

Artificial Intelligence and Intellectual Property

Edited by

JYH-AN LEE, RETO M HILTY,
AND KUNG-CHUNG LIU

Author Copy – Subject to Licence – 02/03/2021

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© The several contributors 2021

The moral rights of the authors have been asserted

First Edition published in 2021

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Crown copyright material is reproduced under Class Licence
Number C01P0000148 with the permission of OPSI
and the Queen's Printer for Scotland

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2020944786

ISBN 978-0-19-887094-4

DOI: 10.1093/oso/9780198870944.001.0001

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

The variety of these proposals is appropriate: different applications of AI pose distinct harms that demand distinct legal responses. But despite wreaking dramatically different harms through dramatically different mechanisms, worrisome applications of AI frequently share a common feature: they rely on unauthorized uses of copyrighted ‘training data.’⁶ In major jurisdictions like the US, very few of these applications of AI to copyrighted data are on firm legal ground.⁷ This uncertainty jeopardizes the development of artificial intelligence technology. It jeopardizes the rights of deserving potential plaintiffs and over-deters law-abiding potential defendants. And if left unclarified, this uncertainty also threatens to undermine the purpose and administrability of copyright law.

To the extent lawmakers and commentators conceive of AI as a challenge for intellectual property (IP) law, the focus is often on issues distinct from protected training data, such as copyright in computer-generated works and potential IP protections for algorithms, software, or trained models.⁸ Those that do examine training data, in turn, typically characterize today’s legal uncertainties as a deficiency in copyright’s exceptions and limitations.⁹ If we could only strike the right balance in our systems of exceptions and limitations, the thinking goes, we could resolve our current predicament.

This chapter argues that, in fact, the current predicament is a product of systemic features of the copyright regime that, when coupled with a technological environment that turns routine activities into acts of authorship, have caused an explosion of media subject to broad, long, and federated ownership claims. Thus, equilibrating the economy for human expression in the AI age requires a solution that focuses not only on exceptions to *existing* copyrights, but also on the doctrinal features that determine the ownership and scope of copyright entitlements *at their inception*.

Because the most pressing issues in AI frequently implicate unauthorized uses of copyrighted training data, this chapter taxonomizes different applications of machine learning according to their relationships to their training data. Four categories emerge: (1) public-domain training data, (2) licensed training data, (3) market-encroaching uses of copyrighted training data, and (4) non-market-encroaching uses of copyrighted training data.

Analysing AI in this way illuminates a conundrum. Copyright is paradigmatically an economic entitlement. Thus, copyright primarily regulates

⁶ Anthony Man-Cho So, Chapter 1 in this volume.

⁷ Benjamin LW Sobel, ‘Artificial Intelligence’s Fair Use Crisis’ (2017) 41 Columbia Journal of Law & the Arts 45, 66–7 (hereafter Sobel, ‘Artificial Intelligence’s Fair Use Crisis’).

⁸ Jyh-An Lee, Chapter 8 in this volume.

⁹ 2019 OJ (L 130/92) 94 (‘In certain instances, text and data mining can involve acts protected by copyright, by the sui generis database right or by both, in particular, the reproduction of works or other subject matter, the extraction of contents from a database or both which occur for example when the data is normalised in the process of text and data mining. *Where no exception or limitation applies, an authorisation to undertake such acts is required from rights holders*’) (emphasis added).

market-encroaching uses of data—that is, uses of copyrighted expression that endanger the market for that very expression, rather than for some non-expressive aspect of a source work.¹⁰ However, market-encroaching uses represent just a narrow subset of AI applications. Moreover, copyright’s economic focus makes it a poor fit for redressing some of the most socially harmful uses of copyrighted materials in AI, like malicious deepfakes. Even more paradoxically, copyright’s property-style remedies are ill-suited to addressing market-encroaching uses, and are in fact much more appropriate remedies for the categories of AI that inflict dignitary harms that fall outside copyright’s normative mandate. In other words, copyright’s property-style remedies are inappropriate for the AI applications that copyright can appropriately regulate. Meanwhile, those remedies *are* appropriate for the AI applications that it would be substantively *inappropriate* for copyright to regulate. Identifying this mismatch helps explain why some commentators have been eager to apply copyright where it does not belong, and eager to dismiss copyright where it rightly governs.

Finally, this chapter discusses a variety of remedies to the ‘AI problems’ it identifies, with an emphasis on facilitating market-encroaching uses while affording human creators due compensation. It concludes that the exception for TDM in the EU’s Directive on Copyright in the Digital Single Market represents a positive development precisely because it addresses the structural issues of the training data problem that this chapter identifies. The TDM provision styles itself as an exception, but it may in fact be better understood as a formality: it requires rights holders to take positive action to obtain a right to exclude their materials from being used as training data. Because of this, the TDM exception addresses a root cause of the AI dilemma rather than trying to patch up the copyright regime post hoc. The chapter concludes that the next step for an equitable AI framework is to transition towards rules that encourage *compensated* market-encroaching uses of copyright-protected training data. Such rules could offer industry a less risky strategy for expansion, facilitate remunerated uses that transaction costs might otherwise impede, and compensate rights holders more proportionally than the outsized, property-style remedies that copyright affords in other contexts.

2. Diagnosing AI’s Copyright Problem and Copyright’s AI Problem

Insofar as lawyers and scholars treat AI as a problem for copyright law to respond to, they often focus on the inadequacy of copyright’s existing exceptions and limitations. However, tensions between copyright entitlements and AI methodologies

¹⁰ See Section 3.3.

are better understood as the results of *systemic* features of global copyright regimes. Machine learning's methods and values clash with broad copyright entitlements that automatically protect even the most banal exertions of human creative effort. Thus, it is inappropriate to blame the friction at the AI-copyright interface on an inadequate exception or limitation to copyright, without interrogating the larger reasons for that safety valve's inadequacy. Accordingly, solutions that present themselves as post hoc safety valves may not fully rectify the problems they purport to address.

Tension between copyright law and AI is in fact a consequence of several interlocking phenomena that should be analysed distinctly from copyright's exceptions and limitations. The root of many copyright-and-AI concerns is the proliferation of bloated copyright entitlements. This phenomenon, in turn, stems from low legal thresholds for protectable originality coupled with information technology that turns routine communications into acts of authorship. In other words, AI has a copyright problem: machine learning may often entail nominal violations of thousands of different copyrights. At the same time, *copyright* has an AI problem: now that nearly every scrap of digital expression is copyright-protected *ab initio*, machine learning technology exposes the practical and theoretical shortcomings of a copyright regime that combines strong rights, low originality standards, and no formalities. This section tracks how copyright law has developed to protect vast amounts of online media and encumber AI development. It also explains why, despite a proliferation of copyright entitlements, copyright does not protect some of the highest-stakes training data.

