

ARTIFICIAL INTELLIGENCE, THE LAW–MACHINE INTERFACE, AND FAIR USE AUTOMATION

Peter K. Yu

INTRODUCTION	188
I. FAIR USE AUTOMATION	191
A. <i>The Standard</i>	192
B. <i>Why Fair Use?</i>	194
C. <i>Resistance Toward Automation</i>	197
1. <i>Backward State of Technology</i>	198
2. <i>Changes in Creative Choices and Practices</i>	200
3. <i>Technological Shortcomings</i>	202
4. <i>Summary</i>	205
II. PATHWAYS FOR LEGAL AUTOMATION.....	206
A. <i>Translation</i>	206
B. <i>Approximation</i>	210
C. <i>Self-Determination</i>	213
D. <i>Summary</i>	214
III. LAW–MACHINE INTERFACE.....	214
A. <i>Allocation of Decision-Making Power</i>	215
B. <i>Hierarchy of Decisions</i>	220
C. <i>Legal Effects of Machine-Made Decisions</i>	223
IV. THE FUTURE	225
A. <i>Law</i>	226
B. <i>Legislature</i>	229
C. <i>Bench</i>	231
D. <i>Bar</i>	234
E. <i>Academe</i>	236
CONCLUSION.....	238

ARTIFICIAL INTELLIGENCE, THE LAW–MACHINE INTERFACE, AND FAIR USE AUTOMATION

Peter K. Yu*

INTRODUCTION

The past decade has seen artificial intelligence (AI) advancing in leaps and bounds, capturing the attention of not only computer experts and academic commentators but also policymakers,¹ the mass media, and the public at large.² In the early 2010s, IBM Watson successfully defeated two noted human champions in the quiz show *Jeopardy!*³ A few years later, Google DeepMind created a “Sputnik moment” in Asia⁴ when it beat the world’s best players in Go, an Asian strategy board game.⁵ In addition, recent research has shown the fast-growing improvements in the performance of artificial intelligence in poker games.⁶ Compared with quiz shows and chess games, these games have been

* Copyright © 2020 Peter K. Yu. Professor of Law, Professor of Communication, and Director, Center for Law and Intellectual Property, Texas A&M University. This Article draws on insights gleaned from the Inaugural HKU Technology Law Symposium organized by the Law and Technology Centre in the Faculty of Law at the University of Hong Kong, the International Law Weekend 2019 at Fordham University School of Law, the Third Annual IP Leaders Roundtable at UIC John Marshall Law School, the *FIU Law Review* Symposium and a presentation for the Intellectual Property Law Society at Florida International University College of Law, the Third Annual Scholarship Retreat at Texas A&M University School of Law, the 17th Annual Works-in-Progress Intellectual Property Colloquium at Santa Clara University School of Law, and a faculty speaker workshop at the University of Kansas School of Law. The discussion of fair use automation is adapted or expanded from the remarks delivered at the *FIU Law Review* Symposium, which was recently published by the *FIU Law Review*. The Author is grateful to Hannah Bloch-Wehba, Daryl Lim, William Magnuson, Milan Markovic, and the participants of these events for their valuable comments and suggestions.

1. For example, the Obama Administration has released a number of documents in the artificial intelligence area, including a strategic plan and a white paper. EXEC. OFF. OF THE PRESIDENT, ARTIFICIAL INTELLIGENCE, AUTOMATION, AND THE ECONOMY (2016); NAT’L SCI. & TECH. COUNCIL, PREPARING FOR THE FUTURE OF ARTIFICIAL INTELLIGENCE (2016) [hereinafter PREPARING FOR THE FUTURE]; NAT’L SCI. & TECH. COUNCIL, THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN (2016).

2. In March 2018, *The Daily Show* featured a segment on robots disrupting the legal system. Ronny Chieng, *Disrupting the Legal System with Robots*, THE DAILY SHOW WITH TREVOR NOAH (Mar. 7, 2018), <http://www.cc.com/shows/the-daily-show-with-trevor-noah/cast/ronny-chieng/b27lei/disrupting-the-legal-system-with-robots>.

3. John Markoff, *Computer Wins on “Jeopardy!”: Trivial, It’s Not*, N.Y. TIMES, Feb. 17, 2011, at A1.

4. See LEE KAI-FU, AI SUPERPOWERS: CHINA, SILICON VALLEY, AND THE NEW WORLD ORDER 3 (2018) (noting that AlphaGo’s victories “turned into China’s ‘Sputnik Moment’ for artificial intelligence”); Paul Mozur, *In Win for A.I., Google Program Humbles Master of a Mind-Boggling Game*, N.Y. TIMES, May 24, 2017, at B3 (describing AlphaGo as “a sort of Sputnik moment” for China).

5. See Choe Sang-Hun & John Markoff, *Machine Masters Man in Complex Game of Go*, N.Y. TIMES, Mar. 10, 2016, at A1 (reporting AlphaGo’s victory over eighteen-time world Go champion Lee Sedol); Mozur, *supra* note 4 (reporting AlphaGo’s victory over Ke Jie, the world’s then best Go player).

6. See Woodrow Barfield, *Towards a Law of Artificial Intelligence*, in RESEARCH HANDBOOK ON THE LAW OF ARTIFICIAL INTELLIGENCE 2, 9 (Woodrow Barfield & Ugo Pagallo eds., 2018) [hereinafter RESEARCH HANDBOOK] (“[A]n artificially intelligent computer designed by computer scientists beat experts in the game

particularly difficult because the poker players’ ability to bluff has created an incomplete information environment.⁷

Given these amazing technological developments, it is no surprise that legal commentators are now actively exploring how artificial intelligence will impact the law.⁸ For instance, Eugene Volokh invited us to join him for a highly provocative thought experiment concerning whether society will be ready to accept robot judges.⁹ Mireille Hildebrandt questioned whether the rapid development of artificial intelligence and smart technologies would undermine or reconfigure the ends of law in a constitutional democracy.¹⁰ Tim Wu discussed whether artificial intelligence would “eat” the law and what the impending “rise of hybrid social-ordering systems” would mean for society.¹¹ Roger Brownsword called for greater attention to the interplay of technology management and legal rules and to its impact on the traditional rules of law.¹²

of poker which required the ability to bluff and to predict whether the opponent was bluffing based on incomplete knowledge of the advisory’s hand.”); Carnegie Mellon University, *AI Beats Professionals in Six-Player Poker*, SCIENCE DAILY (July 11, 2019), <https://www.sciencedaily.com/releases/2019/07/190711141343.htm> (“An artificial intelligence program developed by Carnegie Mellon University in collaboration with Facebook AI has defeated leading professionals in six-player no-limit Texas hold’em poker, the world’s most popular form of poker.”).

7. As a Carnegie Mellon University press release stated:

Games such as chess and Go have long served as milestones for AI research. In those games, all of the players know the status of the playing board and all of the pieces. But poker is a bigger challenge because it is an incomplete information game; players can’t be certain which cards are in play and opponents can and will bluff. That makes it both a tougher AI challenge and more relevant to many real-world problems involving multiple parties and missing information.

Carnegie Mellon University, *supra* note 6.

8. For discussions in this area, see generally Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147 (2017); Rebecca Crotoof, “Cyborg Justice” and the Risk of Technological–Legal Lock-In, 119 COLUM. L. REV. F. 233 (2019); Milan Markovic, *Rise of the Robot Lawyers?*, 61 ARIZ. L. REV. 325 (2019); John O. McGinnis & Russell G. Pearce, *The Great Disruption: How Machine Intelligence Will Transform the Role of Lawyers in the Delivery of Legal Services*, 82 FORDHAM L. REV. 3041 (2014); Andrew C. Michaels, *Artificial Intelligence, Legal Change, and Separation of Powers*, 88 U. CIN. L. REV. 1083 (2020); Frank Pasquale, *A Rule of Persons, Not Machines: The Limits of Legal Automation*, 87 GEO. WASH. L. REV. 1 (2019) [hereinafter Pasquale, *A Rule of Persons*]; Frank Pasquale & Glyn Cashwell, *Four Futures of Legal Automation*, 63 UCLA L. REV. DISCOURSE 26 (2015); Richard M. Re & Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 22 STAN. TECH. L. REV. 242 (2019); Dana Remus & Frank Levy, *Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law*, 30 GEO. J. LEGAL ETHICS 501 (2017); Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87 (2014); Eugene Volokh, *Chief Justice Robots*, 68 DUKE L.J. 1135 (2019); Tim Wu, *Will Artificial Intelligence Eat the Law? The Rise of Hybrid Social-Ordering Systems*, 119 COLUM. L. REV. 2001 (2019).

9. Volokh, *supra* note 8; see also LEE, *supra* note 4, at 115 (noting the “Shanghai-based pilot program that uses data from past cases to advise judges on both evidence and sentencing”); Eric Niiler, *Can AI Be a Fair Judge in Court? Estonia Thinks So*, WIRED (Mar. 25, 2019, 7:00 AM), <https://www.wired.com/story/can-ai-be-fair-judge-court-estonia-thinks-so> (discussing the effort in Estonia “to design a ‘robot judge’ that could adjudicate small claims disputes of less than €7,000 (about \$8,000)”).

10. MIREILLE HILDEBRANDT, SMART TECHNOLOGIES AND THE END(S) OF LAW (2015).

11. Wu, *supra* note 8, at 2001.

12. As he observed:

To the extent that technological management coexists with legal rules, while some rules will be redirected, others will need to be refined and revised. Accordingly, . . . the destiny of legal rules is

Although all of these works carefully reminded us that we are still quite far away from the scenario in which machine-made decisions can provide realistic substitutes to human decisions, it is never too early to think more deeply about the complex questions arising at the intersection of artificial intelligence and the law.

One area that has not received sufficient policy and scholarly attention¹³ concerns the law–machine interface in a hybrid environment in which both humans and intelligent machines will make legal decisions at the same time.¹⁴ Because “human-machine hybrids will [likely] be the first replacement for human-only legal systems,”¹⁵ developing a deeper understanding of this interface is badly needed. Such an understanding will also be important as artificial intelligence technologies continue to improve and as society becomes more comfortable in letting machines take over some decisions that have been traditionally reserved for humans.¹⁶ Indeed, society will be better off if it can achieve an optimal allocation of decision-making power between humans and machines in such a hybrid environment. Such allocation will foster what commentators have referred to as the “new division of labor.”¹⁷

to be found somewhere in the range of redundancy, replacement, redirection, revision and refinement.

ROGER BROWNSWORD, LAW, TECHNOLOGY AND SOCIETY: RE-IMAGINING THE REGULATORY ENVIRONMENT 181 (2019).

13. Commentators have started looking into issues in this area. *See generally* Crootof, *supra* note 8 (discussing the benefits and side effects of hybrid human–AI judicial systems, or “cyborg justice”); Pasquale, *A Rule of Persons*, *supra* note 8 (explaining why complementary legal automation will play a bigger role in the legal profession than substitutive legal automation); Wu, *supra* note 8 (describing the development of hybrid machine–human systems as the “predictable future of legal adjudication” and exploring the prospects and limitations of such development).

14. *See infra* text accompanying notes 150–154.

15. Wu, *supra* note 8, at 2002; *see also* Mariano-Florentino Cuéllar, *A Common Law for the Age of Artificial Intelligence: Incremental Adjudication, Institutions, and Relational Non-Arbitrariness*, 119 COLUM. L. REV. 1773, 1775 (2019) (“[T]he coevolution of human and artificial intelligence—what we could call our dance with machines—is well on its way to becoming routine.”).

16. As Jason Millar and the late Ian Kerr observed:

[W]e will rely on robots without really knowing why—simply because their algorithms provide the greatest number of successful outcomes. We have already seen this in Google’s search approach. Neither Larry [Page] nor Sergey [Brin] (nor any other Google employee) knows exactly why one particular web page is a better result than another. When the click patterns say it is, that’s good enough. No semantic or causal analysis is required. . . . Like the ancients, we will, quite rationally, come to rely upon them, knowing full well that we cannot necessarily explain the reasons for their decisions.

Jason Millar & Ian Kerr, *Delegation, Relinquishment and Responsibility: The Prospect of Expert Robots*, in *ROBOT LAW* 102, 106–07 (Ryan Calo et al. eds., 2016); *see also* WENDELL WALLACH & COLIN ALLEN, *MORAL MACHINES: TEACHING ROBOTS RIGHT FROM WRONG* 40 (2009) (“As people come to trust the advice of a [decision support tool], it can become more difficult to question that advice. There is a danger . . . that [decision support tools] could eventually come to control the decision-making process.”); Anthony J. Casey & Anthony Niblett, *Self-Driving Laws*, 66 U. TORONTO L.J. 429, 435 (2016) [hereinafter Casey & Niblett, *Self-Driving Laws*] (“[A]s more information is generated, and the evolutionary algorithm updates and becomes a better forecaster, we imagine that judges will increasingly rely on the advice of the algorithm.”).

17. FRANK LEVY & RICHARD MURNANE, *THE NEW DIVISION OF LABOR: HOW COMPUTERS ARE CREATING THE NEXT JOB MARKET* (2012); *see also* AJAY AGRAWAL ET AL., *PREDICTION MACHINES: THE*

In a recently published article commissioned for a symposium on artificial intelligence and entertainment law, I identified the pros and cons of using algorithms to automate fair use in U.S. copyright law and called for the development of an enabling environment to facilitate such automation.¹⁸ In this Article, I utilize the case study of fair use automation to explore how legal standards can be automated and what this specific case study can teach us about the law–machine interface. Although this Article utilizes an example generated from a specialized area of the law—namely, copyright or intellectual property law—its insights will apply to other situations involving the interplay of artificial intelligence and the law. As far as these applicable insights are concerned, one should be able to substitute the fair use standard with other legal standards, such as those in criminal, tort, or traffic law.

