Does black-box machine learning shift the US fair use doctrine?

Yangzi Li*

1. Introduction

Machine learning (ML) has been widely used in not only daily life (eg online platforms) but also in professional fields (eg medical care, astronomy). However, it is noteworthy that most ML confronts a common *Black-box Problem*,[1] which is deemed as one of the great policy issues with many ML.[2] In Bathaee's words, the *Black-box Problem* is defined as 'an inability to fully understand an AI's decision-making process and the inability to predict the AI's decisions or outputs.'[3] From a computer scientists' standpoint, the Black-box 'is an algorithm that takes data and turns it into something' and often 'detects patterns without being able to explain their methodology'.[4] Put differently, the black-box decision model produces output without explaining why, because neither the stakeholders nor expert data scientist is capable of understanding the model.[5] Therefore, the two main traits of the black box problem are unpredictability and unaccountability. From a legal perspective, Rieder and Simon opine that the opaque algorithms impede the regulatory bodies to determine whether the ML accesses the protected data reasonably and securely.[6]

The author

- Yangzi Li is a PhD Candidate in Faculty of Law, The Chinese University of Hong Kong. The author would like to thank Professor Jyh-An Lee for his supervision and support and also thank the Editor and anonymous reviewers for their insightful comments and suggestions on this work.

This article

- Machine learning (ML) has become one of the most eye-catching AI technologies in generating creative output. However, it is unable to know why and how the machines make such creative decisions. In other words, there is a black box inside the ML algorithms. Even though the mysterious black box problem engenders opacity in the algorithms, such algorithms have remarkable performance in creative industries because of robust training data sets. In the context of US copyright law, these data sets, which assembling the existing copyrighted works, would implicate copyright infringement without licenses or permission under a general limitation, namely, fair use. Will the creative ML with black box problem shift the current fair use doctrine?
- This article argues that the black box problem shifts the conventional fair use doctrine by breaking the balance between the rights of copyright holders and public interest and augmenting the uncertainty of fair use determination. To make the point, it first focuses on exploring and discussing three creative uses of black-box ML in creative industries and then discerning their nexus with the profound and well-established fair use doctrine.
- Overall, this article aims to fill the gap in the current literature on the relationship between fair use doctrine and black-box ML and make a contribution by providing thoughts for further copyright reform in the digital age.

* Email: audrey021@link.cuhk.edu.hk.

1 Yavar Bathaee, 'The Artificial Intelligence Black Box and the Failure of Intent and Causation' (2018) 31(2) Harvard Journal of Law & Technology 890, 901–05.

2 Stephen McJohn and Ian McJohn, 'Fair Use and Machine Learning' (2020) 12 Ne UL Rev 99, 138.

3 Bathaee (n 1) 905.

4 Matthew Stewart, 'Guide to Interpretable Machine Learning: Techniques to Dispel the Black Box Myth of Deep Learning'(*Medium*, 20 March 2020) <https://towardsdatascience.com/guide-to-interpretable-machine-learning-d40e8a64b6cf> last accessed 23 July 2021.

5 Dino Pedreschi and others, 'Meaningful explanations of black box AI decision systems' (2019) 33(1) AAAI 9780, 9780 <https://www.aaai.org/ojs/index.php/AAAI/article/view/5050> last accessed 23 July 2021.

6 Gernot Rieder and Judith Simon, 'Big Data: A New Empiricism and its Epistemic and Socio-Political Consequences' in Wolfgang Pietsch, Jörg Wernecke and Maximilian Ott (eds), *Berechenbarkeit der Welt?* (Springer 2017) 85–105.

doi:https://doi.org/10.1093/jiplp/jpab118
Advance Access Publication 6 September 2021

To ensure the lawful use of training data, some studies have explored whether the training data with ML falls within the flexible fair use doctrine recognized by s 107 of the 1976 US Copyright Act.[7] As Lemley stated, 'Given the doctrinal uncertainty and the rapid development of ML technology, it is unclear whether machine copying falls within the scope of fair use.'[8] In fact, the issue of whether the open-ended fair use can accommodate new technologies has been in intensive debate for decades. Some scholars hold positive prospect that the flexible fair use can survive and gradually evolve to fit technological change.[9] Others, however, criticize that the fair use is incompatible with international copyright regime and might be overwhelmed by digital technology.[10] Therefore, it is notable to explore whether the black box problem of ML shifts the fair use doctrine. In other words, does the opacity of ML hinder the application of fair use, which may exempt machine copying from copyright infringement?

This article focuses on investigating the application of fair use doctrine towards black-box ML. The second part of this article then illustrates the current application of the black-box ML in copyright-related industries and its correlation with fair use issues. Afterward, section 3 illuminates how the fair use doctrine function in the case of the black-box ML by analysing the research and practices in the digital age in recent years. More importantly, it is noteworthy whether the 'four-factor test' works well when facing the black-box ML scenarios. Lastly, the article draws a conclusion and offers some further thoughts on refining fair use doctrine in the ML-driven era.

2. The creative uses of black-box ML and their fair use issues

Before figuring out the application of black-box ML, it is essential to clarify two types of ML. *Supervised Learning* refers to a scenario in which a training data set includes vital information that is missing in the invisible test data, and it aims to acquire the expertise to forecast the missing information for the test data.[11] For instance, one of the most popular prediction rules in ML today is neural networks.[12] Moreover, *Unsupervised Learning* describes a scenario input data without distinguishing training and test data set to retaining some summary or compressed version of that data.[13] Whatever type of ML, training data are the core of the ML process. More significant, while a machine copies a mass of data for training purposes, copyright infringement may happen if the data set is under copyright protection. Even worse, the opacity of ML process, namely the black box problem, makes such infringement less traceable. Nevertheless, there is no denying that the risk copyright infringement does exist stealthily during ML process. More significantly, such risk further intensifies the tension between the interest of copyright holders and the interest of users. To ease the tension, the US copyright regime provides an open-ended fair use doctrine.

When stepping into the digital era, fair use has been applied in various new areas, such as parody,[14] Internet publication,[15] reverse engineering,[16] text and data mining (TDM),[17] computer code.[18] The flexibility of fair use doctrine is beneficial to accommodate new means of expression and communication of works,[19] and

7 See eg Christian E Mammen and Carrie Richey, 'AI and IP: Are Creativity and Inventorship Inherently Human Activities?' (2020) 14 FIU L Rev 275, 283 (discussing whether training data could be regarded as fair use); Jessica L Gillotte, 'Copyright Infringement in AI-Generated Artworks' (2020) 53 UC Davis L Rev 2655, 2680–85 (arguing that using copyrighted works to train AI is fair use).