2.1 Everyone Owns Something and Someone Owns Everything

Most of the data that train AI are banal. Indeed, some of these data are just 'data exhaust', the information logged as by-products of technology usage.¹¹ This kind of data is not copyrightable. No matter how valuable they are to tech companies, the movements of a cursor on a webpage, or geolocation data, or the patterns in one's Netflix browsing, are not works of authorship. But just as banality does not preclude commercial value, neither does banality preclude copyrightability. Minimally creative emails or chat messages are copyrightable literary works, and it is these messages that train state-of-the-art text-generation AI.¹² The digital photographs that train image recognition networks are almost always protectable, even if they're throwaways. As a result, a great deal of the information that powers state-of-the-art

¹¹ Shoshana Zuboff, *The Age of Surveillance Capitalism* (PublicAffairs 2019) 68.

¹² Anjali Kannan and others, 'Smart Reply: Automated Response Suggestion for Email' in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16*, New York, NY, USA, ACM 2016).

AI is *owned* by ordinary Internet users. Widespread ownership of training data reflects an unexpected convergence of trends that unfolded over the twentieth century: the abolition of formalities for copyright ownership, the lowering of originality requirements for copyright protection, and the proliferation of technologies that fix expressive activity that until then had been unfixed.

2.1.1 Formalities

Copyright protects so much training data today because copyright now vests automatically in a creator. A century or so ago, complying with numerous formal requirements was essential to copyrighting a work in the US. These hurdles arguably helped allocate copyright entitlements to valuable works, while avoiding the social costs of protecting works whose creation was not incentivized by economic rights.¹³ An American novelist in the year 1910 would have forfeited her copyright if, eg, she published her work without a copyright notice.¹⁴ In contrast, the Berne Convention provides that the enjoyment of copyrights ‘shall not be subject to any formality.’¹⁵ Beginning with the Copyright Act of 1976, the US shifted away from formalities, a trend that continued through its accession to the Berne Convention in 1989.¹⁶ US copyright law no longer requires notice or registration as prerequisites to a valid copyright, and foreign authors need not even register their works before bringing suit for infringement.¹⁷ Such ‘unconditional’ copyright means that, by default, anyone who authors a minimally creative email, text message, or photograph becomes its legal owner for decades to come.¹⁸

2.1.2 Low originality standards

The second trend is international gravitation towards a relatively low standard for ‘originality’ in copyright law.¹⁹ Leading up to the twentieth century, American copyright cases recited fairly demanding originality requirements: a work needed to manifest an author’s ‘original intellectual conceptions’ to enjoy protection.²⁰

¹³ To some degree, of course, this system operated at the expense of creators who lacked the sophistication to comply with formalities. Christopher Sprigman, ‘Reform(Aliz)ng Copyright’ (2004–05) 57 *Stanford Law Review* 485, 514 (hereafter Sprigman, ‘Reform(Aliz)ng Copyright’).

¹⁴ An Act to Amend and Consolidate the Acts Respecting Copyright (1909) § 9 <<https://www.copyright.gov/history/1909act.pdf>>.

¹⁵ Berne Convention for the Protection of Literary and Artistic Works, Art 5(2).

¹⁶ Berne Convention Implementation Act of 1988, Pub.L. 100–568.

¹⁷ 17 US Code § 411; *Fourth Estate Pub. Benefit Corp. v Wall-Street.com, LLC*, 139 S. Ct. 881, 891 (2019) (‘Noteworthy, too, in years following the 1976 revisions, Congress resisted efforts to eliminate § 411(a) and the registration requirement embedded in it. In 1988, Congress removed foreign works from § 411(a)’s dominion in order to comply with the Berne Convention for the Protection of Literary and Artistic Works’ bar on copyright formalities for such works. See § 9(b)(1), 102 Stat. 2859. Despite proposals to repeal § 411(a)’s registration requirement entirely, however, see S. Rep. No 100–352, p. 36 (1988), Congress maintained the requirement for domestic works, see § 411(a)’).

¹⁸ Sprigman, ‘Reform(Aliz)ng Copyright’ (n 13) 539.

¹⁹ William W Fisher, ‘Recalibrating Originality’ (2016–17) 54 *Hous L Rev* 437, 438 (hereafter Fisher, ‘Recalibrating Originality’).

²⁰ *Burrow-Giles Lithographic Co. v Sarony*, 111 US 53, 58 (1884).

These standards diminished considerably in the mid-1900s. Today, the prevailing test in the US requires only a ‘minimal degree of creativity’.²¹ As Peter Jaszi canily observes, an expansive conception of copyright ‘offers no sound basis for distinguishing between oil paintings, art reproductions, motion pictures, lamp bases, poems, and inflatable plastic Santa Clauses.’²²

In comparison to the US, continental Europe’s originality requirements are more exacting. However, Europe has relaxed its standards in recent years. Before the harmonization of European copyright law, some jurisdictions imposed stringent originality standards. Austria, eg, required that a copyrightable photograph differ substantially from pre-existing photographs.²³ The standard promulgated by the European Court of Justice ‘is widely seen as an effort to forge a compromise between the more stringent rules previously in force in many continental countries and the more relaxed approach previously in force in Ireland and the United Kingdom.’²⁴ Thus, while copyright today may not protect rote drudgery, it will attach to most other fixations of mundane expression.

2.1.3 New fixation technologies

The third trend is the proliferation of technologies that fix in tangible form expression that would have gone unrecorded in previous eras. Typically, copyright regimes require protected works to be ‘fixed.’²⁵ Today, it is easier than ever to record, store, and transmit large amounts of text, images, video, and audio. These technologies mediate and memorialize interpersonal interactions. Instructions to employees might be fixed in emails instead of announced verbally, flirtatious conversation might take place on a dating app instead of in a bar, and friends might share jokes and skits as videos sent to a group messaging thread. Expressive interactions that were heretofore unfixed are now fixed. Under most copyright regimes, media like this will be both sufficiently original and sufficiently fixed to receive full copyright protections. As Joseph Miller wrote in 2009, ‘technological change is as much a part of copyright’s conquest of daily life as any legal rule. Low-cost computers (with word processing, e-mail, photo, music, drawing, and browsing applications) linked to a global, high-speed communications network routinely transform us into gushing copyright and infringement fountains.’²⁶ If the lowering

²¹ *Feist Publications, Inc. v Rural Tel. Serv. Co.*, 499 US 340, 345 (1991).

²² Peter Jaszi, ‘Toward a Theory of Copyright: The Metamorphoses of “Authorship”’ (1991) 1991 *Duke Law Journal* 455, 485.

²³ Roman Heidinger, ‘The Threshold of Originality under EU Copyright Law’ (Co-Reach Intellectual Property Rights in the New Media, Hong Kong, 20 October 2011); cited in Fisher, ‘Recalibrating Originality’ (n 19) 439.

²⁴ Fisher, ‘Recalibrating Originality’ (n 19) 443.

²⁵ 17 USC § 102(a); Copyright, Designs and Patents Act 1988, s 3(2); Berne Convention for the Protection of Literary and Artistic Works, Art 2(2).