Part I outlines the case study of fair use automation. This Part begins by offering a brief overview of the U.S. fair use standard and explaining why the automation of this standard has been chosen as an illustration. This Part then closely examines three dominant arguments against greater fair use automation. Taking seriously the benefits provided by artificial intelligence, machine learning, and big data analytics, Part II identifies three distinct pathways for legal automation: (1) the translation pathway, which converts legal mandates or analytical approaches into computer code and algorithms; (2) the approximation pathway, which ensures that machine-made decisions closely resemble human decisions; and (3) the self-determination pathway, which enables automated systems to make autonomous decisions.

Part III explores the key questions concerning the law–machine interface, the understanding of which will be important when automated systems are being designed to implement legal standards. Specifically, these questions focus on the allocation of decision-making power, the hierarchy of decisions, and the legal effects of machine-made decisions. Part IV concludes by highlighting the wide-ranging ramifications of artificial intelligence for the law, the legislature, the bench, the bar, and academe. Holistic in scope, this Part focuses on lessons drawn from studying the law–machine interface.

I. FAIR USE AUTOMATION

Although the interplay of artificial intelligence and the law can be analyzed at an abstract level, it will be more instructive to utilize a concrete example that readers can closely examine to evaluate the potential and challenges of legal automation. For coherence and analytical effectiveness, this Article uses the automation of the U.S. fair use standard as an illustrative example throughout.

SIMPLE ECONOMICS OF ARTIFICIAL INTELLIGENCE 53–69 (2018); Re & Solow-Niederman, *supra* note 8, at 282–85.

18. Peter K. Yu, *Can Algorithms Promote Fair Use?*, 14 FIU L. REV. 329 (2020) [hereinafter Yu, *Fair Use*].

Part I.A provides a brief overview of this standard. Part I.B explains why the automation of this standard has been chosen as an illustration. Part I.C explores the ongoing resistance toward such automation. This Subpart analyzes the three dominant arguments questioning the effectiveness and desirability of such automation and offers responses in turn.

A. *The Standard*

In the per curiam decision of *Dellar v. Samuel Goldwyn, Inc.*, the United States Court of Appeals for the Second Circuit described fair use as “the most troublesome in the whole law of copyright.”¹⁹ Historically, this standard can be traced back to the 1841 case of *Folsom v. Marsh*, a case concerning the unauthorized reproduction of President George Washington’s writings, official documents, and private letters that had been extracted from a twelve-volume book set.²⁰ In that case, Justice Joseph Story drew on the traditional English doctrine of fair abridgement to develop the common law doctrine of fair use.²¹ This doctrine was codified a century later when Congress undertook a major overhaul of the copyright statute in 1976.²² Section 107 of the U.S. Copyright Act provides as follows:

Notwithstanding the provisions of sections 106 and 106A, the fair use of a copyrighted work, including such use by reproduction in copies or phonorecords or by any other means specified by that section, for purposes such as criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research, is not an infringement of copyright. 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.²³

19. 104 F.2d 661, 662 (2d Cir. 1939) (per curiam).

20. 9 F. Cas. 342 (C.C.D. Mass. 1841) (No. 4,901).

21. See *id.* at 345–49. For discussions of the traditional English doctrine of fair abridgement, see generally Joseph J. Beard, *Everything Old Is New Again: Dickens to Digital*, 38 LOY. L.A. L. REV. 19, 24–26 (2004); Matthew Sag, *The Prehistory of Fair Use*, 76 BROOK. L. REV. 1371, 1379–93 (2011).

22. See 17 U.S.C. § 107 (codifying the fair use standard).

23. *Id.*

Although this statutory provision enumerates four non-exhaustive factors²⁴ that courts should consider when making fair use determinations, such determinations are made after the fact.²⁵ Because these determinations require a case-by-case balancing of multiple factors, the legal outcomes can vary even for cases involving the same copyrighted work or the same amount of copying.²⁶

Thus far, commentators have widely disagreed over the expediency of the fair use standard.²⁷ Its supporters have argued that this standard is clear and predictable. For instance, Pamela Samuelson observed:

If one analyzes putative fair uses in light of cases previously decided in the same policy cluster, it is generally possible to predict whether a use is likely to be fair or unfair. . . . The only clusters of fair use cases in which it is quite difficult to predict whether uses are likely to be fair is in the educational and research use clusters where judges have tended to take starkly different perspectives on fair use defenses in these settings²⁸

Commentators such as Professor Samuelson and Michael Madison also noted how the use of clusters could help provide the fair use regime with more clarity and predictability.²⁹ By contrast, those critical of fair use took the opposite view. As the Australian Law Reform Commission recounted in its final report on copyright and the digital economy:

24. See *id.* (using the phrase “shall include” when referring to the list of fair use factors); *id.* § 101 (“The terms ‘including’ and ‘such as’ are illustrative and not limitative.”).

25. See Dan L. Burk, *Algorithmic Fair Use*, 86 U. CHI. L. REV. 283, 288 (2019) [hereinafter Burk, *Algorithmic Fair Use*] (“[F]air use carries with it the disadvantage of ex ante uncertainty; no one can be entirely certain in advance how a court will weigh the four factors, and hence there is always some apprehension that a use may be found infringing rather than fair.”); Dan L. Burk & Julie E. Cohen, *Fair Use Infrastructure for Rights Management Systems*, 15 HARV. J.L. & TECH. 41, 61 (2001) (“Under the current conception of fair use, the decision whether or not to use a work is made ex ante by the user—if an infringement suit is brought later, the court may or may not validate the user’s calculus, but penalties, if any, are imposed after the use has been undertaken.”); John S. Erickson & Deirdre K. Mulligan, *The Technical and Legal Dangers of Code-Based Fair Use Enforcement*, 92 PROC. IEEE 985, 992 (2004) (“In the area of copyright law, the evolution of the doctrine of ‘fair use’ is tightly bound to the practice of after-the-fact adjudication.”).

26. For example, the fair use analysis of the unauthorized use of a copyrighted song for parody is significantly different from that of the use of the same song for advertising. See *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 594 (1994) (finding that 2 Live Crew’s parody of Roy Orbison’s “Oh, Pretty Woman” may constitute fair use and remanding the case to the lower court).

27. See generally Australian Law Reform Commission, *Copyright and the Digital Economy* (Discussion Paper No 79, May 2013) 74–76 (discussing the criticism that “[f]air use would create uncertainty and expense”); Peter K. Yu, *The Quest for a User-Friendly Copyright Regime in Hong Kong*, 32 AM. U. INT’L L. REV. 283, 331–34 (2016) [hereinafter Yu, *The Quest*] (discussing the debate on the fair use standard’s lack of clarity and precision).

28. Pamela Samuelson, *Unbundling Fair Uses*, 77 FORDHAM L. REV. 2537, 2542 & n.28 (2009) [hereinafter Samuelson, *Unbundling Fair Uses*].

29. See generally Michael J. Madison, *A Pattern-Oriented Approach to Fair Use*, 45 WM. & MARY L. REV. 1525 (2004) (advancing a pattern-oriented approach to fair use decisions); Samuelson, *Unbundling Fair Uses*, *supra* note 28 (arguing that a focus on common patterns, or what Professor Samuelson called “policy-relevant clusters,” will make fair use law more coherent and predictable than many commentators have perceived).

The opponents of fair use have pointed to research indicating that the outcome of fair use cases is unpredictable. The outcome of litigation is never completely predictable—if it were, the parties would not have commenced litigation, or would likely have settled. This is also true of recent litigation over the fair dealing exceptions and specific exceptions.³⁰

Although the lack of consensus among copyright experts about the clarity and predictability of fair use has foreshadowed the challenge society will have when deploying algorithms and artificial intelligence to automate this legal standard, the widespread concerns about the standard's lack of clarity and predictability also present an immense opportunity—If human decisionmakers have tremendous difficulty making fair use determinations, will intelligent machines do a better job? Regardless of whether these machines can perform better or not, the analysis in this Article will inform research on the interplay of artificial intelligence and the law. Such analysis will also allow us to better understand the importance of providing appropriate interfaces between laws and machines.

B. *Why Fair Use?*

For a cross-cutting project on legal automation and the law-machine interface, choosing an illustration that is familiar to a wide range of readers will be highly important. Other than this Author's specialized expertise in the subject area and his past involvement in global copyright reform,³¹ the case study of fair use automation was chosen for five reasons. First, the topic is familiar to scholars writing in the artificial intelligence area. Many of these scholars already have some expertise in intellectual property law or cyberlaw. As a result, this case study can be easily incorporated into their analyses. In addition, because any analysis in the early days of legal automation is admittedly preliminary, using an example that is well understood by those writing in the area will help foster a productive scholarly dialogue in this fast-evolving area of the law.

Second, the topic is familiar to not only legal scholars but also non-legal researchers. Whether deciding on the use of a quotation in academic research or the copying of an excerpt for classroom teaching, academics frequently have

30. Australian Law Reform Commission, *Copyright and the Digital Economy* (Report No 122, November 2013) 115.

31. For the Author's earlier works on copyright reforms, see generally Peter K. Yu, *The Confuzzling Rhetoric Against New Copyright Exceptions*, in 1 KRITIKA: ESSAYS ON INTELLECTUAL PROPERTY 278 (Peter Drahos et al. eds., 2015); Yu, *Fair Use*, *supra* note 18; Peter K. Yu, *Can the Canadian UGC Exception Be Transplanted Abroad?*, 26 INTELL. PROP. J. 175 (2014); Peter K. Yu, *Digital Copyright Reform and Legal Transplants in Hong Kong*, 48 U. LOUISVILLE L. REV. 693 (2010) [hereinafter Yu, *Digital Copyright Reform*]; Peter K. Yu, *Fair Use and Its Global Paradigm Evolution*, 2019 U. ILL. L. REV. 111, 129–37 [hereinafter Yu, *Global Paradigm Evolution*]; Yu, *The Quest*, *supra* note 27. During the last round of digital copyright reform in Hong Kong, the Author served as a pro bono advisor to Internet user groups and pan-Democrat legislators. *Id.* at 285.

to engage with fair use questions, at times with the help of university librarians. Even if these researchers do not have sufficient copyright expertise, they will have at least some familiarity with this area of the law. Even better, they will have experienced both the benefits of fair use and the potential struggle in drawing precise legal conclusions. Because research in the artificial intelligence area is highly multi- and inter-disciplinary, picking an illustration that is familiar to a wide range of scholars, not just those in the legal discipline, will be conducive to future research.

Third, the longstanding tradition and tremendous complexity of the U.S. fair use standard will allow readers and researchers to see the benefits, drawbacks, and challenges of legal automation. Even better, this standard involves both statutory and case law. While the standard itself has been codified in Section 107 of the Copyright Act,³² its interpretations have evolved over the past century in common law.³³ Moreover, because fair use determinations are made on a case-by-case basis by reference to four statutorily stipulated factors, the study of fair use automation will help illustrate the impact of artificial intelligence on the operation of rules and standards.³⁴ Understanding this impact is important, in view of both the prevailing wisdom that automating rules are easier than standards³⁵ and the recent literature on how legal automation will greatly reduce the trade-offs between legal rules and standards.³⁶

32. 17 U.S.C. § 107.

33. The case law on the uncoded fair use doctrine can be traced back to the 1841 case of *Folsom v. Marsh*, 9 F. Cas. 342, 348 (C.C.D. Mass. 1841) (No. 4,901).

34. As the Australian Law Reform Commission declared in its final report:

The flexibility of fair use largely comes from the fact that it is a standard, rather than a rule. This distinction between rules and standards is commonly drawn in legal theory. Rules are more specific and prescribed. Standards are more flexible and allow decisions to be made at the time of application, and with respect to a concrete set of facts. Further, ‘standards are often based on concepts that are readily accessible to non-experts’.

Rules and standards are, however, points on a spectrum. Rules are ‘not infinitely precise, and standards not infinitely vague’. The legal philosopher H L A Hart wrote that rules have ‘a core of certainty and a penumbra of doubt’. The distinction is nevertheless useful.

Australian Law Reform Commission, *supra* note 30, at 98; *see also id.* at 98–100 (discussing rules and standards in the fair use context).

35. As Dan Burk observed:

[T]he ex ante indeterminacy of a legal standard such as fair use, which in the institutional operation of the law constitutes a benefit, presents a challenge for operational machine coding. Rule-oriented legal imperatives may better lend themselves to automated instructions.

Burk, *Algorithmic Fair Use*, *supra* note 25, at 292 (footnote omitted).

36. As Anthony Casey and Anthony Niblett observed:

[T]echnological advances in predictive and communication technologies will render th[e] trade-off between rules and standards unnecessary. A new form of law, the microdirective, will emerge to provide all of the benefits of both rules and standards without the costs of either. These microdirectives will provide ex ante behavioral prescriptions finely tailored to every possible scenario.

Fourth, although the case study of fair use automation utilizes a legal standard grounded in U.S. law, this standard has received wide and ever-growing international support and recognition.³⁷ At the time of writing, the U.S. fair use standard has been transplanted abroad—in either identical or hybrid form—in a number of jurisdictions, including “Israel, Liberia, Malaysia, the Philippines, Singapore, South Korea, Sri Lanka, and Taiwan.”³⁸ The case study of fair use automation will therefore allow us to think more deeply about the global and cross-jurisdictional impact of legal automation.

Finally, fair use automation is not as far-fetched as other proposals or thought experiments involving legal automation. In the past decade, the copyright industry and their supportive technology platforms have already actively deployed automated copyright enforcement to identify, monitor, filter, and monetize potentially infringing works on digital networks.³⁹ While YouTube’s Content ID system provides a paradigmatic example,⁴⁰ other platforms have deployed similar tools and algorithms to facilitate such enforcement.⁴¹ To the extent these platforms aim to develop automated

Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L.J. 1401, 1403 (2017); *see also* Casey & Niblett, *Self-Driving Laws*, *supra* note 16, at 433 (discussing how the development and automatic updating of micro-directives will move us toward “a world of self-driving laws”).