8 Mark A Lemley and Bryan Casey, 'Fair Learning' (2021) 99 Tex L Rev 743, 746.

9 See eg Pamela Samuelson, 'Unbundling Fair Uses' (2009)77 Fordham L Rev 2537; Pamela Samuelson, 'Possible Futures of Fair Use' (2015) 90 Wash L Rev 815; Martin Senftleben, 'The Perfect Match: Civil Law Judges and Open-Ended Fair Use Provisions' (2017) 33 Am U Int'l L Rev 231.

10 See Ruth Okediji, 'Towards an International Fair Use Doctrine' (2000) 39 Colum J Transnat'l L 75; Lawrence Lessig, *Free Culture: The Nature and Future of Creativity* (Penguin Books 2004); Jerome H Reichman and Ruth L Okediji, 'When Copyright Law and Science Collide: Empowering Digitally Integrated Research Methods on a Global Scale' (2012) 96 Minn L Rev 1362.

11 Shalev-Shwartz Shai and Ben-David Shai, *Understanding Machine Learning: From Theory to Algorithms* (CUP 2014) 4.

12 Anthony M So, 'Technical Elements of Machine Learning for Intellectual Property Law' in J-A Lee, K-C Liu and RM Hilty (eds), *Artificial Intelligence & Intellectual Property* (OUP 2021) 15. (The author gave a rough definition: 'a neural network consists of nodes (representing *neurons*) linked by arrows. Each arrow has a *weight* and connects the output of a node (i.e., the tail of the arrow) to the input of another node (i.e., the head of the arrow).'

13 Shai and Shai (n 11) 4–5.

14 See *Campbell v Acuff-Rose Music, Inc* (1994) 510 US 569 (the US Supreme Court deemed parodies as potential fair use).

15 See eg *Kelly v Arriba Soft Corp*, 336 F.3d 811 (9th Cir 2003) (the Circuit Court concluded that 'Arriba's use of Kelly's images as thumbnails in its search engine is a fair use'); *Lenz v Universal Music Corp*, 801 F.3d 1126 (9th Cir 2015) (the court holding that 'copyright holders must consider fair use in good faith before issuing a takedown notice for content posted on the Internet').

16 *Atari Games Corp v Nintendo of America Inc*, 975 F.2d 832(CAFed (Cal) 1992) (holding that 'reverse engineering of object code to discern unprotectible ideas in computer program is fair use.').

17 See eg *Authors Guild, Inc v HathiTrust*, 902 F.Supp.2d 445(stated that 'the search capabilities of the [HathiTrust Digital Library] have already given rise to new methods of academic inquiry such as text mining').

18 *Google Inc v Oracle Am, Inc* (2021) 593 US <https://www.supremecourt.gov/opinions/20pdf/18-956_d18f.pdf>.

19 Jane Ginsburg, 'Fair Use in the United States: Transformed, Deformed, Reformed?'(2020) Sing JLS 265, 267.

therefore encourage the creation of new technologies. Will the fair use also be a fit for black-box ML in creative industries? This section, therefore, introduces several creative applications of black-box ML, especially deep learning, in music composition, literature creation and facial recognition, and further demonstrates the underlying fair use issues.

2.1. Black-box ML in music composition

In 2020, OpenAI released the latest AI-generated music model called 'Jukebox'. Jukebox is a neural network which 'generates music, including rudimentary singing, as raw audio in a variety of genres and artistic styles'.[20] According to the project team's statement, Jukebox is trained by a vast data set containing 1.2 million songs paired with the corresponding lyrics and metadata from LyricWiki[21] and is finally able to generate music samples.[22] Similarly, MuseNet, also created by OpenAI in 2019, is an online tool using deep neural networks that are trained on a data set of MIDI files collected from a mass of online sources that cover multiple styles of music.[23] Moreover, this unsupervised technology is capable of generating 4-min musical compositions with 10 different instruments and can combine different music styles.[24] For example, it is realistic to compose Adele's *Someone Like You* in the genre of Bollywood. Yet, the OpenAI team admitted two main limitations of MuseNet: (i) the model can generate output deviating from the users' instrument choices[25]; (ii) the model performs less naturally when the pairing of instruments and styles is weird.[26]

Apart from the big tech companies and organizations, some artists, such as Holly Herndon and Dadabots, are also utilizing a similar tool to make their own music albums.[27] Those artists use neural networks, which are trained by data like 'an audio file, a digital image, a melody written out on sheet music, or an autocorrect suggestion for an email or text message',[28] to access 'new aesthetic paradigms beyond the limits of human expression'.[29] It is foreseeable that people without any musical knowledge can produce their own music albums. Afterward, rather than marvelling at the incredible ML magic, it is worth pondering the copyright issues falling within the black-box ML process. More specifically, machine copying, which helps assemble input—training data for ML, will lead to copyright infringement if such action is unauthorized or is not fair use. Yet, the black box problem may impede the evaluation of the nature of machine copying.

Regarding the input of ML, it is a universal view that a grand ML needs a comprehensive data set. In other words, the higher-quality and larger-quantity training data are the more functional and less biased ML can be built. That is to say, a well-built music composition model should extract a mass of music as training data. As we all know, music is one of the copyrightable objectives, and the creators own the copyright of music works. Afterward, here comes the risk of copyright infringement while the machine reproduces the music works. Yet, it is not easy to detect infringement with the new technologies in the digital age. For example, only a limited number of music could be copied 1 time in the magnetic era, when the magnetic tape recording was the lead stringer in the music recording industry. Under that circumstance, if the music recorder were reproducing copyrighted music works, the infringement would be distinct, and the courts could make an uncontested judgment. However, the introduction of digital audio files and automatic recording devices has dramatically changed the way of music reproduction. More specifically, these new techniques lower the cost of reproduction and increase the amount of 1-time copy as well. Thus, the detection of copyright infringement becomes more difficult. Whereas the machine enables an enormous quantity of copying in a short period, it is less practical, if not impossible, to acquire licenses from all the copyright owners. Hence, the tension between the interest of copyright holders and the interest of users becomes intensive.

In order to ease such interference, the US copyright regime has already been equipped with the general

20 Open AI, 'Jukebox' (30 April 2020) <https://openai.com/blog/jukebox/> last accessed 23 July 2021.

21 ibid. (The metadata includes artist, album genre and year of the songs, along with common moods or playlist keywords associated with each song.)