²⁶ Joseph Scott Miller, ‘Hoisting Originality’ (2009) 31 *Cardozo Law Review* 451, 459.

of originality standards democratized authorship, then new fixation technologies and the abolition of formalities democratized *ownership*.

2.2 Copyright as Accidental Obstacle: Some AI Needs neither ‘Works’ nor ‘Authorship’

The confluence of unconditional copyright, low originality standards, and networked technology has precipitated an explosion of owned data. But that ownership is, in many cases, merely a nominal obstacle to using those data in AI. This is because, despite being nominally copyrightable, many of the data that train AI fit uneasily within copyright’s cosmology and do not lend themselves to uses that copyright can control.

Copyright protects ‘works’ by ‘authors.’²⁷ A work connotes a freestanding, ascertainable product of deliberate intellectual labour.²⁸ Michelangelo’s *David* is a work. *Pride and Prejudice* and *Die Zauberflöte* are works—even though, unlike the *David*, they lack a single, authoritative embodiment. A stroke of paint on a canvas can be a work. So can an email to a friend or a text-message to a lover. ‘Authorship’, in turn, suggests some measure of deliberate creative investment, originating from the author. As the US Supreme Court noted in *Feist Publications v Rural Telephone Service*, citing the copyright scholar Melville Nimmer, ‘since an author is the “creator, originator[,]” it follows that a work is not the product of an author unless the work is original.’²⁹ The *Feist* decision remains the authoritative statement that US copyright law requires protectable works to evidence a modest amount of creativity.

Copyright law already has difficulty assimilating information that resembles a work or appears to be authored, but for some reason does not amount to a ‘work of authorship’. Transient performances, eg, often evidence authorial creativity, but that authorship may not easily reduce to a unified work. Unsurprisingly, then, copyright law struggles to protect performance, and often ends up premising protection on the details that make a performance a ‘work’—choices made by a sound engineer recording an improvised solo or a cameraman recording a dance routine—rather than the substantive creativity in the performance itself.³⁰

Many AI training data are an even poorer fit for copyright law because their utility depends *neither* on any authorial qualities *nor* on any resemblance to a

²⁷ US Constitution, Art I, s 8 (referring in pertinent part to protecting the ‘Writings’ of ‘Authors’); Beijing Treaty on Audiovisual Performances, Art 2(a) (“performers” ... perform literary or artistic works or expressions of folklore’); 17 USC § 101.

²⁸ The *Oxford English Dictionary*’s pertinent definitions are: ‘A literary or musical composition, esp. as considered in relation to its author or composer’, and ‘A product of any of the fine arts, as a painting, statue, etc.’ ‘Work, N’, *OED Online* (Oxford University Press, no date).

²⁹ *Feist Publications, Inc. v Rural Tel. Serv. Co.*, 499 US 340, 351–2 (1991).

³⁰ Rebecca Tushnet, ‘Performance Anxiety: Copyright Embodied and Disembodied’ (2013) 60 *Journal of the Copyright Society of the USA* 209, 210.

work. Indeed, some of the richest data for AI are information that copyright law traditionally disfavors, and the copyright-protected status of which is largely accidental. A highly expressive, stylized photograph could well be *less* valuable as training data for a facial recognition algorithm than a full-frontal mugshot. For training voice-synthesis AI, a recording of someone reciting every word in the dictionary, or enunciating every combination of phonemes, would probably be no less useful than a dramatic soliloquy. Are one's involuntary facial movements a work? Are they authored? Is one's gait a work of authorship? What about facial features, speaking cadence, or vocal timbre? It is true, of course, that these attributes *sometimes* evince authorship, but they certainly are not *always* authored. A recording of someone reading a banal sentence, or just enunciating a series of random syllables, is a performance without meaningful indicia of copyrightable authorship. If that recording were made by someone other than the speaker, there would be few, if any, hooks in contemporary doctrine that could give the speaker a copyright interest in the recording itself. In short, it is a legal and technological accident that such large amounts of training data are copyright-protected today.

Moreover, technological progress may obviate the need to fix training data at all, or at least in any form that resembles a work. Already, digital personal assistants learn to recognize our voices and faces without requiring us to supply photographs or sound recordings that we have created ourselves.³¹ Google's 'federated learning' technology trains AI models without reproducing training data on cloud servers.³² Thus, it is neither the qualities of 'authorship' nor of 'works' that make training data valuable for a number of current AI applications. And it is quite easy to imagine that the powerful, valuable training data of the future will be neither authored nor works, and therefore that copyright's somewhat accidental relevance could wane.

3. A Taxonomy of Training Data

Because one of the most distinctive aspects of contemporary AI is its reliance on large quantities of training data, it is illuminating to characterize 'AI problems' in terms of the legal status of their training data.³³ These training data are often copyrighted, and the individuals who develop or deploy an AI system may not be the rights holders, or even the licensees, of the data. As a result, copyright law sometimes seems like a viable counterweight to a number of AI problems. On other occasions, copyright seems like it might stifle salutary applications of AI. Often, copyright's involvement in this process is accidental—a model might learn from

³¹ 'About face ID advanced technology' (*Apple Support*) <<https://support.apple.com/en-us/HT208108>>.

³² 'Federated learning: collaborative machine learning without centralized training data' (*Google AI Blog*, 6 April 2017) <<https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>>.

³³ Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 58.

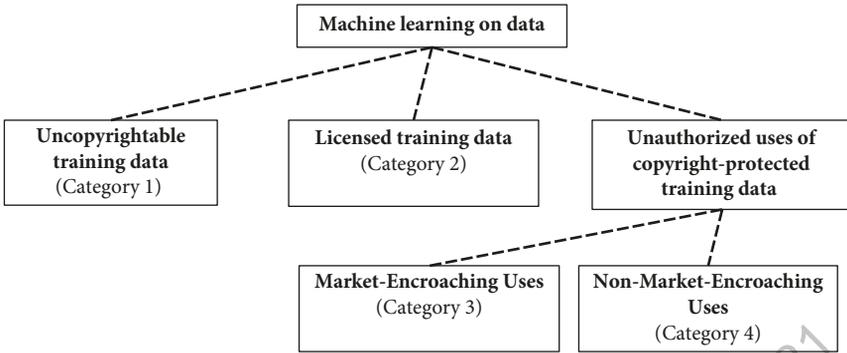


Figure 10.1 Categories of training data used in machine learning

copyrighted data but have little to do with the information that copyright protects, or with copyright’s animating principles.

This section identifies four different categories of machine learning on data, set forth in Figure 10.1: (1) uses involving uncopyrightable training data, including expression that has fallen into the public domain, (2) uses involving copyrightable subject matter that has been released under a permissive licence or licensed directly from rights holders, (3) market-encroaching uses, and (4) non-market-encroaching uses. The first of these categories is straightforward: it encompasses uses of AI that do not implicate copyrighted training data at all. These uses may be worrisome for any number of reasons, but they do not invite control through copyright law. The second category is similar: it comprises AI that trains on data that are authorized for such a use, and therefore raises no copyright issues at the initial stages of AI development.

