Where the user is required to inject some personalized input, then the court will need to consider whether this input constitutes an original work of authorship[29,30].

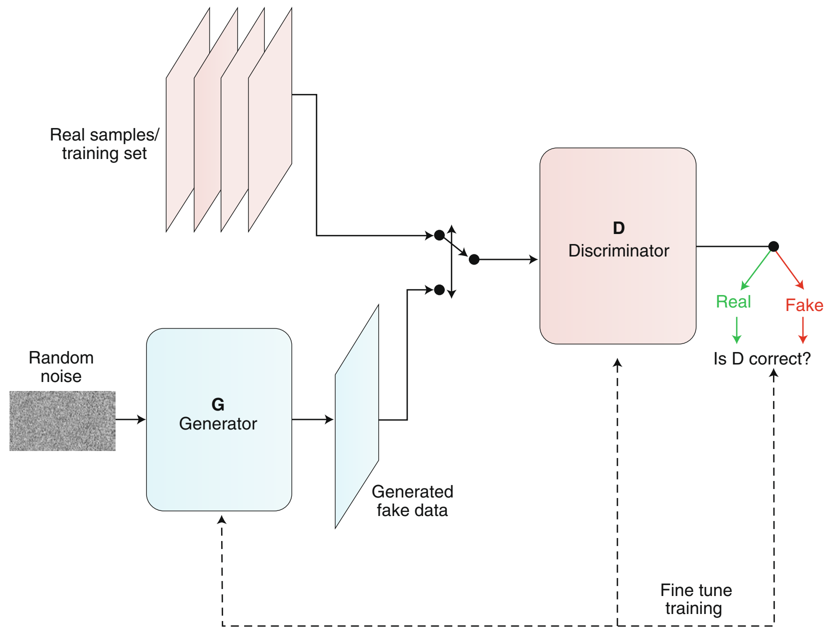

Fig. 2 | Generative adversarial networks. GANs are composed of two distinct networks. For image generation, one network is tasked with producing the image and is called the 'generator'. It accepts random noise as input, and gradually 'learns' how to convert the noise into output images. The other network distinguishes between real and generated images, and is thus called the 'discriminator'. The goal is to train these two networks competitively, such that the generator creates data in a way that the discriminator can no longer tell real and synthesized images apart. This adversary between the two networks is what teaches the generator to create realistic data.

Some past case law concerning user-based input provides helpful guidance here. Where the player of a video game makes an audiovisual recording, it was determined that the player is restricted by the decisions of the game designer and therefore lacks enough creative input[31]. Therefore, the player could not be the author of the recording. In contrast, when Walt Disney used the motion capture program Mova to retarget actor's faces onto CGI characters (used in *Deadpool* and *Beauty and the Beast*), the courts found that the users of the program had performed the "lion's share of creativity"[32]. In this case, the program requires camera footage as input. This camera footage contains within it an unbounded number of creative selections on part of the film crew, such as the actor's performance, facial expressions and lighting setups. These decisions are not delimited by the programmer of Mova, unlike the audiovisual output of a video game[33].

These cases indicate that where a user's influence is bound by the decisions of a programmer, it will be difficult for the user to satisfy the requirements of originality, even in jurisdictions that do not require creativity. If all a user has to do is click a button, this, on its own, is unlikely to be deemed as labour fostering the generation of an original work.

These principles can be carried forward in an AI landscape. For example, NVIDIA's demonstration of the first interactive AI-generated virtual world enables a user to navigate a virtual environment rendered with a GAN[34]. On one hand, the generated graphics are no longer premeditated by the game designer. On the other hand, the contribution of the user is no different to a non-AI generated video game. Under the United States "lion's share of creativity" test, it is improbable that the user will have a stake in ownership as the programmer still dictated how the visuals of a game were to be constructed. The user merely navigates the world designed by the programmer.

Is it possible for a different outcome in a different jurisdiction? In the United Kingdom, Lord Oliver states that the "labour must be of the right kind", which suggests that there is an emphasis placed on quality rather than quantity[35]. Although labour alone is no longer used to find originality in the United Kingdom, when comparing the skilled contributions between the programmer and the user, it is almost certain that the programmer is more likely to have provided the required quality of labour, rather than the user. However, Lord Oliver's obiter statements have been limited to their particular context of copying design drawings, and the Court of Appeal have said this notion should not be broadly applied[36]. Rather, they affirmed that labour, skill and effort can warrant copyright protection. Therefore, the programmer should make an active decision on delimiting user and trainer input that may be demonstrated by the form in which the tool is made available—for example, a custom interface to enable use by non-programmers[37].