37. *See* Yu, *Global Paradigm Evolution*, *supra* note 31, at 129–37 (documenting a growing trend toward the worldwide adoption of the U.S. fair use model and a slowly emerging paradigm evolution of international copyright norms); *see also* Peter K. Yu, *Customizing Fair Use Transplants*, 7 LAWS, no. 1, art. 9, at 3–10 (2018), <http://www.mdpi.com/2075-471X/7/1/9> (discussing the efforts to transplant fair use across the world and the eight different modalities of transplantation that the transplanting jurisdictions have employed). *See generally* JONATHAN BAND & JONATHAN GERAFI, *THE FAIR USE/FAIR DEALING HANDBOOK* (2013), <http://ssrn.com/abstract=2333863> (listing the fair use or fair dealing provisions from around the world).

38. Yu, *Global Paradigm Evolution*, *supra* note 31, at 115.

39. For discussions of algorithmic copyright enforcement, *see generally* Maayan Perel & Niva Elkin-Koren, *Accountability in Algorithmic Copyright Enforcement*, 19 STAN. TECH. L. REV. 473 (2016) [hereinafter Perel & Elkin-Koren, *Algorithmic Copyright Enforcement*]; Maayan Perel & Niva Elkin-Koren, *Black Box Tinkering: Beyond Disclosure in Algorithmic Enforcement*, 69 FLA. L. REV. 181, 189 (2017) [hereinafter Perel & Elkin-Koren, *Black Box Tinkering*].

40. *See generally* *How Content ID Works*, YOUTUBE HELP, <https://support.google.com/youtube/answer/2797370?hl=en> (last visited Sept. 13, 2020) (providing an overview of YouTube’s Content ID system). For discussions of the Content ID system, *see generally* Perel & Elkin-Koren, *Algorithmic Copyright Enforcement*, *supra* note 39, at 509–16; Matthew Sag, *Internet Safe Harbors and the Transformation of Copyright Law*, 93 NOTRE DAME L. REV. 499, 543–60 (2017) [hereinafter Sag, *Internet Safe Harbors*].

41. As Matthew Sag observed:

[D]espite the lack of a de jure obligation to filter under the DMCA [Digital Millennium Copyright Act], many platforms—typically large-scale commercial enterprises—are nonetheless implementing automated copyright enforcement systems. At the present time, platforms using automated copyright enforcement include Scribid, 4shared, Dropbox, YouTube, Facebook, SoundCloud, Twitch, TuneCore, Tumblr, Veoh, and Vimeo. The pressure to adopt automated filtering comes primarily from rightsholders, but these systems also meet some of the business objectives of platforms.

Sag, *Internet Safe Harbors*, *supra* note 40, at 538–39 (footnotes omitted); *see also* NICOLAS P. SUZOR, *LAWLESS: THE SECRET RULES THAT GOVERN OUR DIGITAL LIVES* 72 (2019) (“Automated copyright detection systems have now been built into many other services on the internet. Facebook has developed its own detection systems, and companies like Audible Magic produce software that has been adopted by many platforms.”); Burk, *Algorithmic Fair Use*, *supra* note 25, at 284 (“In the area of copyright, protection of digitized

enforcement systems that are consistent with existing copyright law, it is expected that some form of fair use has already been built into these systems.⁴² Moreover, a growing number of commentators have now called for greater algorithmic deployment to promote fair use in copyright law.⁴³ To them, automation is a much-needed solution demanded by the fast pace of digital dissemination and the exceedingly large volume of distributed content.⁴⁴

C. Resistance Toward Automation

At the time of writing, there have been three dominant arguments against greater fair use automation: (1) the relatively backward state of technology is unable to support satisfactory fair use automation; (2) the development of automated fair use systems will change creative choices and practices; and (3) experts have documented biases, bugs, and other problems in automated systems and artificial intelligence technologies, both within and outside the intellectual property area. This Subpart discusses and responds to each argument in turn in the hope of explaining why greater fair use automation is both urgently needed and socially beneficial.

works is already increasingly mediated by algorithmic enforcement systems that are intended to effectuate the rights of copyright owners while simultaneously limiting the liability of content intermediaries.”).

42. See 17 U.S.C. § 512(c)(3)(A)(v) (requiring “[a] statement that the complaining party has a good faith belief that use of the material in the manner complained of is not authorized by the copyright owner, its agent, or the law”); *Lenz v. Universal Music Corp.*, 572 F. Supp. 2d 1150, 1155 (N.D. Cal. 2008) (“A consideration of the applicability of the fair use doctrine simply is part of that initial review [of the potentially infringing material prior to sending a takedown notice as required by Section 512(c) of the Copyright Act].”).

43. See Timothy K. Armstrong, *Digital Rights Management and the Process of Fair Use*, 20 HARV. J.L. & TECH. 49, 56 (2006) (“[Digital rights management] mechanisms engineered to protect fair use rights are in the long-term interests of both content providers and consumers.”); Niva Elkin-Koren, *Fair Use by Design*, 64 UCLA L. REV. 1082, 1085 (2017) (“[T]he checks that [fair use] intends to create on the rights of authors must . . . be embedded in the design of online systems.”); Sag, *Internet Safe Harbors*, *supra* note 40, at 531–32 (“[T]here is no reason in principle why matching algorithms could not be fine-tuned to identify common situations associated with a higher probability of fair use.”); Peter K. Yu, *Anticircumvention and Anti-anticircumvention*, 84 DENV. U. L. REV. 13, 63 (2006) [hereinafter Yu, *Anticircumvention and Anti-anticircumvention*] (“The fact that the scope and boundaries of [fair use] are uncertain and that software code at the current state of technology may not be able to capture the full range of exceptions and limitations in the copyright system does not mean that we should not build legitimate uses into the [digital rights management] systems.”); Yu, *Fair Use*, *supra* note 18, at 338–50 (building the case for greater algorithmic deployment to promote fair use in U.S. copyright law); Barbara L. Fox & Brian A. LaMacchia, *Encouraging Recognition of Fair Uses in DRM Systems*, COMM. ACM, Apr. 2003, at 61, 63 (“[The limitation on developing a perfect mathematical model of fair use] should not stop us from attempting to identify a useful subset we might approximate in code.”).

44. See Elkin-Koren, *supra* note 43, at 1098 (“The need to address the sheer volume of copyright disputes requires a new approach to fair use that involves rethinking the role of legal oversight in algorithmic adjudication.”); Sag, *Internet Safe Harbors*, *supra* note 40, at 554 (“With over 400 hours of video being uploaded to YouTube every minute, it is hard to imagine that either rightsholders . . . or the platform itself . . . could meaningfully prevent the evisceration of online copyright without relying on automation to some extent.”); *see also id.* at 513 (“In 2016, YouTube users were uploading 400 hours of video content every minute . . .”); Jeff Desjardins, *How Much Data Is Generated Each Day?*, WORLD ECON. F. (Apr. 17, 2019), <https://www.weforum.org/agenda/2019/04/how-much-data-is-generated-each-day-cf4bddf29f> (“By 2025, it’s estimated that 463 exabytes of data will be created each day globally—that’s the equivalent of 212,765,957 DVDs per day!”).

1. *Backward State of Technology*

The first dominant argument concerns our relatively backward state of technology, which commentators believe is inadequate to support satisfactory fair use automation.⁴⁵ At the turn of the millennium, when the copyright and technology industries, policymakers, and legal experts were exploring whether fair use could be built into digital rights management systems, Edward Felten warned us bluntly that we did not yet and might never have a “judge on a chip.”⁴⁶ As he observed at that time: “Fair use is one of the starkest examples of the mismatch between what the law requires and what technology can do. Accurate, technological enforcement of the law of fair use is far beyond today’s state of the art and may well remain so permanently.”⁴⁷ Writing around that time, Dan Burk and Julie Cohen also observed, “At least for now, there is no feasible way to build rights management code that approximates both the individual results of judicial determinations and the overall dynamism of fair use jurisprudence.”⁴⁸

While these scholars were right to identify the technological barriers to developing satisfactory automated fair use systems, it remains debatable whether incremental steps can be taken to build these systems.⁴⁹ After all,

45. See Yu, *Fair Use*, *supra* note 18, at 331–33 (discussing our relatively backward state of technology as a major argument against the satisfactory deployment of algorithms to promote fair use).

46. See Edward W. Felten, *A Skeptical View of DRM and Fair Use*, COMM. ACM, Apr. 2003, at 57, 58 (“A [digital rights management] system that gets all fair use judgments right would in effect be a ‘judge on a chip’ predicting with high accuracy how a real judge would decide a lawsuit challenging a particular use. Clearly, this is infeasible with today’s technology.”); see also Burk & Cohen, *supra* note 25, at 59 (“At present, only human intelligence, reviewing the unique circumstances of a particular use, can determine whether it is likely to be fair.”).

47. Felten, *supra* note 46, at 59; see also JULIE E. COHEN, BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM 192 (2019) (“Automated processes have obvious efficiency advantages, but such processes may not align well (or at all) with applicable legal requirements that are couched in shades of gray.”); Ian R. Kerr et al., *Technical Protection Measures: Tilting at Copyright’s Windmill*, 34 OTTAWA L. REV. 7, 31 (2002) (“[T]he technologies employed by [digital rights management systems] are not yet sufficiently sophisticated to mirror the law of copyright because [technological protection measures] themselves remain incapable of distinguishing between infringing and non-infringing uses of digital works.”); Mark A. Lemley, *Rationalizing Internet Safe Harbors*, 6 J. ON TELECOMM. & HIGH TECH. L. 101, 110–11 (2007) (“Image-parsing software may someday be able to identify pictures or videos that are similar to individual copyrighted works, but they will never be able to determine whether those pictures are fair uses, or whether they are legitimate copies or displays made under one of the many statutory exceptions . . .”).

48. Burk & Cohen, *supra* note 25, at 56.

49. Dan Burk and Julie Cohen expressed concern that the development of automated fair use systems would encourage minimalist interpretations of important safeguards and the establishment of ceilings for these safeguards:

We are . . . skeptical . . . about the ability of negotiated [technical] defaults to capture the full range of social benefit that more flexible legal standards allow. While these defaults sometimes might allow access that would exceed fair use under a judicial determination, the “safe harbor” concept is more likely to tend toward a minimalist view of fair use. We suspect that copyright holders would be willing to concede fair use in only a small fraction of the situations that would constitute fair use—indeed, it was just such insistence upon minimalist guidelines by rights holders that led to the collapse of the [Conference on Fair Use] discussions. Moreover, in the case of the 1976 “safe harbor” guidelines for educational copying, rights holders, content users, and even courts

building technological systems takes time, and there will always be a less-than-ideal transitional period. As Microsoft software architects Barbara Fox and Brian LaMacchia declared in the early 2000s:

[The limitation that no one can mathematically model fair use, as it is understood today,] should not stop us from attempting to identify a useful subset we might approximate in code. That is, we can take a purely pragmatic engineering approach . . . : Focus first on defining and modeling a useful subset of fair use rights in some policy language, then add these expressions to the policy evaluators of [digital rights management] systems.⁵⁰

In an article written in the mid-2000s, I also noted the need to distinguish between limitations and exceptions that can be interpreted by machines from those that cannot.⁵¹ As I explained at that time:

The fact that the scope and boundaries of [fair use] are uncertain and that software code at the current state of technology may not be able to capture the full range of exceptions and limitations in the copyright system does not mean that we should not build legitimate uses into the [digital rights management] systems.⁵²

have shown a deplorable tendency to act as though the guidelines defined the outer limits of fair use. To the contrary, such guidelines were intended to delineate fair use minima: a floor rather than a ceiling. We are consequently reluctant to recommend an infrastructure based solely on the design of similar defaults into self-enforcing “lock-out” systems for fear that the “ceiling” effect could be even more pernicious.

Burk & Cohen, *supra* note 25, at 57 (footnotes omitted); *see also* Elkin-Koren, *supra* note 43, at 1096 (“The main concern is that reducing the four-factor analysis into a simplistic and somewhat rigid set of algorithmic instructions might cause some important aspects of fair use analysis to get lost along the way.”).

50. Fox & LaMacchia, *supra* note 43, at 63. Professor Sag concurred:

The difficulty of completely automating fair use analysis does not suggest . . . that algorithms have no role to play. Experience, common sense, and recent empirical research suggest that there are some objective characteristics that make a finding of fair use more likely, and there is no reason in principle why matching algorithms could not be fine-tuned to identify common situations associated with a higher probability of fair use.

Sag, *Internet Safe Harbors*, *supra* note 40, at 531–32. Likewise, Timothy Armstrong observed:

The flaw in the conclusion that [digital rights management] cannot accommodate fair use is an unduly hasty inductive leap from the specific (the impossibility of modeling the substance of fair use law in machine-administrable form) to the general (the supposed impossibility of protecting fair use at all in [digital rights management] systems). The foreclosure of one avenue for protecting fair use, however, does not imply that all avenues are likewise foreclosed, but only that design principles other than the creation of a perfect “judge on a chip” must be explored.

Armstrong, *supra* note 43, at 88.

51. *See* Yu, *Anticircumvention and Anti-anticircumvention*, *supra* note 43, at 63–73 (discussing the need for such a distinction); *see also* Deirdre Mulligan & Aaron Burstein, *Implementing Copyright Limitations in Rights Expression Languages*, in *DIGITAL RIGHTS MANAGEMENT: ACM WORKSHOP ON DIGITAL RIGHTS MANAGEMENT, DRM 2002, WASHINGTON, DC, USA, NOVEMBER 18, 2002: REVISED PAPERS 137* (Joan Feigenbaum ed., 2002) (discussing ways and challenges to implementing copyright limitations and exceptions in rights expression languages, with a focus on XrML, the eXtensible Rights Markup Language); Fox & LaMacchia, *supra* note 43, at 63 (considering the importance of determining “how to create machine-interpretable expressions that adequately model a set (or subset) of fair use rights”).