The third category is ‘market-encroaching’ uses: uses of copyrighted training data undertaken for a purpose that threatens the market for those data. The fourth category is ‘non-market-encroaching’ uses, which do rely on copyrighted training data, but for purposes unrelated to copyright’s monopoly entitlement. Of these four categories, only market-encroaching uses are properly controlled by copyright law—and yet, paradoxically, copyright law seems worst-equipped to regulate them.

3.1 Uncopyrightable Training Data

Many consequential uses of artificial intelligence have neither a formal nor a substantive relation to copyright law.³⁴ In other words, existing copyright law provides

³⁴ Of course, underlying software or algorithms may themselves be the subject of a copyright interest, as might their outputs, but those relationships are the focus of other chapters of this book.

no ‘hook’ for regulating the technology, and the technology’s application does not relate to the normative purposes of copyright law. This is most likely to be the case for applications that rely on uncopyrightable training data.

AI that trains on uncopyrightable, non-expressive information is probably more significant than AI that trains on expression that is in the public domain. Fraud-detection or creditworthiness models analyse factual or uncreative data to make inferences that impact individuals’ livelihoods. Predictive policing techniques, eg, deploy AI to analyse factual information, like location data, call logs, or criminal records.³⁵ These technologies can threaten civil liberties, and there is every reason to regulate them in some fashion. But there is essentially no reason for *copyright* to be the regulatory mechanism: factual training data do not implicate copyright, and the applications of the trained models, while consequential, usually do not relate to copyright’s purposes.

Public-domain expression might include government work products and out-of-term literary or musical works. Such data have plenty of uses: proceedings of the European Parliament, eg, make for a rich machine-translation training corpus because they offer a voluminous dataset in many parallel translations.³⁶ Out-of-copyright music can train AI to generate novel songs,³⁷ and the long-deceased painter Rembrandt van Rijn’s oeuvre can teach an AI to generate a novel Rembrandt.³⁸

AI trained on material in the public domain may implicate branches of IP other than copyright. The *sui generis* database right, in the jurisdictions that recognize it, might protect certain training corpora.³⁹ Trade secret law, too, could protect valuable troves of data. As mentioned in the introduction above, some IP regimes might protect an algorithm, a software implementation, a weighted model, or the outputs thereof, rather than training data. And a model trained on unprotected materials may nevertheless create outputs that are similar to copyrighted works. But uses of uncopyrightable data and public-domain expression to train artificial intelligence are properly beyond the reach of copyright law.

³⁵ Walter L Perry and others, ‘Predictive policing: forecasting crime for law enforcement’ (*Rand Corporation*, 2013) 2 <https://www.rand.org/pubs/research_briefs/RB9735.html>.

³⁶ ‘Europarl parallel corpus’ <<https://www.statmt.org/europarl/>>.

³⁷ Jukedeck R&D Team, ‘Releasing a cleaned version of the Nottingham dataset’ (*Jukedeck Research*, 7 March 2017) <<https://research.jukedeck.com/releasing-a-cleaned-version-of-the-nottingham-dataset-928cdd18ec68>>.

³⁸ Tim Nudd, ‘Inside “The Next Rembrandt”: how JWT got a computer to paint like the old master’ (*Adweek*, 27 June 2016) <<http://www.adweek.com/brand-marketing/inside-next-rembrandt-how-jwt-got-computer-paint-old-master-172257/>>.

³⁹ Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases [1996] OJ L 77/20.

3.2 Licensed Training Data

Permissively-licensed, in-copyright works—such as those published under Creative Commons licences—offer substantially the same opportunities for training AI as public-domain expressive material. As uses of permissively-licensed works become more prominent, however, litigation may test how the terms of common licences govern certain uses in machine learning. Licensors may also begin to adopt licences that purport to restrict controversial uses, such as the training of facial recognition algorithms.⁴⁰ And, of course, perhaps the most valuable applications of AI are those that train on licensed, non-public corpora of copyrighted data, like Facebook’s trove of Instagram photos or Google’s store of emails, all of which are licensed from end users.⁴¹

Even authorized uses of training data may raise issues for copyright law. A model trained on authorized copies of training data may create output that appears to be substantially similar to prior copyrighted media.⁴² Present-day doctrines concerning ‘improper appropriation’ might underemphasize AI’s ability to appropriate difficult-to-articulate expressive style.⁴³ Existing law might not establish who, if anyone, owns these outputs. And a different paradigm for producing expressive works may require updating the received wisdom about the ways in which copyright incentivizes creative production.⁴⁴ But because these uses raise no copyright issues at the training stage, they are not the focus of this chapter.⁴⁵

3.3 Market-Encroaching Uses of Copyrighted Works

Some uses of machine learning are ‘market-encroaching’: these uses of AI plausibly threaten the market for the copyrighted works that comprise their training data. In market-encroaching uses, economic harms to rights holders predominate over dignitary harms. But not all uses that might diminish the value of their copyrighted training data constitute market-encroaching uses. For example, AI

⁴⁰ ‘Facial recognition’s “dirty little secret”: social media photos used without consent’ (*NBC News*, 17 March 2019) <<https://www.nbcnews.com/tech/internet/facial-recognition-s-dirty-little-secret-millions-online-photos-scraped-n981921>>.

⁴¹ For examples of the licences platforms like Google and Facebook secure from their users, see ‘Google terms of service—privacy & terms—Google’ (*Google*) <<https://policies.google.com/terms?fg=1>>; ‘Terms of use | Instagram help center’ (*Instagram*) <<https://help.instagram.com/581066165581870>>.

⁴² Sobel, ‘Artificial Intelligence’s Fair Use Crisis’ (n 7) 66.

⁴³ Benjamin LW Sobel, ‘Elements of Style: Emerging Technologies and Copyright’s Fickle Similarity Standards’ (unpublished manuscript, 2019).

⁴⁴ Pamela Samuelson, ‘The Future of Software Protection: Allocating Ownership Rights in Computer-Generated Works’ (1986) 47 *University of Pittsburgh Law Review* 1185, 1224.

⁴⁵ Of course, a given licence may not speak unequivocally about whether it would permit certain uses as training data. This uncertainty, however, would be a matter of contractual interpretation rather than of copyright law.

text-mining techniques might help identify a novelist as a plagiarist. Analysing the novels might require reproducing them without permission, and a revelation of plagiarism could damage the market for the author's works. This is not, however, a market-encroaching use.⁴⁶ Rather, a market-encroaching use is one that encroaches upon markets *over which copyright grants a monopoly*: they use protected expression for purposes that usurp the market for that very expression.⁴⁷

Royalty-free or 'stock' music is a vivid example of market-encroaching AI. State-of-the-art technology can generate novel musical works and sound recordings. Startups like Aiva and Jukedeck⁴⁸ advertise AI-generated soundtracks for media producers, with different pricing tiers for non-commercial and commercial uses and different copyright ownership schemes.⁴⁹ Another startup, Melodrive, offers machine learning technology that can dynamically generate soundtracks for video games.⁵⁰ Commercial use of this AI-generated music is more than just academic conjecture. In mid-2019, Jukedeck was reportedly acquired by the Chinese firm Bytedance, a major player in entertainment AI that produces the TikTok app.⁵¹

AI-generated music does not rival the artistry of human composers and performers, but it doesn't need to.⁵² Stock music tracks guard against uncomfortable silences at hotel bars or provide aural backdrops to commercials—they do not deliver artistic revelations on their own. Today's AI-generated music is a perfectly appropriate space-filler for, say, a virtual tour of a real estate listing. In the past, humans would have had to compose, perform, and record this background music, or at least create a composition and input it into sequencing software. It may not require visionary genius to write and record a track called 'Dark Background Piano Tones', but it does demand the sort of expressive act that copyright law protects.