Conversely, transfer learning is more likely to see success in favour of the user. This is a machine learning technique where a model trained on one task is repurposed for a different task. For example, using a GAN developed by a programmer to generate images of dogs, repurposed by a user to instead produce images of cats. Here, the user is not bound to the decisions of the programmer and can apply the network to an infinitely broad range of tasks, akin to the use of Mova mentioned above.

Guiding principles

New artists are often influenced by existing work, and the line between inspiration and appropriation can become unclear. When a network is trained on data from an open, Internet-based body of knowledge, one possible school of thought suggests that copyrighting the output is akin to claiming a proprietary interest in the Internet itself. While the law varies across jurisdictions, in general, copyright infringement can occur where (1) an output is probatively similar to the source material unless it is under fair use (US)[38,39], or (2) a substantial part of the source material is used without consent or a relevant defence (Australia and UK)[40,41]. International copyright law stipulates that adaptations of artistic work shall be protected[42], but the point at which the output becomes 'original' is a fact-specific and jurisdiction-relevant inquiry. This is codified in the US[43], where a work may still be considered original, even if derived from pre-existing materials[44]. It depends on how similar the derivative is to the source[45]. Risk-averse programmers who wish to avoid any uncertainty should rely solely on public domain works, obtain a license to the training set or conduct a reverse image search on the output, which can assist as a quantifiable metric for determining similarity to data in the training set.

On the note of licensing, a common misconception is that open source code implies it is free for use in the public domain. This is not necessarily the case, unless a highly permissive open source licence is attached to the project. Many open source licences are more restrictive, in that the original author of a program may influence how their code may be used, modified and redistributed. The type of licence used will ideally be determined by how the author wishes their code to be used. While some may wish to be attributed in the resultant work, others may view such conditions as restrictive and disincentivizing to downstream contributors. Licensing can very quickly become a complex area, and GitHub assists in navigating some of the most popular licenses by maintaining the Open Source Guide[46].

In addition to being attentive to the licences of their project dependencies, AI artists should document the full creative process from programming to training and parameterizing. It is also important to assess whether the resultant work is sufficiently transformative to mitigate any potential infringement claims, should the work be displayed or sold. The above analysis can be summarized into a set of four guiding principles that any contributor to AI-generated creations should follow:

Programmer: document the full creative and technical process, license software appropriately when making a repository public, and make an active decision on bounding or delimiting user/trainer input.

Trainer: maintain a catalogue of the dataset and its mode of procurement, and document the training process, noting that the more selective the trainer's criteria, or the more distinguishable it is to data from the training set, the more likely it is their contribution will satisfy the indicators of originality.

User: catalogue all runs of the program, as this may indicate selectivity and an element of curation, including any user-based input that was required, such as hyperparameterization.

Output: ensure the generated work does not infringe on the rights of others.

Version control has made the documentation process almost trivial, and reverse image searches can assist with ensuring a generated work invokes no infringements. These principles will only become of increasing importance as AI goes beyond objective artistic skill and towards abstract inventiveness.

Of course, these principles are only applicable for as long as AI does not have legal personhood, the way humans and corporations are afforded[18].

Received: 18 September 2019; Accepted: 11 February 2020;
Published online: 9 March 2020