22 ibid.

23 OpenAI, 'MuseNet' (25 April 2019) <https://openai.com/blog/musenet/> last accessed 23 July 2021.

24 ibid.

25 ibid.

26 ibid.

27 Rob Arcand, 'The Artists Using Artificial Intelligence to Dream Up the Future of Music' (*Spin*, 4 June 2019) <https://www.spin.com/featured/ai-music-artificial-intelligence-feature-holly-herndon-yacht/> last accessed 23 July 2021.

28 ibid. ('The instruments you ask for are strong suggestions, not requirements. MuseNet generates each note by calculating the probabilities across all possible notes and instruments. The model shifts to make your instrument choices more likely, but there's always a chance it will choose something else.')

29 ibid. ('MuseNet has a more difficult time with odd pairings of styles and instruments (such as Chopin with bass and drums). Generations will be more natural if you pick instruments closest to the composer or band's usual style.')

limitation—fair use doctrine, which provides four non-exhaustive factors and aims to maintain the balance between the right of creators and public interest in accessing the copyrighted works.[30] Such flexible fair use may be a feasible approach to deal with the use of copyrighted music by ML. Therefore, it is notable whether the black box problem sharpens the uncertainty on determining fair use doctrine. For example, the black box may hide the proportion of copyrighted music that ML has used, which is an essential criterion in the evaluation of fair use. Only if the use of copyrighted music by black-box ML is fair, the risk of copyright infringement faced by ML programmer declines.

2.2. Black-box ML in literature creation

Aside from AI-generated music, it is exciting to see that the machines with AI are capable of generating literature like novels, poems, speech articles, and scripts. The usual practice is that the computer with black-box algorithms extracts a mass of literature as training data, then 'learns' from those data and finally produces new expressions of literature works. Such new technology has started applying in today's literature industry. The pivotal technique involved in ML is Natural Language Processing (NLP) , which processes and analyses a vast amount of natural language data to train a model to understand human language. And such a model 'can write like a human, but it does not have a clue what it's saying.'[31] In other words, the black box problem engenders the opacity of ML's decision-making and obfuscates whether the ML copy ideas or expression of training data. The following paragraphs introduce two examples to understand how such ML models work to generate new texts.

An example is Deep-speare, a ML program is trained to write sonnets.[32] According to the program team, Deep-speare has been trained by extracting about 2700 sonnets from the online library Project Gutenberg,[33] which includes large amounts of free eBooks mostly published before 1924. Such machine 'independently learned three sets of rules that pertain to sonnet writing: rhythm, rhyme scheme, and the fundamentals of natural language.'[34] The input by Deep-speare is less likely to infringe copyright because most of the data used falls within the scope of the public domain or copyright licenses.

Another example, ScriptBook, is a state-of-the-art ML project which aims at 'creating AI to analyze and comprehend screenplays for decision support and its algorithms are capable of interpreting and understanding a story, this strength allows programmers to carry over key elements from our decision support system to a generative AI, thereby developing an engine that is more capable of generating stories.'[35] From the technical aspect, the operation of ScriptBook algorithms contains three elements: Hindsight (data mining), Foresight (ML, NPL), and Insight (feature engineering).[36] Both data mining and ML require vast quantities of textual data, and the ScriptBook model states that its training data comes from a data set encompassing thousands of scripts that have been released in the past.[37] Yet, whether those scripts in the data set are under copyright protection is unclear. It is likely that such data set consists of both copyrighted and uncopyrighted content. Furthermore, the services offered by ScriptBook are paid[38] and are for commercial purposes.

In brief, the black box problem in these two programs is the reason why the machines produce certain types of works the way they do is unknown. Aside from the opacity of output, the potential copyright issue on the input of black-box ML is whether the use of human-created literature constitutes copyright infringement. Indeed, it is necessary to understand at the outset whether such use is authorized. If such use is unauthorized, it is important to examine whether the use is exempted by the current copyright exceptions. The existing specific exceptions provided by U.S. Copyright Act have not clarified such use. Then, the open-ended fair use may favour this use. Nonetheless, could the open-ended fair use doctrine be feasible defence for the use? For instance, the opacity in such ML algorithms may lead to ambiguity on assessing whether their use of

30 UNCTAD-ICTSD, 'Copyright: Limitations and Exceptions' in *Resource Book on TRIPS and Development* (CUP 2005) 186.

31 Karen Hao, 'The Technology Behind OpenAI's Fiction-Writing, Fake-News-Spewing AI, Explained' (*MIT Technology Review*, February 2019) <https://www.technologyreview.com/2019/02/16/66080/ai-natural-language-processing-explained/> last accessed 23 July 2021.

32 Jey Han Lau and others, 'This AI Poet Mastered Rhythm, Rhyme, and Natural Language to Write Like Shakespeare' (*Spectrum IEEE*, 30 April 2020) <https://spectrum.ieee.org/artificial-intelligence/machine-learning/this-ai-poet-mastered-rhythm-rhyme-and-natural-language-to-write-like-shakespeare> last accessed 23 July 2021.

33 The Project Gutenberg is an online library of over 60 000 free eBooks, which are mostly older literary works published before 1924. See details in <https://www.gutenberg.org> last accessed 23 July 2021.

34 See Lau and others (n 32).

35 ScriptBook, 'Man and Machine: AI as (Co)-Creator in Storytelling' (*ScriptBook Blog*, 5 November 2018) <https://blog.scriptbook.io/man-and-machine-ai-as-co-creator-in-storytelling-537e5995ea88> last accessed 23 July 2021.

36 ScriptBook, 'ScriptBook—Technology' <https://www.scriptbook.io/#!/scriptbook/technology> last accessed 23 July 2021.

37 ScriptBook, 'ScriptBook' <https://www.scriptbook.io/#!/scriptbook> last accessed 23 July 2021.

38 ScriptBook provides three types of service—basic, standard, and premium—for their users in different price. Please see <https://www.scriptbook.io/#!/pricing> last accessed 23 July 2021.

copyrighted content is 'transformative use', which belongs to the first factor of fair use. The current judicial practice has not yet provided any analysis on whether the use of black-box ML can pass the four-factor test.