It is likely that some AI music startups train their models on copyright-protected music, possibly without rights holders' authorization.⁵³ Even if no

⁴⁶ *A.V. ex rel. Vanderhye v iParadigms, LLC*, 562 F.3d 630 (4th Cir. 2009). The author used a variation of this example in comments submitted to the US Patent and Trademark Office on 15 December 2019, and to the World Intellectual Property Organization on 14 February 2020.

⁴⁷ Judge Pierre Leval of the US Court of Appeals for the Second Circuit, a leading jurist on fair use, has in his analysis limited fair use's market-substitution factor to the 'protected aspects' of the works used. *Authors Guild v Google, Inc.*, 804 F.3d 202, 229 (2d Cir. 2015); Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 55, 56.

⁴⁸ Jukedeck went offline sometime between 15 June and 24 June 2019. Compare 'Jukedeck—create unique, royalty-free AI music for your videos' (*Internet Archive*) <<https://web.archive.org/web/20190615121143/https://www.jukedeck.com/>> with 'Jukedeck is offline' (*Internet Archive*) <<https://web.archive.org/web/20190624100130/https://www.jukedeck.com/>>.

⁴⁹ 'AIVA—the AI composing emotional soundtrack music' (*AIVA*) <<https://aiva.ai/>>; 'Jukedeck—create unique, royalty-free soundtracks for your videos' (*Jukedeck*) <<http://www.jukedeck.com/make/>>.

⁵⁰ 'Melodrive | adaptive AI solutions' (*Melodrive*) <<http://melodrive.com/>>.

⁵¹ 'AI-music firm Jukedeck's CEO now runs AI lab of TikTok owner Bytedance' (*Music Ally*) <<https://musically.com/2019/07/23/ai-music-startup-jukedecks-ceo-now-runs-the-ai-lab-of-tiktok-owner-bytedance/>>.

⁵² Here, this chapter refers only to music produced on-demand by an AI engine designed to eschew human input, rather than musicians who use artificial intelligence technologies to achieve particular artistic effects.

⁵³ Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 77–9.

companies in this space currently make unauthorized uses of copyrighted music, they certainly *could* do so without difficulty. To the extent this training requires reproducing source music in order to assemble datasets, it is *prima facie* infringement. Moreover, as this author has argued elsewhere, because such a use of data harnesses expressive works *for expressive purposes*, it is possible that even the US's relatively permissive fair use doctrine would not excuse it.⁵⁴ Jurisdictions that lack an exception or limitation as flexible as fair use would be even less likely to excuse such a market-encroaching use of copyrighted materials.⁵⁵

Thus, copyright may indeed be a barrier to commercial music synthesis, as it would be for other market-encroaching uses of AI. But this is sensible. Whether copyright *ever* ought to impede expressive activity is worth interrogating.⁵⁶ To the extent that copyright should obstruct downstream creativity, however, it is under circumstances more or less like these. AI-generated stock music could very well displace much of the market for human-created stock music. If it does, it will almost certainly be because the technology has appropriated some expressive value from its training data.

3.4 Non-Market-Encroaching Uses of Copyrighted Works

True market-encroaching uses represent a small fraction of AI endeavours. Far more concerning applications of AI rely on data, sometimes copyrighted data, in order to accomplish purposes that have no bearing on rights holders' legitimate markets. Arguably more prominent than market-encroaching machine learning is AI that learns from potentially copyrighted training data for purposes unrelated to the expressive aspects of those data. Facial recognition technology, eg, trains on digital photographs of human faces, which are likely to be copyright-protected.⁵⁷ But the data that facial recognition algorithms analyse are unrelated to the expression in a photograph that copyright protects. Facial geometry, like other biometric data, raises such urgent privacy concerns precisely because it is innate and immutable, *not* authored.⁵⁸ The best facial recognition training data are close-cropped pictures of faces, edited to leave essentially no room for copyrightable expression.⁵⁹ To the extent expressive details like lighting and angle appear in training data, it is

⁵⁴ Ibid, 78, 79.

⁵⁵ Arts 3 and 4 of the recent EU Directive on Copyright in the Digital Single Market may change this calculus. See Section 4.1.2.

⁵⁶ David L Lange and H Jefferson Powell, *No Law: Intellectual Property in the Image of an Absolute First Amendment* (Stanford University Press 2008).

⁵⁷ Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 67. The author used a variation of this example in comments submitted to the US Patent and Trademark Office on 15 December 2019.

⁵⁸ Benjamin LW Sobel, 'Countenancing Property' (unpublished manuscript, 2019) (hereafter Sobel, 'Countenancing Property').

⁵⁹ Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 67, 68.

so that algorithms' performance becomes invariant to differences in the presentation of an underlying human subject. Thus, facial recognition does not encroach upon a copyright-protected interest in source photographs.

Substantially the same analysis applies to 'deepfakes,' AI-synthesized video and audio that reproduce the human likeness.⁶⁰ Deepfake technology allows relatively unsophisticated actors to synthesize verisimilar media using only photographs or videos of a target subject's face, which can then be superimposed over unrelated video footage. Unlike facial recognition, it is likely that *someone* could have a copyright claim arising out of an unauthorized deepfake. But this rights holder, who might own the video into which someone is falsely inserted, is unlikely to be the person most harmed by the synthesized media. The principal victim—the person inserted into falsified media—will have little recourse in copyright. For one, that person may not own the rights to the photos and videos that trained the AI. More fundamentally, deepfakes reproduce immutable, non-authored information about the human likeness. Much like facial recognition, then, deepfakes do not by necessity appropriate expressive information, and may therefore lie outside copyright's scope.⁶¹

It is obvious that facial recognition and deepfakes encroach on legitimate personal interests. Facial recognition facilitates state surveillance and invasive corporate marketing, and deepfakes can fuel abuse, extortion, and fraud. Even in the business-friendly—or, euphemistically, innovation-friendly—US, some academics are calling for a moratorium on facial recognition.⁶² At least two US cities have banned facial recognition use by local government.⁶³ Federal and state legislators in the US have proposed deepfake-specific legislation.⁶⁴ In other words, the public is clamouring for a way to bridle these technologies. Described generally, copyright is a legal entitlement to control the data that train recognition and synthesis models; in these terms, copyright seems like the perfect countermeasure to overreaching facial recognition and deepfake technology.