52. Yu, *Anticircumvention and Anti-anticircumvention*, *supra* note 43, at 63.

Moreover, the landscape of copyright enforcement has changed substantially in the past decade. As noted earlier, the copyright industries and technology platforms have already widely deployed algorithms to facilitate copyright enforcement.⁵³ If fair use is not built, or sufficiently built, into these algorithms—or if we do not develop what Niva Elkin-Koren has coined “fair use by design”⁵⁴—the balance in the copyright system will shift too much toward the interests of copyright holders to the disadvantage of individual users.⁵⁵ Fearing the violation of copyright law, many risk-averse users may forgo their socially productive creative endeavors.⁵⁶ Those who constantly have to test the limits of copyright law may also lose respect for the law,⁵⁷ viewing it instead as an illegitimate product of industry capture.⁵⁸

2. *Changes in Creative Choices and Practices*

The second dominant argument relates to the changes in creative choices and practices that will be generated by the development of automated fair use systems.⁵⁹ In a recent article, Dan Burk expressed fear that algorithmic fair use would create considerable biases, which in turn would affect authorial choices.⁶⁰ As he lamented: “[T]he design values embedded in automated systems become embedded in public behavior and consciousness. Thus, algorithmic fair use carries with it the very real possibility of habituating new media participants to its own biases and so progressively altering the fair use standard it attempts to

53. See sources cited *supra* note 39.

54. Elkin-Koren, *supra* note 43, at 1100.

55. See *id.* (“Fair use by design has become a necessity in an era of algorithmic governance. The need to develop such tools is necessary in order to tilt the copyright balance back to its origin in our robo notice environment.”); Burk, *Algorithmic Fair Use*, *supra* note 25, at 284–85 (“[I]t may seem desirable to incorporate context-specific fair use metrics into copyright-policing algorithms, both to protect against automated over-deterrence and to inform users of their compliance with copyright law.”).

56. See Burk, *Algorithmic Fair Use*, *supra* note 25, at 288 (“Risk averse content users, unable to confidently predict the ultimate decision on their activities, may forgo some socially beneficial uses.”); Elkin-Koren, *supra* note 43, at 1100 (“The high cost and high risk involved in fair use implementation prevents users from taking advantage of productive uses that can foster copyright goals, simply because they fear liability.”); Yu, *Fair Use*, *supra* note 18, at 349 (“If automated fair use determinations can have legal effects—even if only on an interim basis—those determinations can enlarge the creative spaces of risk-averse users, some of whom may fear that their creative endeavors will violate current copyright law.”).

57. Cf. Armstrong, *supra* note 43, at 109 (“Empowering users to exercise their fair use rights without violating the DMCA might . . . increase law-abiding behavior and temper the critical evaluation of the DMCA as a one-sided giveaway to powerful producer cartels.” (footnote omitted)).

58. See generally MONICA HORTEN, A COPYRIGHT MASQUERADE: HOW CORPORATE LOBBYING THREATENS ONLINE FREEDOMS (2013) (discussing how legislative capture by the copyright industries has undermined online freedom); BRINK LINDSEY & STEVEN M. TELES, THE CAPTURED ECONOMY: HOW THE POWERFUL ENRICH THEMSELVES, SLOW DOWN GROWTH, AND INCREASE INEQUALITY 64–89 (2017) (discussing industry capture in the intellectual property area).

59. See Yu, *Fair Use*, *supra* note 18, at 334–35 (discussing the potential changes in creative choices and practices as a major argument against the satisfactory deployment of algorithms to promote fair use).

60. See Burk, *Algorithmic Fair Use*, *supra* note 25, at 285.

embody.”⁶¹ Because of the inevitable entanglement between algorithms and the users’ creative practices, the development of automated fair use systems will cause behavioral changes that will eventually generate new legal norms.⁶² In turn, the development of these new norms and practices will degrade the fair use standard into “an unrecognizable form.”⁶³ Such development will also initiate “a self-reinforcing cycle” in which “[the] increasing use of AI adjudication will foster changes in values that are conducive to even greater use of AI adjudication.”⁶⁴

Professor Burk was right that the development of automated fair use systems will likely foster changes in creative choices and practices, and his observation was well supported by the behavioral changes we have already seen among those Internet and social media users who manipulated or circumvented the algorithms deployed by copyright holders and technology platforms.⁶⁵ However, behavioral changes are inevitable whenever decisions are made. As I noted in a recent symposium, “The key question about automated fair use systems is . . . not whether these systems will make decisions, but whether they will make worse decisions, or make worse decisions more frequently.”⁶⁶ If machine-made decisions are just as good as those made by human decisionmakers, such as judges or law enforcement personnel, the public will find machine-made decisions less problematic even if they are to induce changes in user behavior.

Moreover, there is hitherto insufficient evidence to show whether automated decisions will help creators more than they will hurt them. For risk-averse creators, having low-cost fair use determinations in real time will

61. *Id.*

62. See Tarleton Gillespie, *The Relevance of Algorithms*, in *MEDIA TECHNOLOGIES: ESSAYS ON COMMUNICATION, MATERIALITY, AND SOCIETY* 167, 183 (Tarleton Gillespie et al. eds., 2014) (discussing the entanglement between algorithms and social practices).

63. See Burk, *Algorithmic Fair Use*, *supra* note 25, at 306 (“[A]ttempting to incorporate fair use into enforcement algorithms threatens to degrade the exception into an unrecognizable form. Worse yet, social internalization of a bowdlerized version of fair use deployed in algorithmic format is likely to become the new legal and social norm.”).

64. Re & Solow-Niederman, *supra* note 8, at 247; see also *id.* at 249–52 (discussing how AI-driven developments will affect the ways humans interact with and relate to the law and the judiciary).

65. See Jane Bambauer & Tal Zarsky, *The Algorithm Game*, 94 NOTRE DAME L. REV. 1, 12–14 (2018) (listing avoidance, altered conduct, altered input, and obfuscation among the dominant gaming strategies deployed by users on Internet platforms); Caleb Garling, *Tricking Facebook’s Algorithm*, ATLANTIC (Aug. 8, 2014), <https://www.theatlantic.com/technology/archive/2014/08/tricking-facebooks-algorithm/375801> (discussing the experience of tricking Facebook to elevate the author’s post); Anjana Susarla, *The New Digital Divide Is Between People Who Opt Out of Algorithms and People Who Don’t*, CONVERSATION (Apr. 17, 2019, 6:54 AM), <https://theconversation.com/the-new-digital-divide-is-between-people-who-opt-out-of-algorithms-and-people-who-dont-114719> (“A study of Facebook usage found that when participants were made aware of Facebook’s algorithm for curating news feeds, about 83% of participants modified their behavior to try to take advantage of the algorithm, while around 10% decreased their usage of Facebook.”); Tony Zhou, *Postmortem: Every Frame a Painting*, MEDIUM (Dec. 2, 2017), <https://medium.com/@tonyzhou/postmortem-1b338537fab6>, quoted in Burk, *Algorithmic Fair Use*, *supra* note 25, at 303 (explaining how the author and his partner edited around YouTube’s Content ID system by making trial-and-error adjustments).

66. Yu, *Fair Use*, *supra* note 18, at 354.

likely be highly beneficial.⁶⁷ Because the U.S. fair use system requires courts to make determinations *ex post*, those users who do not have sufficient economic resources to hire copyright lawyers to test the law's boundaries may choose not to make socially productive use of copyrighted works in the first place.⁶⁸ By providing a helpful safe harbor, greater fair use automation can provide important benefits to creators—and, by extension, society.

3. *Technological Shortcomings*

The third dominant argument pertains to the biases, bugs, and other documented problems now found in automated systems and artificial intelligence technologies.⁶⁹ The technological problems in this area are not limited to fair use automation; they have been widely documented outside the intellectual property area. For instance, ProPublica published a widely praised exposé on the racial biases found in COMPAS, the scoring software used by law enforcement and correction personnel to determine risks of recidivism.⁷⁰ As the investigatory report stated, “black defendants were far more likely than white defendants to be incorrectly judged to be at a higher risk of recidivism, while white defendants were more likely than black defendants to be incorrectly flagged as low risk.”⁷¹ In addition, the media provided wide coverage of how Microsoft’s Twitter bot Tay had quickly become sexist and racist because its “algorithms . . . had [the bot] ‘learning’ how to respond to others based on what was tweeted at it.”⁷² Another report stated that Hewlett-Packard’s facial recognition technology had failed to properly recognize African-Americans

67. *See id.* (discussing the benefits of automated systems in providing low-cost fair use determinations); *see also* Burk, *Algorithmic Fair Use*, *supra* note 25, at 289 (“Automated identification and removal, whether accurate or mistaken, is relatively cheap, whereas legal and institutional engagement is comparatively expensive.”); Volokh, *supra* note 8, at 1147 (“Realistically, the only way we are likely to sharply increase access to expensive services, such as lawyering, is through technology.”).

68. *See* LAWRENCE LESSIG, *FREE CULTURE: HOW BIG MEDIA USES TECHNOLOGY AND THE LAW TO LOCK DOWN CULTURE AND CONTROL CREATIVITY* 187 (2004) (“[F]air use in America simply means the right to hire a lawyer to defend your right to create.”).

69. *See* Yu, *Fair Use*, *supra* note 18, at 335–38 (discussing technological shortcomings as a major argument against the satisfactory deployment of algorithms to promote fair use); *see also* ANDREW MCAFEE & ERIK BRYNJOLFSSON, *MACHINE, PLATFORM, CROWD: HARNESSING OUR DIGITAL FUTURE* 53 (2017) (noting the “biases and bugs” in intelligent machines); Burk, *Algorithmic Fair Use*, *supra* note 25, at 285 (listing “ersatz objectivity, diminished decisional transparency, and design biases” among the inherent pitfalls in reliance on algorithmic regulation); Peter K. Yu, *The Algorithmic Divide and Equality in the Age of Artificial Intelligence*, 72 FLA. L. REV. 331, 354–61 (2020) [hereinafter Yu, *Algorithmic Divide*] (discussing algorithmic discrimination and distortion).

70. Jeff Larson et al., *How We Analyzed the COMPAS Recidivism Algorithm*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>. COMPAS stands for “Correctional Offender Management Profiling for Alternative Sanctions.” *Id.*

71. *Id.*

72. LEE RAINIE & JANNA ANDERSON, *CODE-DEPENDENT: PROS AND CONS OF THE ALGORITHM AGE 2* (2017), <http://www.pewinternet.org/2017/02/08/code-dependent-pros-and-cons-of-the-algorithm-age>.

because the “[c]ameras on [its] new . . . computers did not track the faces of Black people in some common lighting conditions.”⁷³

Even worse, commentators have shown that automated systems will “disproportionately affect groups that are already disadvantaged by factors such as race, gender and socio-economic background.”⁷⁴ When learning algorithms—or so-called “learners”⁷⁵—are deployed, the harm to these disadvantaged groups could be even greater, considering that the problematic algorithmic outcomes will be fed back into the automated systems as training data. Such repeated use of data will create self-reinforcing feedback loops that amplify the biases found in the initial algorithms or training data.⁷⁶ Until these biases are corrected, the initial biases will be greatly magnified.⁷⁷

As we build automated fair use systems and make the needed adjustments to improve them, having problems in the transitional period is inevitable. The fact that we have problems in the current iterations of the automated systems does not mean that we should refrain from using these systems in the first place. It only means that we have to be careful about such usage, be active in

73. Christian Sandvig et al., *When the Algorithm Itself Is a Racist: Diagnosing Ethical Harm in the Basic Components of Software*, 10 INT’L J. COMM. 4972, 4973 (2016) (citations omitted).

74. Kate Crawford & Ryan Calo, *There Is a Blind Spot in AI Research*, NATURE (Oct. 13, 2016), <https://www.nature.com/news/there-is-a-blind-spot-in-ai-research-1.20805>; see also VIRGINIA EUBANKS, AUTOMATING INEQUALITY: HOW HIGH-TECH TOOLS PROFILE, POLICE, AND PUNISH THE POOR 12 (2017) (lamenting how “[a]utomated decision-making shatters the social safety net, criminalizes the poor, intensifies discrimination, and compromises our deepest national values”); CATHY O’NEIL, WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY 8 (2016) (noting that algorithm-driven automated systems “tend to punish the poor . . . because they are engineered to evaluate large numbers of people”); RAINIE & ANDERSON, *supra* note 72, at 63–65 (surveying views on whether the disadvantaged will lag behind even further in this algorithmic age). See generally SAFIYA UMOJA NOBLE, ALGORITHMS OF OPPRESSION: HOW SEARCH ENGINES REINFORCE RACISM (2018) (discussing how search engines promote racism and sexism).

75. See PEDRO DOMINGOS, THE MASTER ALGORITHM: HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD 6 (2015) (“Learning algorithms—also known as learners—are algorithms that make other algorithms. With machine learning, computers write their own programs, so we don’t have to.”).

76. As Ronald Yu and Gabriele Spina Ali observed:

[T]here is a strong risk that AI may reiterate and even amplify the biases and flaws in datasets, even when these are unknown to humans. In this sense, AI has a self-reinforcing nature, due to the fact that the machine’s outputs will be used as data for future algorithmic operations.

Ronald Yu & Gabriele Spina Ali, *What’s Inside the Black Box? AI Challenges for Lawyers and Researchers*, 19 LEGAL INFO. MGMT. 2, 4 (2019) (footnote omitted); see also Sofia Grafanaki, *Autonomy Challenges in the Age of Big Data*, 27 FORDHAM INTELL. PROP. MEDIA & ENT. L.J. 803, 827 (2017) (noting that “algorithmic self-reinforcing loops are now present across many spheres of our daily life (e.g., retail contexts, career contexts, credit decisions, insurance, Google search results, news feeds)”); Sonia K. Katyal, *Private Accountability in the Age of Artificial Intelligence*, 66 UCLA L. REV. 54, 69 (2019) (“Bad data . . . can perpetuate inequalities through machine learning, leading to a feedback loop that replicates existing forms of bias, potentially impacting minorities as a result.”); *Digital Decisions*, CTR. FOR DEMOCRACY & TECH., <https://cdt.org/files/2018/09/Digital-Decisions-Library-Printer-Friendly-as-of-20180927.pdf> (“Unreliable or unfair decisions that go unchallenged can contribute to bad feedback loops, which can make algorithms even more likely to marginalize vulnerable populations.”).