2.3. Black-box ML in facial recognition and image generation

Facial recognition, facilitated by deep convolutional neural networks (CNNs), has become one of the most significant computer vision tasks and has been applied to a wide range of fields, including inter alia portrait generation, finance and criminal identification.[39] The facial recognition model is run under the decision-based black-box scenario by extracting face features in images.[40] Like most AI models, a synthetic data set is required to train the targeted model.[41] For example, ImageNet is an early online image data set, inspired by the need for more data in the image and vision research field, providing a mass of sorted images for researchers and educators around the world.[42] Notably, ImageNet clarifies that it only gives access to the images rather than owning their copyrights.[43] Therefore, the users, who want to download or copy the original images for training their ML models, may face a risk of copyright infringement if they have not got permission from the copyright holders. Afterward, several new public data sets, such as Microsoft's COCO, Google's Open Images and YouTube-8M, have come up to the stages to meet the need for feeding the emerging neural networks and other ML algorithms.[44] Yet, all of the aforementioned public image sets still do not get rid of the copyright constraints. Before the release of the open-source data sets, those giant techs should acquire copyright licences if the data is under copyright protection. Without licensing, the unauthorized uses of copyrighted contents may lead to copyright infringement. Indeed, the researchers' reproductions are partially exempted by both the open-ended fair use and designated exceptions when certain conditions are fulfilled.[45] However, what about commercial users? The current fair use doctrine is to favour those ML models for commercial purposes if such use passes the four-factor test. The black box problem, however, brings vagueness to determining whether the purpose of ML is commercial or non-commercial. This problem also influences the determination of 'transformative use', which will be discussed in the next section. Imagine an ML model is built up for research purpose initially, while it finally gets involved in commercial domain. Does such use belong to fair use?

The answer is unclear in both a legal and practical sense. For instance, a facial-recognition-based deep learning algorithm, which was trained on scanned data from the works created by the 17-century Dutch painter Rembrandt van Rijn, produced the first AI-generated 3D-printed portrait named 'The Next Rembrandt' in 2016.[46] It is said that the AI-generated painting is composed of over 148 million pixels and is based on 168 263 painting fragments from Rembrandt's oeuvre.[47] Furthermore, to train the ML algorithm, an extensive database of over 150 GB data was built by gathering and analysing the contents of Rembrandt's paintings pixel by pixel.[48] Beyond its technical breakthrough, the 'The Next Rembrandt' project has earned a media value of €12.5 million.[49] It is not deniable that the machine-generated models can create considerable commercial

39 Yinpeng Dong and others, 'Efficient Decision-based Black-box Adversarial Attacks on Face Recognition' (2019) arXiv 1904.04433, 1 <https://arxiv.org/pdf/1904.04433.pdf> last accessed 23 July 2021.

40 ibid 1–2.

41 Nicolas Papernot and others, 'Practical Black-Box Attacks against Machine Learning' (ASIA CCS '17: Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security 2017) 508 <https://doi.org/10.1145/3052973.3053009> last accessed 23 July 2021.

42 'ImageNet' <https://image-net.org/about.php> last accessed 23 July 2021.

43 ibid.

44 Arjan Wijnveen, 'How Copyright Is Causing a Decay in Public Imagesets' (*Medium*, 28 November 2016) <https://medium.com/@arjanwijnveen/how-copyright-is-causing-a-decay-in-public-datasets-f760c5510418> last accessed 23 July 2021.

45 See eg 17 USC s 107 (2018): In determining whether the use made of a work in any particular case is a fair use the factors to be considered shall include—(1) the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes; (2) the nature of the copyrighted work; (3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and (4) the effect of the use upon the potential market for or value of the copyrighted work.The fact that a work is unpublished shall not itself bar a finding of fair use if such finding is made upon consideration of all the above factors. 2019 DSM Directive (n 123) arts 3 and 4 provide specific TDM exceptions: art 3 allows TDM for the purposes of scientific research; art 4 permitts the TDM exception or limitation shall apply on condition that the use of works and other subject matter referred to in that paragraph has not been expressly reserved by their rightholders in an appropriate manner.

46 Maria Iglesias, Sharon Shamulia and Amanda Anderberg, 'Intellectual Property and Artificial Intelligence: A Literature Review' (EUR 30017 EN, Publications Office of the European Union, Luxembourg, 2019) <https://publications.jrc.ec.europa.eu/repository/handle/JRC119102> last accessed 23 July 2021; 'The Next Rembrant' <https://www.nextrembrandt.com/> accessed 23 July 2021; 'The Next Rembrant: Data's New Leading Edge Role in Creativity' (*Newsroom*, 17 June 2016) <https://thenextrembrandt.pr.co/130454-the-next-rembrandt> last accessed 23 July 2021.

47 ibid.

48 ibid.

49 Dutch Digital Design, 'The Next Rembrandt: Bringing the Old Master Back to Life' (*Medium*, 24 January 2018) <https://medium.com/@DutchDigital/the-next-rembrandt-bringing-the-old-master-back-to-life-35dfb1653597> last accessed 23 July 2021.

benefits. Is the Next Rembrandt Project for commercial or non-commercial use? Will such AI-generated artworks engender negative market effect on the original Rembrandt works? The black-box ML itself is unable to answer this question. Of course, this project is less likely to infringe copyright. The replication of Rembrandt's works is indeed not copyright infringement because his works are in the public domain. Yet, uses and reuses of copyrighted paintings for AI generating purposes require permissions from copyright holders unless specific exceptions allow such uses. Even though the flexible fair use might favour such use, it is no evidence that black-box ML has yet passed the fair use test.

2.4. Summary

This section has chosen three common applications of black-box ML in creative industries and discussed how these black-box ML algorithms bring out underestimated fair use issues. As discussed, the opacity of ML process hinders the determination of the nature and purpose of machine copying. Does the black box problem become an obstruction to extend fair use doctrine to emerging ML in the creative domains? The analysis now investigates this issue.

3. How does black-box ML affect the assessment of fair use doctrine?

Section 2 has introduced some predominant creative applications of black-box ML and figures out the potential issues related to the system of limitations and exceptions; this section illuminates whether the black-box ML can be exempted by the fair use doctrine. According to 17 USC s 107, the application of this doctrine should consider four factors:

(1) the purpose and character of the use, including whether such use is of a commercial nature or is for non-profit educational purposes;

(2) the nature of the copyrighted work;

(3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and

(4) the effect of the use upon the potential market for or value of the copyrighted work.[50]

To understand whether the use of black-box ML should be regarded as fair, it is vital to first assess how these factors favour or disfavour such use. More importantly, it is also necessary to investigate whether the black box problem hinders the evaluation of fair use.