But a recent controversy over facial recognition data illustrates that copyright is the wrong vehicle for constraining facial recognition. Some background: insofar as facial recognition software's training data mis- or under-represent particular groups of people, the resulting algorithms are liable to perform less accurately on

⁶⁰ Danielle Keats Citron, 'Sexual Privacy' (2019) 128 Yale Law Journal 1870, 1922.

⁶¹ In a recent analysis of a deceptive use of likeness, albeit a less technologically sophisticated one, the US Court of Appeals for the Ninth Circuit reasserted copyright's role as an economic entitlement, rather than a privacy protection. *Garcia v Google, Inc.*, 786 F.3d 733, 745 (9th Cir. 2015).

⁶² Woodrow Hartzog, 'Facial recognition is the perfect tool for oppression' (*Medium*, 2 August 2018) <<https://medium.com/s/story/facial-recognition-is-the-perfect-tool-for-oppression-bc2a08f0fe66>>.

⁶³ Tim Cushing, 'Somerville, Massachusetts becomes the second US city to ban facial recognition tech' (*Techdirt*, 1 July 2019) <<https://www.techdirt.com/articles/20190627/21011442491/somerville-massachusetts-becomes-second-us-city-to-ban-facial-recognition-tech.shtml>>.

⁶⁴ H.R. 3230, 116th Cong. (2019); S. 3805, 115th Cong. (2018).

those groups.⁶⁵ Indeed, flagship facial recognition algorithms may be notably worse at recognizing people who are not white men than they are at recognizing white men.⁶⁶ In early 2019, IBM published a dataset designed to mitigate these biases and increase diversity in facial recognition training data.⁶⁷ IBM's dataset comprised images uploaded to the Flickr photo-sharing service under permissive Creative Commons (CC) licences. Some Flickr users were surprised and dismayed to learn that their photographs ended up in the IBM dataset; one news article quoted several photographers who expressed frustration that IBM had not given them notice before using their CC-licensed photographs in a facial recognition tool.⁶⁸

If using a photograph to train facial recognition aligned well with the interests that copyright protects—and the photographers on Flickr understood the terms of their CC licences—we would not expect photographers to bristle at IBM's use of their photographs. But the interests at stake here do not align with copyright at all, and have much more to do with the privacy of the photographs' subjects than with the economic interests of the photographers. Addressing the kerfuffle, Creative Commons' then-CEO Ryan Merkeley wrote, 'copyright is not a good tool to protect individual privacy, to address research ethics in AI development, or to regulate the use of surveillance tools employed online.'⁶⁹ Merkeley is correct both normatively and descriptively. Copyright law as it exists today is the wrong tool to further the urgent and legitimate goal of regulating facial recognition. To begin with, blackletter copyright law would not entitle most authors to reserve rights to control how photographs are used for facial recognition, because facial recognition is not a market-encroaching use.⁷⁰ The technology most likely does not implicate protected expression in the source photographs that train algorithms, which means that it may not even amount to a *prima facie* infringement. And even if using photographs to train facial recognition does nominally encroach upon an exclusive right, that use would likely be excused as fair use in the US.

In sum, the most worrisome non-market-encroaching AI uses training data that copyright law disfavors, for purposes outside copyright's ambit. Images and videos that train facial recognition or deepfake models need not be expressive. In fact, the more these media evidence 'facts' of a person's physical appearance, rather than a photographer's expressive contributions, the better. Synthesizing someone's

⁶⁵ Steve Lohr, 'Facial recognition is accurate, if you're a white guy' (*The New York Times*, 11 February 2018) <<https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html>>.

⁶⁶ Claire Garvie and others, 'The Perpetual Line-Up' (18 October 2016).

⁶⁷ Michele Merler and others, 'Diversity in Faces' (2019) arXiv:1901.10436 [cs.CV].

⁶⁸ 'Facial recognition's "dirty little secret": social media photos used without consent' (*NBC News*, 17 March 2019) <<https://www.nbcnews.com/tech/internet/facial-recognition-s-dirty-little-secret-millions-online-photos-scraped-n981921>>.

⁶⁹ Ryan Merkeley, 'Use and fair use: statement on shared images in facial recognition AI' (*Creative Commons*, 13 March 2019) <<https://creativecommons.org/2019/03/13/statement-on-shared-images-in-facial-recognition-ai/>>.

⁷⁰ Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 67, 68.

voice does not require a recording of that person ‘performing’ a ‘work’. A recording of a banal, extemporaneous conversation could work just as well as a recording of a dramatic reading. Moreover, none of these data need to be fixed under the authority of the people whose privacy and dignitary interests they implicate—and, at least in the case of photographs and videos, they often are not. These applications of AI engender wrongs that compel some sort of legal intervention. But the substantive law of copyright would be difficult to deploy to redress these harms, and copyright’s normative underpinnings suggest that this difficulty is no accident.

4. Solving AI’s Copyright Problem and Copyright’s AI Problem

Copyright may not be the appropriate means of redressing the bulk of our concerns about AI—even AI that trains on copyrighted data. Copyright can, however, mitigate the injuries presented by market-encroaching uses of copyrighted works. This section describes several ways of addressing the tensions between present-day copyright doctrine and commercial realities. It concludes that, of the regimes practically available, the EU’s new TDM exceptions do a surprisingly good job of balancing the various interests that market-encroaching uses implicate. Nevertheless, this section also describes farther-reaching reforms that could address the gaps in the EU exception. Finally, even though copyright seems like an obvious means of controlling harmful, non-market-encroaching machine learning, this section reminds readers why deploying copyright to address these problems would be both doctrinally awkward and normatively unwarranted.

4.1 Market-Encroaching Uses

Copyright should support robust technological progress without steamrolling legitimate, present-day legal entitlements. Copyright should *not*, however, serve to wring a few more years of unnecessary profit from otherwise obsolete endeavours. In other words, it is not copyright’s place to stop AI from eating the stock music business; rather, copyright’s place is to ensure that AI’s human trainers get their due as AI begins to displace them.

An ideal system, then, would give the appropriate creators some appropriate compensation for their service to market-encroaching AI, and it would do so with minimal transaction costs. Such a system is easy to write about and difficult to implement. As this chapter argued above, duly compensating expressive activity without throttling innovation would require reexamining fundamental features of copyright protection in addition to redrawing existing exceptions and limitations. Without progress in both arenas, it is unlikely that these problems will be resolved.

Indeed, the differences between US and EU law illustrate different ways of failing to address the AI problem. To the extent that EU member states impose slightly higher bars to initial copyright interests in the form of higher originality requirements, those jurisdictions thereby mitigate aspects of the copyright regime that may inhibit the progress of AI. At the same time, the EU offers only a few rigid carveouts for unauthorized uses of copyrighted media, which means that it is likely to over-deter machine learning that does not threaten rights holders' legitimate economic interests. The US, on the other hand, encumbers training data with its rock-bottom criteria for copyright protection. Yet it shows a greater solicitude for uses that do not prejudice rights holders' economic interests because of its flexible fair use doctrine.