77. See Yu, *Algorithmic Divide*, *supra* note 69, at 359 (“As time passes, the biases generated through these loops will become much worse than the biases found in the original algorithmic designs or the initial training data.”).

undertaking cost-benefit analyses, and be ready to provide updates or corrections when problems arise.⁷⁸

For example, in view of the problems found in automated systems and artificial intelligence technologies, commentators have called for efforts to make algorithmic designs more transparent by requiring audits⁷⁹ or regulatory oversight.⁸⁰ Such transparency is badly needed considering that the algorithms involved are often locked in so-called “black box” systems.⁸¹ Commentators have also noted the importance of human intervention.⁸² Even though

78. See BROWNSWORD, *supra* note 12, at 297 (calling for “the regulatory framework [to] provide for the correction of the malfunction” in the technology); Yu, *Algorithmic Divide*, *supra* note 69, at 379–80 (calling for the development of a “notice and correct” mechanism to address problems generated by automated systems).

79. As the Center for Democracy and Technology noted:

Audits are one method to provide explanations and redress without compromising the intellectual property behind the business model. Designing algorithmic systems that can be easily audited increases accountability and provides a framework to standardize best practices across industries. While explanations can help individuals understand algorithmic decision making, audits are necessary for systemic and long-term detection of unfair outcomes. They also make it possible to fix problems when they arise.

Digital Decisions, *supra* note 76; see also Deven R. Desai & Joshua A. Kroll, *Trust but Verify: A Guide to Algorithms and the Law*, 31 HARV. J. L. & TECH. 1, 37–42 (2017) (discussing ways to test and evaluate algorithms); Pauline T. Kim, *Auditing Algorithms for Discrimination*, 166 U. PA. L. REV. ONLINE 189 (2017) (discussing the use of audits as a check against discrimination); Yu, *Algorithmic Divide*, *supra* note 69, at 380–82 (discussing the need for algorithmic audits).

80. See Yu, *Algorithmic Divide*, *supra* note 69, at 380 (discussing the need for institutional oversight); see also INST. ELEC. & ELEC. ENG’RS, ETHICALLY ALIGNED DESIGN: A VISION FOR PRIORITIZING HUMAN WELL-BEING WITH AUTONOMOUS AND INTELLIGENT SYSTEMS 70 (2017) (“An independent, internationally coordinated body . . . should be formed to oversee whether [autonomous and intelligent systems] actually meet ethical criteria, both when . . . deployed, and considering their evolution after deployment and interaction with other products.”); Frank Pasquale, *Restoring Transparency to Automated Authority*, 9 J. ON TELECOMM. & HIGH TECH. L. 235, 247 (2011) (“[P]erhaps a trusted advisory committee within the Federal Trade Commission could help courts and agencies adjudicate coming controversies over search engine practices.”).

81. See EUBANKS, *supra* note 74, at 5 (“[T]hat’s the thing about being targeted by an algorithm: you get a sense of a pattern in the digital noise, an electronic eye turned toward *you*, but you can’t put your finger on exactly what’s amiss.”); FRANK PASQUALE, THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION 3 (2015) (“[W]orkings [in black box systems] are mysterious; we can observe [their] inputs and outputs, but we cannot tell how one becomes the other.”); RAINIE & ANDERSON, *supra* note 72, at 19 (“There is a larger problem with the increase of algorithm-based outcomes beyond the risk of error or discrimination—the increasing opacity of decision-making and the growing lack of human accountability.” (quoting Marc Rotenberg, Executive Director, Electronic Privacy Information Center)). For book-length treatments of the problems generated by “black box” algorithms, see generally EUBANKS, *supra* note 74; O’NEIL, *supra* note 74; PASQUALE, *supra*.

82. Professors Casey and Niblett, for example, noted the continuous role of humans in algorithmic development:

Algorithmic decision-making does not mean that humans are shut out of the process. Even after the objective has been set, there is much human work to be done. Indeed, humans are involved in all stages of setting up, training, coding, and assessing the merits of the algorithm. If the objectives of the algorithm and the objective of the law are perfectly aligned at the ex ante stage, one must ask: Under what circumstances should a human ignore the algorithm’s suggestions and intervene *after* the algorithm has made the decision?

Anthony J. Casey & Anthony Niblett, *A Framework for the New Personalization of Law*, 86 U. CHI. L. REV. 333, 354 (2019) [hereinafter Casey & Niblett, *A Framework*]; see also Council Regulation 2016/679, art. 22(3), 2016 O.J. (L 119) 1, 46 (requiring data controllers to “implement suitable measures to safeguard the data subject’s

Professors Burk and Cohen were skeptical of the successful development of automated fair use systems, they advanced a proposal calling for “the introduction of an external [human] decisionmaker into the process for obtaining access to technologically secured works.”⁸³ In a proposal advanced more than a decade ago, I also advocated the “technology first, courts later” approach to enable courts to step in to provide the needed human intervention.⁸⁴

4. Summary

In sum, the development of automated fair use systems is still fraught with problems. Nevertheless, remedies do exist to address some of these problems. Moreover, technology will continue to improve. Compared with the turn of this century when commentators were actively debating whether fair use could be built into digital rights management systems, the technology and data that have become available today to build automated fair use systems are already very different. As Professor Elkin-Koren reminded us:

Overall, th[e] concerns regarding the limitations of algorithmic fair use overlook recent developments in Artificial Intelligence . . . and machine learning capabilities. AI has already been applied in very sophisticated contexts: physicians use algorithms to guide their diagnoses; banks use them to decide when to approve a loan; security agencies use AI to identify risks; lawyers use them to perform due diligence; and even courts rely on algorithms for sentencing, by scoring the risk of the offender committing future crimes. AI has already been applied for decision-making processes in contexts that are

rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest [a decision based solely on automated processing, including profiling]”); Peter K. Yu, *Beyond Transparency and Accountability: Three Additional Features Algorithm Designers Should Build into Intelligent Platforms*, 13 NE. U. L. REV. (forthcoming 2020) (calling on technology platforms to build intervenability into algorithmic designs and operations). *See generally* Aziz Z. Huq, *A Right to a Human Decision*, 106 VA. L. REV. 611 (2020) (discussing whether individuals have a “right to a human decision”); Meg Leta Jones, *The Right to a Human in the Loop: Political Constructions of Computer Automation and Personhood*, 47 SOC. STUD. SCI. 216 (2017) (tracing the historical roots of “[t]he right to a human in the loop” back to rights that protect the dignity of data subjects).

83. Burk & Cohen, *supra* note 25, at 59; *see also* Armstrong, *supra* note 43, at 75 (recognizing the need to include “human involvement” to facilitate the consideration of “a greater level of complexity in the circumstances”); Dan L. Burk, *Legal and Technical Standards in Digital Rights Management Technology*, 74 FORDHAM L. REV. 537, 551 (2005) (“[T]echnological controls tend to be relatively blunt instruments for control of digital content, unable to accommodate copyright fair use without the re-introduction of human discretion.”).

84. *See* Yu, *Anticircumvention and Anti-anticircumvention*, *supra* note 43, at 73 (“[A] two-step approach—technology first, then courts—seems to be the best compromise we can have today, and it is worth considering developing such a system as we explore the next generation of [digital rights management] systems.”). Niva Elkin-Koren outlined a similar approach: “Algorithmic fair use could . . . involve a two-tier review. First, algorithmic screening would be performed and second, for cases which were flagged by the system, but were inconclusive, human review would be conducted.” Elkin-Koren, *supra* note 43, at 1098.

far more complex than fair use, involving critical issues of life and death, health, financial risks, and national security.⁸⁵

II. PATHWAYS FOR LEGAL AUTOMATION

In addressing the three dominant arguments against fair use automation, the previous Part has shown the benefits of greater legal automation in this area. If we are to proceed with such automation, we will need to think about the different paths that can be taken to automate the fair use standard. Although many pathways for legal automation exist, three stand out: (1) translation; (2) approximation; and (3) self-determination. While the first two pathways are built upon the existence of and reliance on human decisions, the last pathway allows for autonomous determinations, which can take place regardless of the existence or volume of human decisions. This Part discusses each pathway in turn and ties the discussion to the ongoing developments in artificial intelligence, machine learning, and big data analytics.

A. Translation

The first pathway for automating the fair use standard is translation. The scholarly engagement with the need to translate legal standards into computer code and algorithms is nothing new. When the Internet first entered the mainstream in the mid-1990s, a sizeable literature quickly emerged to discuss ways to faithfully translate laws in physical space to cyberspace. For instance, Lawrence Lessig reminded us that “code is law” and that algorithms could be built to reflect or ignore our constitutional values.⁸⁶ Focusing on what he coined “lex informatica,” the late Joel Reidenberg also called on policymakers to pay greater attention to the development of technology rules and to encourage such development.⁸⁷

While the discussion of the need for translation in the artificial intelligence context is a logical extension of this earlier cyberlaw debate, tremendous difficulties remain in the efforts to translate legal mandates into computer code and algorithms.⁸⁸ As far as legal automation is concerned, the developers of automated fair use systems have the daunting task of figuring out how to build

85. Elkin-Koren, *supra* note 43, at 1096–97 (footnotes omitted).

86. LAWRENCE LESSIG, CODE: VERSION 2.0, at 1 (2006).

87. Joel R. Reidenberg, *Lex Informatica: The Formulation of Information Policy Rules Through Technology*, 76 TEX. L. REV. 553 (1998).

88. See Perel & Elkin-Koren, *Algorithmic Copyright Enforcement*, *supra* note 39, at 486 (“Translating doctrinal law and policy into code may result in significant, albeit unintentional, alterations of meaning, partly because the artificial languages intelligible to computers have a more limited vocabulary than human languages.” (footnote omitted)).

legal rules and outcomes into these systems.⁸⁹ As Maayan Perel and Niva Elkin-Koren observed, “[T]ranslating legal mandates into code inevitably embodies particular choices as to how the law is interpreted, which may be affected by a variety of extrajudicial considerations, including the conscious and unconscious professional assumptions of program developers, as well as various private business incentives.”⁹⁰ In their earlier work, Dan Burk and Julie Cohen also expressed skepticism that “system designers will be able to anticipate the range of access privileges that may be appropriate for fair uses to be made of a particular work . . . [as well as] the types of uses that would be considered fair by a court.”⁹¹

Even worse, for a legal standard that courts will only interpret *ex post*, such as the U.S. fair use standard, computer programmers will have to determine in advance how the law will affect the outcome—often by making educated guesses. While adjudicated cases and their related fact patterns can provide helpful guidance, many situations will be of first impression and will therefore present substantial translational challenges and complications. A case in point is an interesting empirical experiment conducted by Lisa Shay, Woodrow Hartzog, John Nelson, and Gregory Conti.⁹² When they brought together three teams of computer programmers to translate a subset of the New York State traffic law into computer code for the purpose of determining traffic violations based on real-world driving data, they found wide variances in cited violations and citation frequency depending on whether the group followed the letter of the law, the intent of the law, or additional guidance and instructions from the experiment’s designers.⁹³

One solution that can help alleviate this type of translation-induced problem is to conduct periodic audits—both internally and externally—to determine whether the laws have been faithfully translated.⁹⁴ Such audits reflect the best practices advocated by the technology community. Principle 7 of the *ACM Statement on Algorithmic Transparency and Accountability* declared, “Institutions should use rigorous methods to validate their models and document those methods and results.”⁹⁵ The *FAT/ML Principles for Accountable*

89. See Lisa A. Shay et al., *Confronting Automated Law Enforcement*, in *ROBOT LAW*, *supra* note 16, at 257–59 (discussing the legal integration of algorithms); Antje von Ungern-Sternberg, *Autonomous Driving: Regulatory Challenges Raised by Artificial Decision-Making and Tragic Choices*, in *RESEARCH HANDBOOK*, *supra* note 6, at 251, 262–64 (discussing the need to translate law into algorithm).

90. Perel & Elkin-Koren, *Black Box Tinkering*, *supra* note 39, at 189; see also Shay et al., *supra* note 89, at 257 (“[T]hose who specify and implement the code base of a system will likely make their own interpretations of legal and illegal behavior, perhaps without any legal training.”).

91. Burk & Cohen, *supra* note 25, at 55.

92. Lisa A. Shay et al., *Do Robots Dream of Electric Laws? An Experiment in the Law as Algorithm*, in *ROBOT LAW*, *supra* note 16, at 274.

93. See *id.*

94. See sources cited *supra* note 79.

95. U.S. PUB. POL’Y COUNCIL, ASS’N FOR COMPUTING MACHINERY, STATEMENT ON ALGORITHMIC TRANSPARENCY AND ACCOUNTABILITY (2017), [hereinafter *ACM STATEMENT*].

Algorithms and a Social Impact Statement for Algorithms also called for impact assessment “(at least) three times during the design and development process: design stage, pre-launch, and post-launch.”⁹⁶ As Lorna McGregor, Daragh Murray, and Vivian Ng explained:

During the design and development stage, impact assessments should evaluate how an algorithm is likely to work, ensure that it functions as intended and identify any problematic processes or assumptions. This provides an opportunity to modify the design of an algorithm at an early stage, to build in . . . compliance—including monitoring mechanisms—from the outset, or to halt development if . . . concerns cannot be addressed. Impact assessments should also be conducted at the deployment stage, in order to monitor effects during operation. . . . [T]his requires that, during design and development, the focus should not only be on testing but steps should also be taken to build in effective oversight and monitoring processes that will be able to identify and respond to [problems] once the algorithm is deployed.⁹⁷

To promote transparency, commentators have called for greater disclosure of not only algorithms but also of training data and algorithmic outcomes.⁹⁸ While such disclosure will certainly help those who are technology savvy, it is often insufficient, especially for those who have difficulty understanding the computer code, training process, or selected data involved.⁹⁹ When learning algorithms are deployed, closely scrutinizing the initial algorithms alone is

96. Nicholas Diakopoulos et al., *Principles for Accountable Algorithms and a Social Impact Statement for Algorithms*, FAT/ML, <https://www.fatml.org/resources/principles-for-accountable-algorithms> (last visited Sept. 12, 2020). FAT/ML stands for “Fairness, Accountability, and Transparency in Machine Learning.” FAT/ML, <https://fatml.org> (last visited Sept. 12, 2020).