3.1. US fair use doctrine and black-box ML

As Samuelson states, 'fair use has taken on an increasingly important role in enabling copyright law to adapt to new technological challenges not contemplated by the legislature.'[51] Through two notable cases, *Hathitrust(2014)*[52] and *Google(2015)*,[53] courts have established that conducts of web and TDM are transformative and, thus, the fair use doctrine is applicable, irrespective of whether they are conducted for commercial purposes.[54] These decisions offer guidance for further determination of whether ML's reproduction is fair use.

Yet, the answer to that question is still blurred on account of two main barriers. One of the significant barriers is the fundamental difference between TDM and ML. Specifically, TDM is run only according to the rules set by engineers, while ML is capable of learning and performing tasks beyond human expectations.[55] In other words, ML's black-box problem leads to unpredictable output, whereas TDM is more transparent and certain on the results.

The other is that the fair use doctrine is a case-by-case approach rather than a 'one-size fits all' one.[56] Though fair use is an open-ended L&E, which has adapted to a mass of fields in the digital age, this doctrine seems to be 'less available as a legal defence for end-users to assert' because of new features of some technologies.[57] It is critical that the black boxes are not only opaque but also leading to potential biases hidden in the algorithms.[58] Furthermore, the opacity

50 17 USC s 107 (2018).

51 Pamela Samuelson, 'Justifications for Copyright Limitations and Exceptions' in Ruth Okediji (ed), *Copyright Law in an Age of Limitations and Exceptions* (CUP 2017) 46.

52 *Authors Guild v Hathitrust*, 755 F.3d 87, 99 (2d Cir 2014).

53 *Authors Guild v Google, Inc*, 804 F.3d 202, 225 (2d Cir 2015).

54 Thomas Margoni and Giulia Dore 'Why We Need a Text and Data Mining Exception (But It Is Not Enough)' (2016) <https://interop2016.github.io/pdf/INTEROP-13.pdf> last accessed 23 July 2021. *Authors Guild, Inc v Google, Inc*, 954 F. Supp 2d 282, 291 (SDNY 2013); Aff'd 2015 2d Circuit; *Authors Guild v HathiTrust*, 755 F.3d 87 (2d Cir 2014); see in general the study of the US Association of Research Libraries (ARL, 2015).

55 Lemley and Casey (n 8) 753.

56 See *Campbell* (n 14) 578 (the court held that: '[T]he task is not to be simplified with bright-line rules, for the statute, like the doctrine it recognizes, calls for a case-by-case analysis'); *United States v ASCAP*, 599 F Supp 2d 415, 423(the court held that: 'courts must undertake a case-by-case analysis to determine whether a given secondary use of a copyrighted work is fair use.')

57 Ruth Okediji, 'Creative Markets and Copyright in the Fourth Industrial Era: Reconfiguring the Public Benefit for a Digital Trade Economy' (2018) ICTSD vii.

58 Pedreschi and others (n 5) 9780.

and possible biases may have an inestimable impact on the fairness and accuracy of the decisions.[59] Hence, the black box problem may become a major hurdle on the smooth evaluation of the 'four-factor' test. This section discusses the interaction between the black-box ML and fair use doctrine by analysing it factor by factor to ultimately determine whether the black-box ML shifts the evaluation of fair use doctrine.

3.1.1. Black box problem and Factor 1

As Lemley noted, 'Fair use in the machine learning context, for example, should be sensitive to the purpose of the ML system and what it eventually produces as output.'[60] In other words, it is important to clarify whether the purpose of the ML and its output are productive when assessing the first factor.

When evaluating the first factor, the core sub-factor 'transformative use' should be emphasized. Since the 1990s, Judge Leval had encouraged the courts to pay more attention to examining if the defendant's use was 'transformative' when analysing the fair use Factor 1.[61] Some courts have further stated that 'transformative use' is 'the heart of fair use inquiry'[62] and it 'forms the basis of the entire fair use analysis'.[63] A recent empirical study has shown that transformative use is still playing a vital role in an increasing number of fair use decisions.[64] Since its first occurrence in *Campbell v Acuff-Rose* (1994), a 'transformative use' was defined as one that adds 'something new, with a further purpose or different character, altering the first [work] with new expression, meaning, or message.'[65] This means that, irrespective of the purpose (commercial or non-commercial), a use is deemed fair when it is transformative use. Moreover, the term has been expanded to 'technological transformation' to 'cover technologically empowered acts of verbatim copying, insofar as these acts are able to give the secondary work a new purpose or new meaning, despite the fact that no creative input is involved.'[66] However, some courts engage in different approaches to interpret 'transformative'.[67] For example, in *Abilene Music, Inc v Sony Music Enter, Inc*, the court assessed the 'transformative' based on audience's reaction:

> [T]he question is not whether Ghostface Killah intended The Forest purely as a parody of Wonderful World, but whether, considered as a whole, The Forest "differs [from the original] in a way that may reasonably be perceived as commenting, through ridicule, on what a viewer might reasonably think" is the unrealistically uplifting message of Wonderful World.[68]

But in *Blanch v Koons*, the court adhered to authorial intent when evaluating 'transformative' by acknowledging that the defendant had a 'clear conception of his reasons for using [the photograph], and his ability to articulate those reasons, ease [the court's] analysis.'[69] Accordingly, the transformative inquiry often considers the authorial intent or audience's reaction.

All the above said, the black-box problem impedes programmers to sufficiently explain the reasons for using copyrighted works to train algorithms and to generate creative output. It is less likely to prove the legitimacy of authorial intent without adequate and clear evidence. Besides, such black box problem may further support the conclusion that the transformative inquiry should focus on readers' response rather than authorial intention.[70] In other words, the black box problem, in some vein, shifts the approaches to evaluating 'transformative'.

In addition, one of the key elements of determining 'transformative' is to understand the further purpose or different character of the use. From the perspective of ML, it seems to be less likely, if not impossible, to detect the intent of the learning activities and decision-making process when the machine has the black-box algorithm.[71] In other words, it is unclear whether the purpose expressed by the users is the same as that of computers because a human can express their intentions while the machine cannot.

59 ibid.

60 Lemley and Casey (n 8) 748.

61 Pierre N Leval, 'Toward a Fair Use Standard' (1990) 103 Harv L Rev 1105, 1111.

62 *On Davis v Gap, Inc*, 246 F.3d 152, 174 (2d Cir 2001).

63 *Hofheinz v Discovery Commc'ns, Inc*, No 00-3802, 2001 WL 1111970 (SDNY 20 September 2001) 3.