4.1.1 Fair use is ill-equipped to address market-encroaching uses

The US's fair use doctrine,⁷¹ and regimes like it, set forth a flexible, case-by-case standard that allows courts to immunize certain unauthorized uses of copyrighted works. Fair use's open-endedness permits the doctrine to adapt to technological change more nimbly than a closed-list approach might.⁷² Indeed, other contributions to this volume advocate more open-ended exceptions frameworks for precisely this reason.⁷³

It is difficult to deny that more rigid exceptions regimes restrict the salutary, non-market-encroaching uses of artificial intelligence already flourishing under the US's fair use regime.⁷⁴ With respect to *market-encroaching* uses of AI, however, fair use is not ideal. Fair use is all-or-nothing: either the defence succeeds and the use in question is unreservedly legal, or the defence fails and leaves the defendant liable for infringement. Neither situation makes sense for market-encroaching uses of AI because of the way AI alters traditional balances of equities.⁷⁵ Treating market-encroaching AI just like other forms of copyright infringement would obstruct technological progress by offering rights holders property remedies like statutory damages and injunctions, which would far overcompensate for the harms authors are likely to suffer from having their work included in training datasets. At the same time, categorizing market-encroaching AI as fair use would fail to compensate rights holders for valuable commercial uses of their *expression*. A more equitable approach would encourage—and perhaps even reward—authorized uses of copyrighted works in market-encroaching AI. The EU's new TDM exception, while far from perfect, represents a tentative step towards such a system.

⁷¹ 17 USC § 107.

⁷² See, eg, Ian Hargreaves, 'Digital Opportunity: A Review of Intellectual Property and Growth' (May 2011) 43–4 (comparing and contrasting the fair use doctrine with the European enumerated-exceptions approach in the TDM context) (hereafter Hargreaves, 'Digital Opportunity').

⁷³ Tianxiang He, Chapter 9 in this volume.

⁷⁴ Hargreaves, 'Digital Opportunity' (n 72) 43, 44.

⁷⁵ Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 79–82.

4.1.2 The EU TDM exception: better for market-encroaching uses, worse for non-encroaching uses

The most significant recent development in the AI-and-copyright field is the EU's Digital Single Market (DSM) Directive of 2019. The DSM Directive mandates two pertinent exceptions to member state copyright laws. The first, in Article 3 of the Directive, requires member states to permit certain reproductions of copyrighted materials by research and cultural heritage organizations, undertaken for the purposes of text and data mining (TDM) research.⁷⁶ The second, in Article 4, extends that same exception to *any entity* seeking to perform TDM.⁷⁷ However, the Article 4 exception does not apply when rights holders expressly reserve their TDM rights.⁷⁸

Despite calling itself an 'exception,' Article 4 appears to operate more like a formality.⁷⁹ The language of exceptions makes sense, because framing Article 4 as a formality would invite scrutiny as a possible violation of the Berne Convention. Nonetheless, because Article 4 prescribes steps an author must take in order to possess a particular exclusive right, it redraws the default bundle of copyright entitlements in a way that an exception does not. Like a formality, and unlike an exception, Article 4 focuses on an author's behaviour, rather than on an evaluation of a particular use. Article 4 thus subverts unconditional copyright by making a TDM right conditional on an owner's express reservation. In other words, Article 4 requires an author affirmatively to reserve a right to exclude uses of her works to train AI, if she wishes to exercise that right.

In spite of, or maybe because of, its possible incongruence with the Berne Convention's prohibition on formalities, Article 4 represents a positive development for copyright law and for AI. Indeed, Article 4's formality-like qualities allow it to operate more effectively than a pure exception. Article 4 solves part of the AI-and-copyright problem by facilitating the confident development of artificial intelligence using data that are not subject to reserved rights.

Article 4 falls short, however, in establishing an effective scheme for authors who do not object to their works training AI, but who want to be compensated. Presumably, a sizeable portion of rights holders would be willing to license their works for use as training data, but unwilling to extend the gratis licence that Article 4 creates by default. In theory, a centralized registration regime, coupled with collective licensing arrangements, could lubricate this market. However, AI training combines massive corpora of works with low per-use payoffs, which together mean

⁷⁶ 2019 OJ (L 130/92) 113.

⁷⁷ *Ibid.*, 113, 114.

⁷⁸ *Ibid.*

⁷⁹ A blog post by Neil Turkewitz brought Art. 4's formality-like qualities to the author's attention. Neil Turkewitz, 'Sustainable text and data mining: a look at the recent EU Copyright Directive' (*Medium*, 16 May 2019) <https://medium.com/@nturkewitz_56674/sustainable-text-and-data-mining-an-look-at-the-recent-eu-copyright-directive-9ea13ba05f60>.

that incentives to register works will probably be too low to sustain a conventional government-run registration system. The private sector may be able to fill this lacuna: private entities like the PLUS registry organize media and licences to lower search and transaction costs.⁸⁰ What may distinguish AI training data from these other licensing systems, however, is the premium on secrecy that data command. A company's AI can only be as good as the data on which it trains. Thus, there will be few parties that have unilateral control over large datasets *and* an incentive to license those datasets widely. This suggests that compensated uses of authored training data—other than by platforms that receive licences from their end users—will be unlikely to flourish without some form of government intervention to lubricate the market. Finally, there is a risk that Article 4 unduly burdens *non-market-encroaching* uses to the extent that its opt-out regime empowers rights holders to exclude others from making *non-market-encroaching* TDM uses of copyrighted materials.

4.2 Beyond Exceptions and Limitations

Resolving the disharmony between artificial intelligence and intellectual property regimes requires more than an exception or limitation to copyright. This subsection briefly introduces three doctrinal areas that more ambitious reforms might target: originality, formalities, and remedies.

4.2.1 Tweaking originality

Redrawing copyright's originality requirements could mitigate the potential copyright liabilities associated with some machine learning. Delineating copyright in a manner that unambiguously excluded emails, throwaway photographs, or short comments left on an Internet forum might alleviate some worries about impeding the progress of TDM. Moreover, originality-based reforms would not jeopardize the economic protections that more creative endeavours would receive. As a practical matter, originality-oriented reforms are particularly appealing because they are comparatively unlikely to conflict with international treaties. The US, eg, is not party to a treaty that places a 'ceiling' on the stringency of copyright's originality requirements.⁸¹ However, heightened originality requirements would also eliminate a rare source of leverage for the humans whose run-of-the-mill expression currently trains the AI that may replace them in the workplace.⁸² Copyright may not be the best tool to ensure distributive equity in the age of artificial intelligence,

⁸⁰ 'PLUS Registry: about the Registry' (*Plus Registry*) <<https://www.plusregistry.org/cgi-bin/WebObjects/PlusDB.woa/2/wo/Z2vbrwAumsxHHQToC5sz1w/0.111.27>>.