97. Lorna McGregor et al., *International Human Rights Law as a Framework for Algorithmic Accountability*, 68 INT’L & COMPAR. L.Q. 309, 330 (2019).

98. See O’NEIL, *supra* note 74, at 229 (“We have to learn to interrogate our data collection process, not just our algorithms.”); Anupam Chander, *The Racist Algorithm?*, 115 MICH. L. REV. 1023, 1024–25 (2017) (“What we need instead is a *transparency of inputs and results*, which allows us to see that the algorithm is generating discriminatory impact.”); Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633, 641 (2017) (“[W]ithout full transparency—including source code, input data, and the full operating environment of the software—even the disclosure of audit logs showing what a program did while it was running provides no guarantee that the disclosed information actually reflects a computer system’s behavior.”).

99. See RAINIE & ANDERSON, *supra* note 72, at 19 (“Only the programmers are in a position to know for sure what the algorithm does, and even they might not be clear about what’s going on. In some cases there is no way to tell exactly why or how a decision by an algorithm is reached.” (quoting Doc Searls, Director, Project VRM, Berkman Klein Center For Internet & Society, Harvard University)); Chander, *supra* note 98, at 1040 (“[T]he algorithm may be too complicated for many others to understand, or even if it is understandable, too demanding, timewise, to comprehend fully.”); Kroll et al., *supra* note 98, at 638 (“The source code of computer systems is illegible to nonexperts. In fact, even experts often struggle to understand what software code will do, as inspecting source code is a very limited way of predicting how a computer program will behave.”); Guido Noto La Diega, *Against the Dehumanisation of Decision-Making: Algorithmic Decisions at the Crossroads of Intellectual Property, Data Protection, and Freedom of Information*, 9 J. INTELL. PROP. INFO. TECH. & ELEC. COM. L. 3, 23 (2018) (suggesting that “a technical document which includes the algorithm used and the mere explanation of the logic in mathematical terms will not in itself meet the legal requirement [for the right to explanation]” and that this requirement “should be interpreted as the disclosure of the algorithm with an explanation in non-technical terms of the rationale of the decision and criteria relied upon”).

unlikely to reveal the full extent of any problems that the automated fair use systems may encounter.¹⁰⁰ As Kartik Hosanagar and Vivian Jair observed:

[M]achine learning algorithms—and deep learning algorithms in particular—are usually built on just a few hundred lines of code. The algorithm[']s logic is mostly learned from training data and is rarely reflected in its source code. Which is to say, some of today’s best-performing algorithms are often the most opaque.¹⁰¹

Given these disclosure-related challenges, commentators have called for the development of explainable artificial intelligence to help document the algorithmic analysis and the training process and to enhance human understanding of the algorithmic operation.¹⁰² As Pauline Kim explained:

When a model is interpretable, debate may ensue over whether its use is justified, but it is at least possible to have a conversation about whether relying on the behaviors or attributes that drive the outcomes is normatively

100. As Yu and Spina Ali observed:

Deep learning machines can self-reprogram to the point that even their programmers are unable to understand the internal logic behind AI decisions. In this context, it is difficult to detect hidden biases and to ascertain whether they are caused by a fault in the computer algorithm or by flawed datasets.

Yu & Spina Ali, *supra* note 76, at 5; *see also* Chander, *supra* note 98, at 1040 (“[I]n the era of self-enhancing algorithms, the algorithm’s human designers may not fully understand their own creation: even Google engineers may no longer understand what some of their algorithms do.”). Likewise, Joshua Kroll and his collaborators explained:

Machine learning . . . is particularly ill-suited to source code analysis because it involves situations where the decisional rule itself emerges automatically from the specific data under analysis, sometimes in ways that no human can explain. In this case, source code alone teaches a reviewer very little, since the code only exposes the machine learning method used and not the data-driven decision rule.

Kroll et al., *supra* note 98, at 638 (footnote omitted).

101. Kartik Hosanagar & Vivian Jair, *We Need Transparency in Algorithms, but Too Much Can Backfire*, HARV. BUS. REV. (July 23, 2018), <https://hbr.org/2018/07/we-need-transparency-in-algorithms-but-too-much-can-backfire>; *see also* Daniel Gervais, *Exploring the Interfaces Between Big Data and Intellectual Property Law*, 10 J. INTEL. PROP. INFO. TECH. & ELEC. COM. L. 3, 5 (2019) (“[A]ny human contribution to the output of deep learning systems is ‘second degree.’”).

102. *See* ACM STATEMENT, *supra* note 95, Principle 4 (“Systems and institutions that use algorithmic decision-making are encouraged to produce explanations regarding both the procedures followed by the algorithm and the specific decisions that are made.”); INST. ELEC. & ELEC. ENGRS, *supra* note 80, at 68 (recommending software engineers to “document all of their systems and related data flows, their performance, limitations, and risks,” with emphases on “auditability, accessibility, meaningfulness, and readability”); Diakopoulous et al., *supra* note 96 (“Ensure that algorithmic decisions as well as any data driving those decisions can be explained to end-users and other stakeholders in non-technical terms.”). As Yu and Spina Ali recounted:

[A] team at Microsoft is trying to teach AI to show how it weighted every single variable in evaluating mortality risk factors. Similarly, a team at Rutgers University is working on a deep neural network that provides users with examples that demonstrates why it took a specific algorithmic decision. Another project at the University of Berkeley involves lashing two neural networks together, tasking one to describe the inner procedures running inside the other. Finally, an international team consisting, among the others, of researchers from Facebook, Berkeley and the University of Amsterdam has taught an image recognition software to show the evidence he relied upon to reach its decisions.

Yu & Spina Ali, *supra* note 76, at 7 (footnotes omitted).

acceptable. When a model is not interpretable, however, it is not even possible to have the conversation.¹⁰³

B. *Approximation*

The second pathway for automating the fair use standard is approximation. It differs from the translation pathway in that its primary goal is not to convert legal mandates or analytical approaches into computer code and algorithms, but to approximate those decisions that have already been made, or are to be made, by humans—whether in a courtroom, as part of law enforcement, or through ordinary day-to-day practice.¹⁰⁴ Because of the primary focus on end results and their correlation to human decisions, algorithm designers are free to come up with methods or strategies to facilitate legal automation, including those that judges, lawyers, law enforcement personnel, and other human decisionmakers have not traditionally used. The additional freedom in this pathway will also allow algorithm designers to take full advantage of the technological potential provided by deep learning, neural networks, and other advances in artificial intelligence.¹⁰⁵

For illustrative purposes, consider the different methods used to determine fair use in these two pathways. In the translation pathway, computers will be trained, most likely under the supervision of computer programmers,¹⁰⁶ to conduct fair use analysis based on the factors stipulated in Section 107 of the

103. Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & MARY L. REV. 857, 922–23 (2017).

104. See Burk & Cohen, *supra* note 25, at 57–58 (“Judicial determinations and negotiated minimum standards are not the only possible measures of current fair use practice; arguably, the more accurate measure of fair use is the daily behavior of ordinary users.”). See generally ASS’N OF INDEP. VIDEO & FILMMAKERS ET AL., DOCUMENTARY FILMMAKERS’ STATEMENT OF BEST PRACTICES IN FAIR USE (2005), http://archive.cmsimpact.org/sites/default/files/fair_use_final.pdf (stating the best practices in fair use for documentary filmmakers).

105. As a government report on artificial intelligence explained:

Deep learning uses structures loosely inspired by the human brain, consisting of a set of units (or “neurons”). Each unit combines a set of input values to produce an output value, which in turn is passed on to other neurons downstream. For example, in an image recognition application, a first layer of units might combine the raw data of the image to recognize simple patterns in the image; a second layer of units might combine the results of the first layer to recognize patterns-of-patterns; a third layer might combine the results of the second layer; and so on.

PREPARING FOR THE FUTURE, *supra* note 1, at 9. For discussions of deep learning, see generally ETHEM ALPAYDIN, MACHINE LEARNING: THE NEW AI 104–09 (2016); JOHN D. KELLEHER, DEEP LEARNING (2019); JOHN D. KELLEHER & BRENDAN TIERNEY, DATA SCIENCE 121–36 (2018); THIERRY POIBEAU, MACHINE TRANSLATION 181–95 (2017).

106. Machine learning generally can be separated into supervised and unsupervised learning, with the latter having no predefined output. See generally ALPAYDIN, *supra* note 105, at 38–42, 111–18 (discussing supervised and unsupervised learning); KELLEHER, *supra* note 105, at 26–30 (discussing supervised, unsupervised, and reinforcement learning). Supervision, in this case, will be to set parameters for the algorithmic operation or to add predefined outputs to constrain that operation. Although unsupervised learning has become increasingly attractive due to its unlimited potential, most artificial intelligence systems combine supervised and unsupervised learning techniques. See generally David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653 (2017) (providing an accessible overview of machine learning for lawyers).

Copyright Act.¹⁰⁷ The way these automated systems undertake individual factor-based analyses will likely mirror those taken by human decisionmakers. While these systems may end up generating different decisions, human experience largely informs the analytical processes that have been coded into the systems. In fact, past human decisions, including but not limited to those handed down by courts, will be used to train the automated systems to make future decisions.

By contrast, the approximation pathway allows algorithm designers—and, in the deep learning world, also artificial intelligence systems themselves—to freely determine the methods used to approximate decisions made by judges, lawyers, law enforcement personnel, and other human decisionmakers. As these methods and strategies are deployed, adjustments will be continuously made, utilizing new training data while relying on some or all algorithmic outputs as feedback data. As Professor Burk described:

One can imagine that a neural network or other machine learning system could detect these or other patterns in the data surrounding past cases, matching them to similar patterns in the data surrounding future fair use incidents, situations, and scenarios without formal programming definition of the fair use factors.¹⁰⁸

Professor Elkin-Koren noted the scenario in which “AI and machine learning would make it difficult for courts to check the rules embedded in the system, since these systems may not explicitly demonstrate the legal specifications of the four factors of fair use.”¹⁰⁹

In short, if automated systems are able to come up with decisions that have a strong correlation to human decisions—for example, with a ninety percent match (or whatever percentage society prefers)—that process may be deemed satisfactory even if it relies mostly on pattern recognition, as opposed to automated legal analyses based on the four statutorily stipulated fair use factors.¹¹⁰ After all, the primary focus of the approximation pathway is not on whether the automated systems have faithfully translated legal principles and techniques, but whether the decisions generated by those systems approximate human decisions.

One could certainly debate whether such approximation could provide an acceptable pathway for legal automation.¹¹¹ After all, thinking like a lawyer is

107. 17 U.S.C. § 107.

108. Burk, *Algorithmic Fair Use*, *supra* note 25, at 293.

109. Elkin-Koren, *supra* note 43, at 1099.

110. *Cf.* Volokh, *supra* note 8, at 1192 (“We should focus on the quality of the proposed AI judge’s product, not on the process that yields that product.”).

111. *See* Yu, *Fair Use*, *supra* note 18, at 347 (“While one could argue that a proper fair use analysis must be conducted the same way as how judges would, one cannot help but wonder whether society would find it acceptable to have automated fair use determinations that generate outcomes that have high correlations to the outcomes of judge-made decisions.”).

what law schools try to instill in future members of the legal profession.¹¹² Nevertheless, the benefit of this alternative pathway can be quite significant, especially considering the growing evidence that intelligent machines can perform quite well when left to their own devices.¹¹³ To be sure, human decisionmakers remain superior in making judgment calls,¹¹⁴ especially with respect to circumstances that have not arisen before.¹¹⁵ However, there is sufficient evidence to show that intelligent machines can compensate for these shortcomings by performing well on matters involving variables or hidden relationships that human decisionmakers often overlook. Because humans can make certain decisions better than machines, and vice versa, the best-case scenario is when the legal system can take full advantage of the superior performance of both types of decision-making.¹¹⁶

There are some significant drawbacks, however. Automated fair use systems could consider factors that are highly problematic in democratic society and that Congress and courts have treated as protected classes in the anti-discrimination context,¹¹⁷ such as the race, color, religion, or sex of the author or user. In their effort to approximate human decisions, these systems may also introduce new factors that the statute and case law have not mentioned or anticipated. While the creation of these new factors could spark helpful insights and research—on factors that are more predictive of fair use outcomes, perhaps—making decisions based on factors that courts do not use or anticipate is inherently problematic from a rule-of-law standpoint.¹¹⁸

112. The literature on how to think like a lawyer is vast. *See, e.g.*, BENJAMIN N. CARDOZO, *THE NATURE OF THE JUDICIAL PROCESS* (1921); WARD FARNSWORTH, *THE LEGAL ANALYST: A TOOLKIT FOR THINKING ABOUT THE LAW* (2007); OLIVER WENDELL HOLMES, JR., *THE PATH OF THE LAW* (1897); KARL N. LEWELLYN, *THE BRAMBLE BUSH: ON OUR LAW AND ITS STUDY* (1951); FREDERICK SCHAUER, *THINKING LIKE A LAWYER: A NEW INTRODUCTION TO LEGAL REASONING* (2012).

113. *See infra* text accompanying notes 126–131.

114. As Rebecca Crotof observed:

[T]he judgment we value in a common law process is a distinctively human skill. Human judges are sensitive to context, both to extenuating circumstances in individual cases and shifts in social norms over time, and can flexibly apply legal rules. While human contextualization may be incorporated during the design or training of an AI system, that is hardly the same as having human contextualization at the time the algorithmic rule is applied, especially as that application may occur in a temporally, geographically, and culturally different context. AI may be consistent, but it is “brittle”: “[I]t lacks] the flexibility humans have to step outside their instructions and apply ‘common sense’ to adapt to novel situations.”

Crotof, *supra* note 8, at 238 (footnotes omitted).

115. *See* AGRAWAL ET AL., *supra* note 17, at 59 (noting the weaknesses of machines in making predictions “when there is too little data” and concerning “events that are not captured by past experience”); Pasquale, *A Rule of Persons*, *supra* note 8, at 53 (“Many past efforts to rationalize and algorithmatize the law have failed, for good reason: there is no way to fairly extrapolate the thought processes of some body of past decisionmaking to all new scenarios.” (emphasis omitted)).