64 See eg Jiarui Liu, 'An Empirical Study of Transformative Use in Copyright Law' (2019) 22 Stan Tech L Rev 163, 163 ('More importantly, of all the dispositive decisions that upheld transformative use, 94% eventually led to a finding of fair use.'); Clark D Asay, Arielle Sloan and Dean Sobczak, 'Is Transformative Use Eating the World?' (2020) 61 Boston Coll L Rev 905, 940. ('transformative use may be eating the fair use world by playing a role in increasingly more fair use opinions.')

65 *Campbell* (n 14) 569.

66 Maurizio Borghi and Stavroula Karapapa, 'Technological Transformative Use' in *Copyright and Mass Digitization* (OUP 2013) 20. doi: <10.1093/acprof:oso/9780199664559.003.0002> last accessed 23 July 2021.

67 Laura A Heymann, 'Everything Is Transformative: Fair Use and Reader Response' (2008) 31 Col J L Arts 445, 449.

68 *Abilene Music, Inc v Sony Music Enter, Inc* (2003) 320 F. Supp 2d 84, 89–90

69 *Blanch v Koons*, 467 F.3d 244, 255 (2d Cir 2006).

70 See eg Heymann (n 67) 445–66; Samuelson, 'Unbundling Fair Uses' (n 9) 2553–55; Lyrissa Barnett Lidsky, 'Nobody's Fools: The Rational Audience as First Amendment Ideal' (2010) 2010 U Ill L Rev 799, 805–09.

71 Bathaee (n 1) 893. ('If an AI program is a black box, it will make predictions and decisions as humans do, but without being able to communicate its reasons for doing so.')

Take the aforementioned music-generated model 'Jukebox' as an example. The artists or programmers extract a mass of copyrighted music to train the machine to create a new version of Adele's *Someone Like You* in the style of Bollywood. If we only consider human intention, the conduct of such ML is probably transformative. However, there is another concern relating to whether machine intention should also be considered. If the answer is yes, the black box problem, indeed, impedes the cognition of the machine's actual purpose during the learning process. Therefore, the courts may confront difficulties in identifying transformative use while ML has a black-box algorithm. Furthermore, as the *Google (2015)* court addressed, the 'transformative use' should not be abused while understanding the factors of fair use.[72] It is also essential to analyse whether the transformative use factor can be justified under the scenario of black-box ML. Last but not least, Samuelson opines that the *HathiTrust (2014)* ruling 'is not presuming, expressly or implicitly, that non-transformative uses are unfair'.[73]

Indeed, the 'transformative use' is one of the most significant elements of fair use towards technological changes. The black-box ML, on the one hand, should satisfy the transformative test, which means the output generated by the black-box ML should be regarded as transformative use of the original copyright work. On the contrary, the uncertainty of the black box problem may, to some extent, impedes the application of 'transformative use' on consideration of the unaccountability of ML's decision-making process. It is wise to address the latter issue in the future ruling related to ML.

3.1.2. Black box problem and Factor 2

In terms of the second factor, Judge Leval illustrated that this factor should require consideration of whether the copyrighted work is 'of the creative or instructive type that the copyright laws value and seek to foster.'[74] Copyright law predominantly protects valuable copyrighted works. In accordance with *Hathitrust (2014)* and *Google (2015)*, Factor 2 shows less correlation with the reproduction conducted by new digital technologies.[75] Hence, the black box problem is less likely to impact the identification of the nature of copyrighted works, which is the key element when examining the second factor. In other words, the nature of copyrighted works does not significantly affect the assessment of whether the uses of copyrighted works by black-box ML consist of fair use. Yet, the evaluation of training data should consider whether the copyrighted contents used as training data are deserving of copyright protection.

3.1.3. Black box problem and Factor 3

The third factor considers the amount and substantiality of the portion used of the whole copyrighted work. It is noteworthy that the *Harper & Row, Publishers, Inc v Nation Enterprises*[76] demonstrated that both the quantity and quality should be considered while evaluating this factor.[77] In other words, the use of a large number of insignificant parts in the copyrighted work may be fair, whereas the use of a snippet of 'heart' of copyrighted work may not be in favour of fair use.[78] Then, *Campbell* established a reasonableness approach to analysing Factor 3.[79] More specifically, the *Campbell* court implied that the evaluation of the third factor should govern the answer to the question whether the amount the defendant used is reasonable by considering her purpose of use.[80] In addition, R Anthony Reese's empirical research has shown that the reasonableness approach is important in determining whether the use of the entire copyrighted work is fair when the technological changes make the total-copying easier and cheaper.[81] Therefore, it is essential to make sure the causation between the amount and proportion of copyrighted works used by black-box ML and the purpose of its use should fit the reasonableness approach to decide whether such use weighs against Factor 3.

72 *Authors Guild v Google, Inc*, 804 F.3d 202, 225 (2d Cir 2015) (the court stated: 'The word "transformative" cannot be taken too literally as a sufficient key to understanding the elements of fair use; it is rather a suggestive symbol for a complex thought, and does not mean that any and all changes made to an author's original text will necessarily support a finding of fair use.')

73 Samuelson, 'Possible Futures of Fair Use' (n 9) 855–56.

74 Leval (n 61) 1117.

75 See eg *Authors Guild v Hathitrust* (2014) 98 (the court held that the second factor may be of limited usefulness in the case and the analysis should focus on the other three factors.) *Authors Guild v Google, Inc* (2015) 292 (note 4 states: 'The parties agree that the second factor plays little role in the ultimate fair use determination.')

76 *Harper & Row, Publishers, Inc v Nation Enterprises*, 105 S.Ct. 2218 (Supreme Ct 1985)

77 ibid 2233–34.

78 R Anthony Reese, 'How Much Is Too Much? Campbell and the Third Fair Use Factor' (2015) 90 Wash L Rev 755, 761.

79 *Campbell* (n 14) 586.

80 ibid. See also Reese (n 78) 812.

81 Reese (n 78) 812–13. (In this study, Reese analyzed seven cases involved a variety of different uses of different kinds of works: *Authors Guild, Inc v HathiTrust*, 755 F.3d 87 (2d Cir 2014); *Swatch Grp Mgmt Servs Ltd v Bloomberg LP*, 756 F.3d 73, 78 (2d Cir 2014); *AV ex rel Vanderhye v iParadigms*, LLC, 562 F.3d 630 (4th Cir 2009); *Perfect 10, Inc v Amazon.com, Inc*, 508 F.3d 1146 (9th Cir 2007); *Kelly v Arriba Soft Corp*, 336 F.3d 811 (9th Cir 2003); *BMG Music v Gonzalez*, 430 F.3d 888 (7th Cir 2005); *A&M Records, Inc v Napster*, Inc, 239 F.3d 1004 (9th Cir 2001). Reese states that: '[A]ll of these cases suggest that, in recent years at least, new digital technologies have often enabled new uses of copyrighted works that allow or require using the entire work.')