References

1. Lin, Y. & Pan, D. Z. in *Machine Learning in VLSI Computer-Aided Design* (eds Elfadel, I. M., Boning, D. S. & Li, X.) 95–115 (Springer, 2018).
2. Chui, M. et al. *Notes from the AI frontier: Insights from Hundreds of Use Cases* (McKinsey Global Institute, 2018).
3. Goodfellow, I. et al. Generative adversarial nets. In *Advances in Neural Information Processing Systems* 2672–2680 (NeurIPS, 2014).
4. Radford, A., Metz, L. & Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. In *Int. Conf. Learning Representations* (ICML, 2015).
5. Barrat, R. *art-DCGAN* (GitHub, accessed 18 January 2019); https://github.com/robbiebarrat/art-DCGAN
6. Chintala, S. *DCGAN* (GitHub, accessed 10 October 2019); https://github.com/soumith/dcgan.torch
7. Aplin, T. & Pasqualetto, G. *Regulating Industrial Internet through IPR, Data Protection and Competition Law* (eds Ballardini, R. M. et al.) Ch. 5 (Wolters Kluwer, 2019).
8. Guadamuz, A. Do androids dream of electric copyright? Comparative analysis of originality in artificial intelligence generated works. *Intell. Prop. Q.* **2**, 169–186 (2017).
9. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
10. Berne Convention for the Protection of Literary and Artistic Works: Paris Act of 24 July 24 1971, as amended on 28 September 1979, Article 2, paragraph 2.
11. Sterling, J. A. L. *World Copyright Law* 198–201 (2018).
12. *Feist Publications, Inc. v. Rural Telephone Service Co.* 499 US340 (1991).
13. *Ladbroke (Football) Ltd v. William Hill (Football) Ltd* 1 WLR 273 (HL) (1964).
14. *Infopaq International A/S v. Danske Dagblades Forening* C-5/08 (2009).
15. *CCH v. Law Society of Upper Canada* 1 SCR 339, 2004 SCC 13, 236 DLR (4th) 395, 30 CPR (4th) 1, 247 FTR 318 CHH (2004).
16. *Telstra Corporation Ltd v. Phone Directories Company Pty Ltd* FCAFC 149 (2010).
17. *Compendium of U. S. Copyright Office Practices* 3rd edn § 306 (US Copyright Office, 2017).
18. *Warner-Lambert Co Ltd v. Generics (UK) Ltd* (2018).
19. Copyright, Designs and Patents Act 1988 (UK), s9(3).
20. Copyright and Related Acts, 2000 (Ireland), s.21(f).
21. New Zealand Copyright Act 1994 s5(2)(a).
22. *SAS Institute Inc. v. World Programming Ltd* EWCA Civ 1482, 2014. RPC 8 (2013).
23. Samuelson, P. Allocating ownership rights in computer-generated works. *U. Pitt. L. Rev.* **47**, 1185–1228 (1985).
24. Copyright Act of 1976 (USA) s.503.03(a).
25. Copyright Act 1968 (Australia) s.32.
26. Perry, M. & Margoni, T. From music tracks to Google Maps: who owns computer generated works? *Comp. Law Security Rev.* **26**, 621–629 (2010).
27. *Information Sheet G089v4: Music: DJs* (Australian Copyright Council, 2014).
28. Hugenholtz, P. B. in *Kritika: Essays on* Intellectual Property (eds Ulrich, H. et al.) Ch. 3 (Edward Elgar, 2018).
29. Gatys, L. A., Ecker, A. S. & Bethge, M. Image style transfer using convolutional neural networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 2414–2423 (IEEE, 2016).
30. Isola, I., Zhu, J. Y., Zhou, T. & Efros, A. Image-to-image translation with conditional adversarial networks. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition* 1125–1134 (IEEE, 2017).
31. *Nova Productions Ltd v. Mazooma Games Ltd* RPC 379 (UK) (2006).
32. *Rearden, LLC v. Walt Disney* no. 17-cv-04006-JST, 2018 WL 1000350 (2018).
33. Ginsburg, J. C. & Budiardjo, L. A. Authors and machines. *Berkeley Technol. J.* **34**, 343–448 (2019).
34. Wang, T.-C. et al. Video-to-video synthesis. In *Advances in Neural Information Processing Systems* 1152–1164 (NeurIPS, 2018).
35. *Interlego AG v. Tyco Industries Inc* AC 217, 263 (1989).
36. *Sawkins v. Hyperion Records Ltd* EWCA Civ 565, 2005. 1 WLR 3281 (2005).
37. Simon, J. *Artbreeder* http://www.artbreeder.com (2019).
38. *Positive Black Talk Inc. v. Cash Money Records Inc.* 394 F.3d 357 (2003).
39. Sobel, B. Artificial intelligence's fair use crisis. *Colum. J. Law Arts* **41**, 45 (2017).
40. Copyright Act 1968 (Australia) s.14.
41. Copyright, Designs and Patents Act 1988 (UK), s16(3)(a).
42. Berne Convention for the Protection of Literary and Artistic Works, Paris Act of 24 July 1971, as amended on 28 September 1979 Article 2, paragraph 3.
43. U.S. Code § 106.
44. *LB (Plastics) Ltd v. Swish Products Ltd* RPC 551 at 56 per Whitford J (1979).
45. de Castilho, R. E., Dore, G., Margoni, T., Labropoulou, P. & Gurevych, I. A legal perspective on training models for natural language processing. In *Proc. Eleventh International Conference on Language Resources and Evaluation* 1267–1274 (ACL, 2018).
46. The legal side of open source. Section 5. *Open Source Guides* https://opensource.guide/legal/#which-open-source-license-is-appropriate-for-my-project (2019).

Acknowledgements

Thanks to X. Feng, G. Zarrella and G. Cohen, for their invaluable discussions in the development of this manuscript, and the National Science Foundation who supported the organization of the Telluride Neuromorphic Cognition Engineering Workshop, which inspired this Perspective.

Competing interests

The author declares no competing interests.

Additional information

Correspondence should be addressed to J.K.E.

Reprints and permissions information is available at www.nature.com/reprints.