116. *See* sources cited *infra* note 150.

117. *Cf.* 42 U.S.C. § 2000e-2(a) (prohibiting workplace discrimination based on “race, color, religion, sex, or national origin”).

118. *See supra* text accompanying notes 165–167. In defense of automated fair use systems, the use of the words “shall include” in Section 107 of the Copyright Act indicates that the statute provides a

C. Self-Determination

The final pathway for automating the fair use standard is self-determination—that is, the automated systems will make autonomous decisions. While the starting point for the translation and approximation pathways is, respectively, to imitate methods or strategies used by humans or to approximate decisions they have already made, the self-determination pathway places emphasis on independent decision-making.

In this pathway, automated systems will make decisions that, in their views, will best promote creativity and serve the goals designated by computer programmers—in this case, the goals of copyright. They will make fair use determinations based on what they believe will fulfill those designated goals, as opposed to the goal of faithfully translating legal norms into computer code and algorithms or the goal of approximating human decisions.

Providing automated systems with wide autonomy will allow them to generate new fair use decisions that differ significantly from those that have already been, or are to be, handed down by courts. While such a pathway would be highly problematic from a stare decisis standpoint, especially in a common law jurisdiction like the United States, that pathway could help generate new solutions that may initially sound counterintuitive to human decisionmakers but that can in the end be proven to better promote creativity. If the goal of these automated systems is to improve the creative environment that copyright law supports, the latter can be as appealing as, if not more appealing than, the former.

Indeed, outside the area of fair use and intellectual property law, commentators have already documented how computers and artificial intelligence can generate seemingly counterintuitive decisions that are ultimately superior to human decisions.¹¹⁹ Even more complicated, human decisionmakers, due to their own cognitive barriers, may not always be able to fully appreciate the merits of these seemingly counterintuitive decisions. As Professors Casey and Niblett reminded us:

Algorithms will often identify counterintuitive connections that may appear erroneous to humans even when accurate. Humans should be careful in those cases not to undo the very value that was added by the algorithm’s ability to recognize these connections. This is especially true when the benefit of the algorithm was that it reduced human bias and behavioral errors.¹²⁰

non-exhaustive list of factors for courts to consider when making fair use determinations. *See supra* text accompanying note 24. Thus, when the automated systems introduce new factors that the statute and case law have not mentioned, these factors will not precipitate a direct conflict with the fair use provision.

119. *See* Millar & Kerr, *supra* note 16, at 120–22 (discussing the time when expert robots get better decisions than humans). *See generally id.* at 117–24 (discussing human–robot disagreement).

120. Casey & Niblett, *A Framework*, *supra* note 82, at 354; *see also* RAINIE & ANDERSON, *supra* note 72, at 40 (“People often confuse a biased algorithm for an algorithm that doesn’t confirm their biases. If

Compared with the translation and approximation pathways, the self-determination pathway will minimize these situations by ensuring that the automated systems will not immediately discard those machine-made decisions that do not correspond well to preexisting human decisions. Nevertheless, because this pathway may generate decisions that differ significantly from those preexisting decisions, a society that chooses the self-determination pathway should put in place mechanisms to address potential conflicts between human and machine-made decisions.¹²¹

D. Summary

Even though this Part has focused on three distinct pathways for legal automation, it is important to keep in mind that hybrid routes can be developed to incorporate more than one pathway. Indeed, the choice over the best mix of pathways will lead algorithm designers to ask some key questions concerning how best to automate legal standards and how to address the law-machine interface. The next Part will discuss these design questions in greater detail.

As time passes, and as artificial intelligence technologies continue to improve, new pathways may also emerge while some existing ones may become obsolete. Should we reach the technological state at which machine-made decisions are always preferable to human decisions—a scenario that would admittedly be very far away¹²²—the starting point for making legal decisions may be intelligent machines, not human decisionmakers. If so, the translation and approximation pathways would seem somewhat misguided, as they privilege human decisions over machine-made decisions. Those two pathways would also become increasingly impractical. After all, machines, not humans, would make the majority of decisions, and there might not be enough human decisions for machines to translate from or approximate.

III. LAW-MACHINE INTERFACE

The previous Part has identified three distinct pathways for legal automation that can help enlist algorithms and artificial intelligence to

Facebook shows more liberal stories than conservative, that doesn't mean something is wrong. It could be a reflection of their user base, or of their media sources, or just random chance." (quoting an anonymous principal consultant of a consulting firm)); Harry Surden & Mary-Anne Williams, *Technological Opacity, Predictability, and Self-Driving Cars*, 38 CARDOZO L. REV. 121, 158 (2016) ("[I]t is not uncommon for pilots in the cockpit to be surprised or confused by an automated activity undertaken by an autopilot system."). See generally Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 FORDHAM L. REV. 1085 (2018) (documenting the limitations of intuition while noting the need to address inscrutability).

121. See discussion *infra* Part III.B, III.C.

122. See, e.g., Wu, *supra* note 8, at 2004 ("[F]or the foreseeable future, software systems that aim to replace systems of social ordering will succeed best as human-machine hybrids, mixing scale and efficacy with human adjudication for hard cases.").

modernize the legal system. This Part turns to a key issue that most commentators have overlooked: the law–machine interface. To illustrate the different questions on algorithmic design that will emerge in relation to this interface, this Part focuses on three distinct issues: (1) the allocation of decision-making power; (2) the hierarchy of decisions; and (3) the legal effects of machine-made decisions. The more algorithm designers think through questions involving these issues, the more success they will likely have in charting an effective path toward legal automation.

A. Allocation of Decision-Making Power

When machine-made decisions are inferior to human decisions, it is logical that technology will be used only, or mostly, to assist humans in making decisions. By default, decision-making power resides in humans. However, as artificial intelligence technologies continue to improve and as intelligent machines become capable of making better decisions—at least in select areas¹²³—questions will arise over the allocation of decision-making power.¹²⁴ Should machines at least make some decisions?¹²⁵ If so, what are those decisions? Should those machine-made decisions receive deference in the legal system?

With growing evidence on the machines’ ability to outperform humans in select areas, answering these questions has become increasingly challenging. For instance, researchers have documented the advantage of using learning algorithms to diagnose cancer and to perform other tasks in the health area.¹²⁶

123. See Millar & Kerr, *supra* note 16, at 117 (“Once there are expert robots, it will be easier to argue in some instances that they *ought* to be used to their full potential, because the evidence will suggest that in those instances they will, on average, deliver better results than human experts.”).

124. For example, Tim Wu asked: “Just when and why are decisions brought to human attention, and who decides when a human should decide?” Wu, *supra* note 8, at 2027. Viktor Mayer-Schönberger and Thomas Ramge asked a similar question: “Which decisions should we reserve for ourselves and which should we delegate?” VIKTOR MAYER-SCHÖNBERGER & THOMAS RAMGE, *REINVENTING CAPITALISM IN THE AGE OF BIG DATA* 219 (2018).

125. As Mayer-Schönberger and Ramge observed:

If data-driven adaptive systems will offer us better answers to questions such as which school we should send our kids to or which hospital an ambulance should take us to in case of an emergency, then should we delegate that decision to the machines or retain it as the exclusive province of human responsibility? What are we aiming for in decisions, anyway—getting the correct answer or the one that makes us happy (after all, we, not the machines, must live with the consequences)? Until now we rarely faced such choices, but in the future we routinely will. Developing a good, solid sense of how to choose is a core competency we’ll have to develop and maintain.

This ability to choose what to choose is fundamentally empowering to humans. It preserves our chance to contribute to the fate of the universe and may ensure us an enduring seat at the table of evolution.

MAYER-SCHÖNBERGER & RAMGE, *supra* note 124, at 219–20.

126. See ERIC J. TOPOL, *DEEP MEDICINE: HOW ARTIFICIAL INTELLIGENCE CAN MAKE HEALTHCARE HUMAN AGAIN* 117–18 (2019) (discussing the impressive progress in algorithmic image processing); Jonathan Guo & Li Bin, *The Application of Medical Artificial Intelligence Technology in Rural Areas of Developing Countries*, 2 *HEALTH EQUITY* 174, 175 (2018) (noting research showing that systems using deep

Commentators have also noted that algorithms “are better and faster than humans at detecting credit card fraud,”¹²⁷ not to mention that “[m]achines can pool their resources in ways that humans cannot.”¹²⁸ In addition, the performance of intelligent machines will not be affected by emotion, exhaustion, stress, or other cognitive barriers.¹²⁹ These machines “can [also] be tested and [therefore] improved.”¹³⁰ Should errors be found and corrected, the machines “are unlikely to make the same mistake[s] again.”¹³¹

Given such superior performance, one cannot help but wonder whether machines, as opposed to humans, should make more decisions. In several narrow areas that require instantaneous responses, such as those involving the application of emergency brakes in automobiles, we have already given machines significant power to make those decisions.¹³²

For illustrative purposes, consider the automated analysis of the four statutorily stipulated fair use factors. While an automated system may find it challenging to analyze the first factor concerning “the purpose and character of

convolutional neural networks are “able to classify skin cancer at a comparable level to dermatologists” and “could improve the speed, accuracy, and consistency of diagnosis [of breast cancer metastasis in lymph nodes], as well as reduce the false negative rate to a quarter of the rate experienced by human pathologists”).

127. *Digital Decisions*, *supra* note 76.

128. RAY KURZWEIL, *THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY* 261 (2005) [hereinafter KURZWEIL, *THE SINGULARITY IS NEAR*].

129. See Karni Chagal-Feferkorn, *The Reasonable Algorithm*, 2018 U. ILL. J.L. TECH. & POL’Y 111, 144 (“Unlike humans, algorithms do not have self-interests affecting their judgement, they do not omit any of the decision-making stages or base their decisions on heuristics or biases, and they are not subject to human physical or emotional limitations such as exhaustion, stress or emotionality.” (footnotes omitted)); Crotoof, *supra* note 8, at 236 (noting that a “judge’s sensitivity to context and penchant for leniency may vary dramatically with whether they are hungry, tired, bored, overworked, overwhelmed, or otherwise distracted”); Daryl Lim, *AI & IP: Innovation & Creativity in an Age of Accelerated Change*, 52 AKRON L. REV. 813, 834 (2018) (“AI does not suffer from perceptual limitations the way that humans do.”); Ozkan Eren & Naci Mocan, *Emotional Judges and Unlucky Juveniles* (Nat’l Bureau of Econ. Rsch., Working Paper No. 22,611, 2016), <https://www.nber.org/papers/w22611.pdf> (documenting the surprising impact of unexpected outcomes of football games on the type and length of sentences handed down by juvenile court judges); Kurt Kleiner, *Lunchtime Leniency: Judges’ Rulings Are Harsher When They Are Hungrier*, SCI. AM. (Sept. 1, 2011), <https://www.scientificamerican.com/article/lunchtime-leniency> (“Judges granted 65 percent of requests they heard at the beginning of the day’s session and almost none at the end. Right after a snack break, approvals jumped back to 65 percent again.” (citing a study at Ben Gurion University in Israel and Columbia University examining more than 1,000 decisions by eight Israeli judges who ruled on convicts’ parole requests)).

130. MCAFEE & BRYNJOLFSSON, *supra* note 69, at 53.

131. *Id.* As the authors observed, “[I]t is a lot harder to get humans to acknowledge their biases (how many avowed racists or sexists do you know?), let alone do the hard work required to overcome them.” *Id.*

132. See AGRAWAL ET AL., *supra* note 17, at 112 (“Carmakers in the United States have reached an agreement with the Department of Transportation to make automatic emergency braking standard on vehicles by 2022.”); Millar & Kerr, *supra* note 16, at 118 (“[C]ases that are time-sensitive—critical emergency room admissions, perhaps, or cases where [Google driverless cars] need to make split-second decisions about how best to navigate rapidly evolving traffic situations—might afford human experts the time to disagree with the robot, but little or no time to evaluate the underlying rationales to come to anything resembling a meaningful conclusion about the sources of disagreement.”).

the use”¹³³ of the copyrighted work,¹³⁴ it may find the analysis of other factors easier. A case in point is the analysis of the third factor, which focuses on “the amount and substantiality of the portion used in relation to the copyrighted work as a whole.”¹³⁵ When analyzing this factor, courts usually engage in both quantitative and qualitative analyses.¹³⁶ For computers, quantitative analyses are the easiest.¹³⁷ In fact, any judge wanting to make an efficient and effective comparison will deploy computers to undertake some of the comparative tasks, such as counting the number of words in the original work and the potentially infringing work.¹³⁸

By contrast, qualitative analyses seem to be much more challenging. After all, how can an automated system know which part of the copyrighted work is highly important, especially considering that computers and robots are notorious for their lack of emotion and empathy?¹³⁹ Indeed, determining what

133. 17 U.S.C. § 107.

134. See Burk, *Algorithmic Fair Use*, *supra* note 25, at 292 (“[C]oncepts like ‘educational use’ or ‘news reporting’ might be unexpectedly tricky to reduce to computable code. But one can, for example, imagine programming a system to determine, perhaps on the basis of geolocational data and scraped calendaring or advertising data, whether a nondramatic musical work is being performed at an agricultural fair.”); Felten, *supra* note 46, at 58 (identifying the “[l]ack of knowledge about the circumstances” of the use as one of the two key reasons why fair use cannot be built into digital rights management systems (emphasis omitted)). However, Professor Elkin-Koren disagreed:

[One] concern [regarding the limitations of algorithmic fair use] is that algorithms that analyze fair use will fail to process information that is external to the content itself. For instance, determining the nature of use may require external information and additional analysis of facts. Yet, algorithms could be programmed to extract and analyze data from external sources. For instance, educational use might be determined based on tagging the nature of the user. A program could detect the type of user (e.g., educational institution, governmental agency) based on the domain name (e.g., .edu, .gov) or by checking registration in external databases. Another indication for the nature of use could be the type of tagging selected by the party that uploads the work (educational, commercial, personal/private use). The commercial nature of use might actually be determined by the presence of advertisements, or other means of monetizing the content. External information might also be used to determine “the effect of the use upon the potential market” for the copyrighted work, using the commercial nature of use as a proxy.