According to Bathaee, the current causation test in every field of law may fail if an AI with black box gets involved.[82] Moreover, 'when an AI is a black box, causation doctrines, such as proximate cause, fail because the causation inquiry will focus on what is foreseeable to the creator or user of the AI.'[83] As discussed, the evaluation of Factor 3 requires examining the causation between the defendant's conduct and purpose. However, the work is used by the machine, and the reasons why the machine creates such kind of output are unknown and unpredictable. Hence, it is less likely, if not impossible, for the courts to consider the third factor through a regular approach under black-box ML cases.

3.1.4. Black box problem and Factor 4

The fourth factor, 'effect on the market', had been described as 'undoubtedly the single most important element of fair use' in *Harper& Row* ruling.[84] According to the empirical study conducted by Barton Beebe on US copyright fair use opinions, it is still commonly believed that 'the fourth factor analysis remains the most influential on the outcome of the overall test.'[85] Moreover, unlike the other factors, Factor 4 does not have any real sub-factor on account of the synthetic and dispositive nature of the four-factor analysis.[86] The *Campbell* court further opined that Factor 4 'requires courts to consider not only the extent of market harm caused by the particular actions of the alleged infringer, but also whether unrestricted and widespread conduct of the sort engaged in by the defendant ... would result in a substantially adverse impact on the potential market' for the original.'[87] However, there is still no unified threshold to evaluate the fourth factor. Put differently, the judgement on the market impact of the alleged copyright infringement is based on the court's discretion.

When it comes to Factor 4 analyses towards new technologies, the *Google (2015)* court concentrated on 'whether the copy brings to the marketplace a competing substitute for the original.'[88] The court also opined that 'the close linkage between the first and fourth factors, in that the more the copying is done to achieve a purpose that differs from the purpose of the original, the less likely it is that the copy will serve as a satisfactory substitute for the original.'[89] Then, the court recognized that Google's reproduction could result in some loss of sale instead of 'a meaningful or significant effect "upon the potential market for or value of the copyrighted work".'[90] This said, some commentators submit that even though the output created by black-box ML, like MuseNet, fulfils the 'transformative use' factor, the defendant may fail the fair use test while such output harms the plaintiff's core market.[91] Hence, the fourth factor does favour the defendant's TDM activity. However, it is notable whether the black-box ML could be a parallel case to TDM. If so, the fair use assessment in Google case may be applied to the black-box ML as well.

Furthermore, several studies have acknowledged that ML could have negative effect on the application of the fourth factor. Sobel notes that, if the extraction of copyrighted recordings as training data by AI composers, such as MuseNet, were unlimited, the market for compositions would undeniably become robotized.[92] Moreover, the so-called unrestricted expressive ML would further 'deprive authors of markets they currently exploit.'[93] Sobel thus puts forward that 'expressive machine learning shifts the balance of fair use's fourth factor because it could substitute for the individual works, on which it trains and for the authors of those works.'[94] Put differently, the expressive black-box ML may damage the market for the original works, which are utilized as training data to build up the algorithms. Besides, Pasquale submits that the evaluation of Factor 4 is indeterminate and useless when facing initial archival copies by machines in the information-overloaded era.[95] However, these scholars have not specifically observed the impact of ML's black box problem on the evaluation of Factor 4.

82 Bathaee (n 1) 922.

83 ibid.

84 *Harper & Row* (n 76) 2233.

85 Barton Beebe, 'An Empirical Study of U.S. Copyright Fair Use Opinions, 1978-2005' (2008)156 U Pa L Rev 549, 617.

86 ibid 618.

87 *Campbell* (n 14) 590.

88 *Authors Guild v Google, Inc,* 804 F.3d 202, 223 (2d Cir 2015).

89 ibid.

90 ibid 224.

91 Lemley and Carsy (n 8) 777–778. (The author explained that: 'many of the works created by systems like MuseNet will be transformative uses that society values. But it makes the fair use case closer, because the output of the ML's learning competes with the plaintiff's core market. And some purposes—say, a system designed to write a new pop song in the style of Taylor Swift—seem more substitutive than transformative, so that if they run afoul of the ever-broadening definition of similarity in music, fair use is unlikely to save them.')

92 Benjamin Sobel, 'Artificial Intelligence's Fair Use Crisis' (2017) 41 Colum JL & Arts 45, 79.

93 ibid.

94 ibid.

95 Frank Pasquale, 'Copyright in an Era of Information Overload: Toward the Privileging of Categorizers' (2007) 60 Vand L Rev 135, 151–52 (the author explains that: 'the fourth fair use factor—the effect on the market—is entirely indeterminate. A court might find that Google's failure to pay licensing fees for the right to archive the books is a grievous financial loss to the copyright holders. Or a court might find that such a licensing market is not "traditional, reasonable, or likely to be developed," and that the archival copies, standing alone, pose no threat to the commercial interests of copyright holders. Given the equivocal nature of the other three factors, the futility of the fourth factor's analysis makes fair use analysis of the initial archival copies a black box.')

As discussed in the introduction part, the black box problem leads to the inability to predict the machine-generated output. Further, such problem influences the stability of the quality of output. For example, an ML algorithm may generate a new work similar to the original human-created content at 1 time, and produce an output entirely different from the original work at the other. Faced with such situation, it is challenging to figure out whether the use, which consists of unpredictable and various expressions, does harm the market of the original work.

In addition, a recent empirical study shows that individuals are unable to accurately differentiate human-created artworks from machine-generated artworks by doing a survey experiment on Qualtrics.[96] In some vein, such finding indicates that the quality and condition of the machine-generated output are capable of mimicking those of the human-created artworks. From an economic perspective, the machine-generated artistic works are likely to rival human-created artworks, and, therefore, impact the traditional art market. However, the black box problem impedes the identification of the nature of market effect, which is brought by the secondary use conducted by ML. The opacity of decision-making process leads to difficulties in analysing whether the output has positive or negative value of market effect on the original works. While the use by black-box ML engenders negative effect on the market of the original works, such use should not be favoured by the fair use doctrine. In brief, the black box problem inside ML becomes a barrier for the courts to clarify whether the ML does harm to the copyright holders' markets and produces substitutes that replace the original works.