⁸¹ Fisher, 'Recalibrating Originality' (n 19) 457, 461.

⁸² Sobel, 'Artificial Intelligence's Fair Use Crisis' (n 7) 97.

and low originality requirements may never have been wise policy. But raising originality standards now, without social support for the workers who may soon be displaced by AI, could exacerbate the difficulties some human labourers face in an automated economy.

4.2.2 Formalities

Prior sections of this chapter have suggested that formalities regimes could mitigate risks associated with commercial AI without abridging the rights of sophisticated authors. Indeed, despite its framing as an exception, Article 4 of the EU DSM Directive may also be understood as a formalities regime. While Article 4 does not condition copyright protection on a formality, it does establish a formal prerequisite to asserting a right to exclude uses in text and data mining. As a practical matter, traditional formalities regimes would be difficult to implement in a manner consistent with major international treaties. However, it is worth noting that the technological architecture of the internet already has a well-established permissions standard that augurs well for a TDM-focused formality. The robots.txt standard for webpages sets forth the instructions that automated software should observe when indexing or archiving a given page.⁸³ Because this convention is fundamental to the Internet, a robots.txt-like standard dictating permissible uses for AI might be more straightforward to implement than other copyright formalities regimes, and less likely to penalize the legally unsavvy.

4.2.3 Rethinking remedies

Copyright punishes infringers with stiff penalties that can include statutory damages and injunctive relief. These remedies may be appropriate to deter some unauthorized uses of copyrighted materials, but they are utterly inapt for addressing unauthorized uses of copyrighted materials in training data. Machine learning technology learns from hundreds of thousands of data points at a time, each of which makes some small contribution to a trained model. Injunctive relief and high damages awards will over-deter technological development and overcompensate plaintiffs. Indeed, these remedies would be much more appropriate for the types of non-market-encroaching uses that lie beyond copyright's scope, such as the creation of simulated nonconsensual pornography using deepfakes technology.⁸⁴

In contrast to conventional property-style entitlements, liability-style rules in copyright can promote commerce that would otherwise be too costly, and promote access to expressive works that a market system might not otherwise provide.⁸⁵

⁸³ 'Robots Exclusion Standard', *Wikipedia* (2019).

⁸⁴ Representative Clarke's DEEPFAKES bill, H.R. 3230, 116th Cong. (2019), proposes a slate of remedies strikingly similar to those codified in the United States Copyright Act of 1976, 17 USC § 504. For further discussion of the resemblances between some deepfakes-related privacy proposals and moral rights, see Sobel, 'Countenancing Property' (n 58).

⁸⁵ Jacob Victor, 'Reconceptualizing Compulsory Copyright Licenses' (2020) 72 *Stanford Law Review* 915, 920.

Indeed, a liability-style rule may be the only way of fulfilling the function that Article 4 fails to fulfil: facilitating compensated, market-encroaching uses of copyrighted training data. Authors who would otherwise opt out of the Article 4 exception might instead opt into a licensing regime that offers their works for market-encroaching AI uses at a set licensing rate.

4.3 What Copyright Cannot Do

This chapter has shown that many worrisome uses of AI depend on copyright-protected training data. However, it has also shown that only a significant revision of contemporary copyright law would allow it to regulate the applications of AI that implicate non-economic interests. Attempts to deploy copyright in service of non-economic interests like privacy, dignity, and reputation are nothing new. In the nineteenth century, an English publisher surreptitiously obtained etchings created by Queen Victoria and Prince Albert. When the publisher released a pamphlet describing the etchings, Prince Albert invoked a ‘common law right to the copy’ to enjoin the pamphlet’s publication.⁸⁶ In fact, it was partially from this jurisprudence of common-law copyright that Samuel Warren and Louis Brandeis famously distilled a ‘right to privacy’ in Anglo-American law.⁸⁷ This privacy-protective, dignity-protective conception of copyright transposes to contemporary concerns. For instance, Kim Kardashian reportedly used a copyright takedown to remove a deepfake video depicting her.⁸⁸

But the preceding sections of this chapter have described numerous reasons why copyright probably cannot redress the harms that non-market-encroaching uses of AI entail. First, many of these uses require neither works nor authorship. Instead, these uses implicate data that are not protected by copyright in the first place, or whose protection is merely incidental to a fixation step that future technologies can probably eschew. Second, even if these uses do rely on copyrighted training data, they use those data in a manner unrelated to any expressive information in the data. Indeed, it is *precisely* because these uses implicate facts of personal identity—rather than expression—that they can be so harmful.

Mechanisms like the TDM opt-out found in Article 4 may allow rights holders to control even non-expressive uses of information in copyrighted works. To the extent that a copyright regime restricts non-expressive uses, however, it at best

⁸⁶ *Prince Albert v Strange* (1849) 2 De G. & Sm. 652, 64 ER 293.

⁸⁷ Samuel D Warren and Louis D Brandeis, ‘The Right to Privacy’ (1890–91) 4 Harvard Law Review 193.

⁸⁸ ‘Kim Kardashian deepfake taken off of YouTube over copyright claim’ (*Digital Trends*, 17 June 2019) <<https://www.digitaltrends.com/social-media/kim-kardashian-deepfake-removed-from-youtube/>>; Tiffany C Li, ‘This backdoor approach to combating deepfakes won’t work’ (*Slate Magazine*, 18 June 2019) <<https://slate.com/technology/2019/06/deepfake-kim-kardashian-copyright-law-fair-use.html>>.

stifles the flow of non-proprietary information, and at worst, it encroaches upon the functions served by different causes of action designed to protect privacy and dignitary interests. Claims designed to prevent harmful uses of non-expressive information should be founded in tort or privacy law, not in the economic rights that copyright grants authors.

5. Conclusion

AI has a copyright problem: valuable business practices implicate the unauthorized reproduction of countless copyrighted works. AI's copyright problem, in turn, exposes copyright's AI problem: vast amounts of digital media are copyright-protected essentially due to historical accident. Despite adding uncertainty to AI development, the copyright protections attached to many training data cannot, and should not, regulate many applications of AI. Only those applications of AI that encroach upon the market for their copyrighted training data—such as royalty-free music generation—are properly within copyright's ambit.

This is difficult news for two reasons. First, copyright seems at first glance like it could help prevent harmful uses of AI that involve copyrighted training data, but that do not threaten the market for those data. But copyright's economic focus makes it a poor vehicle for redressing harms to privacy or dignity. Second, copyright's remedies are a bad fit for the AI-related economic injuries that copyright *can* appropriately regulate.

Thus, copyright's role in addressing AI is limited but nevertheless significant. Successful interventions will focus not simply on exceptions to copyright entitlements, but also on the nature and scope of copyrights in the first instance. The EU's recent Directive on the Digital Single Market begins to effectuate some of these reforms, although it is drafted in a way that obscures its own significance. But the problem of *compensating* authors for market-encroaching uses will persist after the Directive's implementation. The international copyright system will have to reshape itself if it is to address this problem. Whether it can do so remains to be seen.