In general, the black box problem inside the ML algorithm impedes the interpretation and assessment of the four factors, which are traditionally used to evaluate the use of copyrighted works by human beings. The opaque and unpredictable decision-making process reveals the machine generates creative output with little human intervention. It is unclear whether the human-based fair use could be applied to actions conducted by machines. If so, the current criteria for fair use may be unable to fit the machine-based use of copyrighted works.

4. Conclusion and future direction

Training data sets is 'fuel' to the operation of ML algorithms. Yet, the gathering and curation of data for generating creative works often require using a large number of copyrighted contents. Some giant techs, like Google,[97] are able to build up various data sets by acquiring licenses from the copyright holders and open access to some of their data sets under open-source licences (eg Creative Commons licences). However, individuals and start-ups may lack the capacity to do so. More important, the opacity of the black-box ML may have negative impacts on copyright holders and their ability to enforce their own exclusive rights, and also add difficulties for courts when determining whether the ML, as users of copyrighted works, can be exempted by the existing the fair use doctrine.

In the wake of rapid technological advancement, the discussion of whether the fair use has 'gone too far' has intensified.[98] Unfortunately, the evolution of legislation always falls behind the pace of technology. On the one hand, courts should take responsibility and may apply flexible fair use to solve those problems.[99] On the other hand, fair use has been criticized for its unpredictability and uncertainty.[100] Likewise, the black-box ML also has an opacity problem, which is described as 'the person affected by a decision can hardly ever comprehend how or why a certain input of data has been categorized and produced a certain output.'[101] The black-box ML lacks transparency, and such a feature may lead to negative impacts on the enforcement of the traditional copyright rules.[102] Lastly, the double blows of indeterminacy may increase difficulties for courts and thus shift the existing

96 Harsha Gangadharbatla, 'The Role of AI Attribution Knowledge in the Evaluation of Artwork' (2021) Empirical Stud Arts 1, 1.

97 For example, Google's Project Magent has released a MIDI and Audio Edited for Synchronous Tracks and Organization (MAESTRO) data set for public use.

98 Irene Segal Ayers, 'The Future of Global Copyright Protection: Has Copyright Law Gone Too Far' (2000) 62 U Pitt L Rev 49; Samuelson (n 78) 839–853; Patricia Aufderheide and Peter Jaszi, *Reclaiming Fair Use: How to Put Balance Back in Copyright* (2nd edn, The University of Chicago Press 2011).

99 *Twentieth Century Music Corp v Aiken*, 422 US, 167, 95 S Ct, 2049 ('[t]here can be no really satisfactory solution to the problem presented here, until Congress acts.') See also *Sony Corp of Am v Universal City Studios, Inc* (1984) 464 US 417, 500, 104 S Ct 774, 818, 78 L Ed 2d 574 ('Like so many other problems created by the interaction of copyright law with a new technology, (dissenting opinion). But in the absence of a congressional solution, courts cannot avoid difficult problems by refusing to apply the law. We must "take the Copyright Act . . . as we find it," *Fortnightly Corp. v. United Artists*, 392 U.S. 390, 401–402, 88 S.Ct. 2084, 2090, 20 L.Ed.2d 1176 (1968), and "do as little damage as possible to traditional copyright principles . . . until the Congress legislates." ibid, at 404, 88 S.Ct., at 2091')

100 See eg McJohn and McJohn (n 2) 161. ('Fair use law is often considered contradictory, vague, and unpredictable.')

101 Jenna Burrell, 'How the Machine "Thinks": Understanding Opacity in Machine Learning Algorithms' (2016) Big Data & Soc'y 1 <https://doi.org/10.1177/2053951715622512> last accessed 23 July 2021.

102 See Maayan Perel and Niva Elkin-Koren, 'Accountability in Algorithmic Copyright Enforcement' (2016) 19 Stan Tech L Rev 473, 487–88 ('even if self-learning algorithms can be created to engage in case-by-case applications of legal standards, the lack of transparency remains a serious problem'); see also McJohn and McJohn (n 2) 161. ('Decision trees may be simpler than other approaches in some respects but could work on smaller data sets (addressing one of the data issues above) and provide something that machine learning problematically lacks: transparency.')

fair use doctrine. The case-by-case nature of fair use requires judges to examine the four factors based on different scenarios.[103] The opacity of black-box ML may, to some extent, increase difficulties when evaluating those factors.

Aside from the impact on fair use doctrine, the black-box decision-making ML algorithms have been applied in a wide range of areas encompassing some sensitive domains like crime prediction and image classification. Some researchers demonstrate that black-box algorithms may result in unfair or even wrong decisions on account of the inadequate and biased training data.[104] From a technological perspective, it is an emerging opinion that building explainable AI is necessary to solve the black box problem. Some commentators also evince that the opening of black box is indispensable, and the future transparency regulation on AI may benefit the legal systems[105] Additionally, some researchers suggest that a black box explanation framework should be 'model-agnostic, logic-based, both local and global, and high-fidelity.'[106] It the foreseeable that the appearance of explainable ML will mitigate the opacity of the decision-making process and make AI more responsible and human-centric. Yet, it is unclear what the impact of opening the black box in ML up would be on the copyright framework and whether the open-ended fair use should only favor the use conducted by explainable ML.

Above all, the black-box ML has shifted the balance in the current fair use doctrine. Thus, it is urgent to refine the four-factor test to strike that balance back. The evolution of the copyright regime is an on-going mission and is impacted significantly by the rapid development of ML. It is essential to consider the unique technical features of ML while reconfiguring the fair use doctrine.

103 Maurizio Borghi and Stavroula Karapapa (n 61) 22. ('Since there is no mandate that the four fair use factors should weigh equally, the outcome of which facts of each case will be stressed, and which will not have a determinative value, will depend upon judicial examination.')

104 Pedreschi and others (n 5) 9780; see also Dino Pedreschi and others, 'Open the Black Box Data-driven Explanation of Black Box Decision Systems' (2018) arXiv 1806.09936. <https://arxiv.org/pdf/1806.09936.pdf> last accessed 23 July 2021.

105 Thomas Wischmeyer, 'AI and Transparency: Opening the Black Box' in Thomas Wischmeyer and TimoRademacher (eds), *Regulating Artificial Intelligence* (Springer 2020) 4, 23.

106 Pedreschi and others (n 5) 9782.