Elkin-Koren, *supra* note 43, at 1095–96. Whether the automated system can extract and analyze data from external sources, as Professor Elkin-Koren proposed, will depend largely on whether an enabling environment exists to allow for such extraction and analysis. See Yu, *Fair Use*, *supra* note 18, at 351–63 (underscoring the need to build this enabling environment).

135. 17 U.S.C. § 107.

136. See *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 587 (1994) (“[The third] factor calls for thought not only about the quantity of the materials used, but about their quality and importance, too.”); *Harper & Row, Publishers, Inc. v. Nation Enters.*, 471 U.S. 539, 583 n.6 (1985) (Brennan, J., dissenting) (“The inquiry into the substantiality of appropriation has a quantitative and a qualitative aspect.”).

137. See Elkin-Koren, *supra* note 43, at 1096 (“Some fair use considerations might be relatively easy to automate, such as the amount copied from the original work. For instance, a program could give a higher fair use score based on similarity of less than 10 percent.”).

138. See Yu, *Fair Use*, *supra* note 18, at 344 (“[A]ny judge seeking to undertake a quick quantitative analysis will likely rely on computer assistance to count words or compare sizes.”).

139. See LEE, *supra* note 4, at 142 (“Taking the next step to emotionally intelligent robots may require self-awareness, humor, love, empathy, and appreciation for beauty. These are the key hurdles that separate what AI does today—spotting correlations in data and making predictions—and artificial general intelligence.”); MCAFEE & BRYNJOLFSSON, *supra* note 69, at 123 (“[T]he ability to work effectively with

courts have referred to as the “heart” of the work¹⁴⁰ will likely require the professional judgment of human decisionmakers.¹⁴¹ Nevertheless, Amazon now has a large trove of data concerning which pages or sentences of a book Kindle users have highlighted.¹⁴² Netflix also has substantial, and at times shocking, data about which part of a movie or a TV program its subscribers have paused or viewed repeatedly.¹⁴³ In fact, with the deployment of big data analysis and the utilization of external market data, the automated system may be able to generate some useful predictions on which part of the copyrighted work will likely be popular or commercially successful. Even though these indicators alone may not show what courts would consider as the heart of the copyrighted work, the increased availability of these indicators does suggest the machines’ growing ability to make automated fair use determinations.

Similar to the third-factor analysis, computers can also analyze quite well the fourth factor, which concerns “the effect of the use upon the potential market for or value of the copyrighted work.”¹⁴⁴ On its face, analyzing the actual or potential market of a copyrighted work will require professional expertise. As Professor Felten observed more than a decade ago, “[T]he fourth factor in the [fair use] test . . . requires reasoning about the economics of a particular market, a task even well-trained humans find difficult.”¹⁴⁵ In reality, computers and artificial intelligence have already been actively deployed to provide predictive analyses in many areas that are far more complex, challenging, and volatile than predicting the market of a copyrighted work.¹⁴⁶ In the financial industry, for example, it is increasingly common to find computers making

people’s emotional states and social drives will remain a deeply human skill for some time to come.”); TOPOL, *supra* note 126, at 290 (“[H]uman empathy is not something machines can truly simulate, despite ongoing efforts to design sociable robots or apps that promote empathy.”); Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231, 1269–71 (1992) (discussing the lack of capacity in artificial intelligence for feelings).

140. See, e.g., *Campbell*, 510 U.S. at 588–89 (discussing the use of “the ‘heart’ of the original” in the parody context); *Harper & Row*, 471 U.S. at 600 (analyzing whether the defendant magazine “had taken ‘the heart of the book’”).

141. See Burk, *Algorithmic Fair Use*, *supra* note 25, at 292 (noting the difficulty in programming an automated fair use system “to determine . . . whether an excerpt from the work is so significant as to constitute the ‘heart’ of an author’s creation”).

142. See *Viewing Popular Highlights on Kindles*, EBOOK READER (Feb. 15, 2018), <https://blog.the-ebook-reader.com/2018/02/15/viewing-popular-highlights-on-kindles> (“Popular Highlights show the most highlighted passages that readers have added to Kindle books. . . . Amazon also displays how many times each passage has been highlighted.”).

143. See Kal Raustiala & Christopher Jon Sprigman, *The Second Digital Disruption: Streaming and the Dawn of Data-Driven Creativity*, 94 N.Y.U. L. REV. 1555, 1587 (2019) (“Some parameters that Netflix tracks include, but are likely not limited to, pause/rewind/fast-forward behavior; day of the week; date of viewing; time of viewing; zip code; preferred devices; completion rate; user ratings; user search behavior; and browsing and scrolling behavior.”).

144. 17 U.S.C. § 107.

145. Felten, *supra* note 46, at 58.

146. See Elkin-Koren, *supra* note 43, at 1097 (“AI has already been applied for decision-making processes in contexts that are far more complex than fair use, involving critical issues of life and death, health, financial risks, and national security.”).

predictions about stock values or prices.¹⁴⁷ There is also a fast-growing literature on the use of artificial intelligence in finance.¹⁴⁸

In sum, society will continue to rely on humans to make certain decisions, especially those involving judgment calls or those lacking in historical data. Meanwhile, machines can be utilized to make other decisions. Even if those machine-made decisions are not better than human decisions, the machines' ability to provide decisions in real time, or close to real time, will make the former highly appealing.¹⁴⁹ For creative projects that do not involve substantial investments, many users will likely find instantaneous fair use determinations more useful than time-delayed decisions rendered by professional experts, as long as there is no significant variation in quality.

Because humans and machines can make better decisions in different areas, commentators have started to highlight the importance of enabling two types of decisions to complement each other.¹⁵⁰ For instance, Lee Kai-fu provided “a blueprint for human coexistence with AI.”¹⁵¹ Frank Levy and Richard Murnane discussed the importance of a “new division of labor” that aims to maximize the comparative advantage of both humans and machines.¹⁵² Mary Gray and

147. See Chagal-Feferkorn, *supra* note 129, at 116 (“In finance, algorithms are used for assessing credit risks and mortgage risks, pricing complex insurance products, stocks ranking, or in general, creating financial forecasts.”).

148. The literature emergent in this area is vast and fast-growing. See generally Tom C.W. Lin, *Artificial Intelligence, Finance, and the Law*, 88 *FORDHAM L. REV.* 531 (2019) (discussing the risks and limitations of financial artificial intelligence); William Magnuson, *Artificial Financial Intelligence*, 10 *HARV. BUS. L. REV.* 337 (2020) (discussing the dangers and real-world limitations of deploying artificial intelligence in finance); Dirk A. Zetzsche et al., *Artificial Intelligence in Finance: Putting the Human in the Loop* (Univ. of Hong Kong Fac. of L. Working Paper, Paper No. 2020/006, 2020) <https://ssrn.com/abstract=3531711> (discussing the increasing role of artificial intelligence in finance, with a focus on human responsibilities).

149. See Yu, *Fair Use*, *supra* note 18, at 346 (“[Automated fair use systems] will be able to draw conclusions more quickly than humans, and will thereby facilitate real-time market analysis that will be both costly and time-consuming when conducted manually.”).

150. As Ajay Agrawal, Joshua Gans, and Avi Goldfarb observed in the context of cancer diagnostics: The human and the machine are good at different aspects of prediction. The human pathologist was usually right when saying there was cancer. It was unusual to have a situation in which the human said there was cancer but was mistaken. In contrast, the AI was much more accurate when saying the cancer wasn't there. The human and the machine made different types of mistakes. By recognizing these different abilities, combining human and machine prediction overcame these weaknesses, so their combination dramatically reduced the error rate.

AGRAWAL ET AL., *supra* note 17, at 65; see also Re & Solow-Niederman, *supra* note 8, at 282 (“[H]uman and AI judges might collaborate by operating in tandem at specified stages of the judicial process, either by functioning with a human in-the-loop or by preserving an extra measure of human oversight and involvement at particular points.”); Paul Scharre, *Centaur Warfighting: The False Choice of Humans vs. Automation*, 30 *TEMP. INT'L & COMPAR. L.J.* 151, 151 (2016) (“[I]n many situations, human-machine teaming in engagement decisions will not only be possible but preferable. Hybrid human-machine cognitive architectures will be able to leverage the precision and reliability of automation without sacrificing the robustness and flexibility of human intelligence.”); Wu, *supra* note 8, at 2005 (“[W]hen it comes to systems that replace the law, designers should be thinking harder about how best to combine the strengths of humans and machines, by understanding the human advantages of providing a sense of procedural fairness, explainability, and the deciding of hard cases.”).

151. See LEE, *supra* note 4, at 197–225.

152. LEVY & MURNANE, *supra* note 17.

Siddharth Suri documented the large pool of humans performing “ghost work” that is indispensable to advances in the field of artificial intelligence.¹⁵³ Sarah Roberts provided an important ethnographic study of human commercial content moderators, who work behind the scenes to screen and remove content, enforce policies on online platforms, and improve the outcomes of automated moderation.¹⁵⁴ Thus, should humans and machines be making decisions at the same time, it will be highly important to decide how to allocate decision-making power between humans and machines. For those with economic acumen, it will also be fruitful to find ways to maximize the optimality of such allocation.

B. *Hierarchy of Decisions*

Once we have decided how to allocate decision-making power, the next key design question concerns the hierarchy of decisions—or the establishment of a set of decisional rules.¹⁵⁵ For the foreseeable future, human decisions will trump machine-made decisions in most, if not all, cases. However, as society becomes more accustomed to artificial intelligence and more willing to trust machine-made decisions, the latter will receive more deference—either in select areas or be given more weight in the overall decisions. As a result, the hierarchy of decisions may begin to shift away from a hegemony of human decisions.

Consider, for instance, the context of automated copyright enforcement, which provides one of the most widely used examples of automated legal systems.¹⁵⁶ While machines have been used to identify potential infringing

153. *See generally* MARY L. GRAY & SIDDHARTH SURI, GHOST WORK: HOW TO STOP SILICON VALLEY FROM BUILDING A NEW GLOBAL UNDERCLASS (2019). As they explained:

Beyond some basic decisions, today’s artificial intelligence can’t function without humans in the loop. Whether it’s delivering a relevant newsfeed or carrying out a complicated texted-in pizza order, when the artificial intelligence . . . trips up or can’t finish the job, thousands of businesses call on people to quietly complete the project. This new digital assembly line aggregates the collective input of distributed workers, ships pieces of projects rather than products, and operates across a host of economic sectors at all times of the day and night.

Id. at ix–x. Falling within “ghost work” are such tedious tasks as content classification, image tagging, photo comparison, video screening, and data cleaning. *See id.* at x–xxiii. The book further discussed the need for human workers to develop datasets that are used for training artificial intelligence and how the new advances, in turn, have generated new cycles that require even more human workers to complete intervening tasks. They described these cycles as “the paradox of automation’s last mile”: “Humans trained an AI, only to have the AI ultimately take over the task entirely. Researchers could then open up even harder problems. . . . These problems needed yet more training data, generating another wave of ghost work.” *Id.* at 8.

154. *See generally* SARAH T. ROBERTS, BEHIND THE SCREEN: CONTENT MODERATION IN THE SHADOWS OF SOCIAL MEDIA (2019). As she observed, “Issues of scale aside, the complex process of sorting user-uploaded material into either the acceptable or the rejected pile is far beyond the capabilities of software or algorithms alone.” *Id.* at 34.

155. This hierarchy of decisions immediately brings to mind Isaac Asimov’s Second Law of Robotics: “A robot must obey the orders given it by human beings except where such orders would conflict with the First Law [which states that a robot may not injure a human being or, through inaction, allow a human being to come to harm].” ISAAC ASIMOV, *Runaround*, in 1, ROBOT 25, 37 (Del Rey, reprint ed. 2008).

156. *See* sources cited *supra* note 39.

materials, human oversight has been built into the systems to ensure verification before takedown requests are sent to online service providers or platforms.¹⁵⁷ Indeed, when incorrect requests have been made, the copyright holder or its supportive industry group often explains away the mistake by showing how the human involved has failed to properly verify the alleged infringement.¹⁵⁸

In recent years, however, we have seen the growing use of robo notices, automatic takedown notices that are being sent out by computers to online service providers or platforms with no or insufficient human oversight.¹⁵⁹ Part of the reason for the popularity of these robo notices is their ability to respond to the unmanageable volume of copyrighted works that are now being disseminated and the exceedingly large amount of potential infringement that is being found on the Internet.¹⁶⁰ Another key reason is that economics favor the use of such automated notices, especially when there is no penalty for sending out incorrect notices.¹⁶¹ Indeed, commentators have lamented the growing impact of a large volume of robo notices that has now been sent to online service providers and platforms without human oversight.¹⁶² In short,

157. *Cf.* *Lenz v. Universal Music Corp.*, 572 F. Supp. 2d 1150, 1155 (N.D. Cal. 2008) (“The DMCA already requires copyright owners to make an initial review of the potentially infringing material prior to sending a takedown notice; indeed, it would be impossible to meet any of the requirements of Section 512(c) without doing so.”).

158. *See, e.g.*, Declan McCullagh, *RIAA Apologizes for Threatening Letter*, ZDNET (May 13, 2003, 11:42 GMT), <https://www.zdnet.com/article/riaa-apologizes-for-threatening-letter> (reporting the claim of the Recording Industry Association of America that the failure of a temporary employee to follow its established protocol was the reason behind a wrongful takedown notice sent to Penn State University that had almost caused the departmental server to shut down during the final examination period).

159. *See generally* Joe Karaganis & Jennifer Urban, *The Rise of the Robo Notice*, COMM. ACM, Sept. 2015, at 28 (expressing concern about the growing use of robo notices to remove potentially infringing copyrighted materials).

160. *Cf. id.* at 28 (noting “the adoption of automated notice-sending systems by rights holder groups responding to sophisticated infringing sites”); Perel & Elkin-Koren, *Black Box Tinkering*, *supra* note 39, at 190–91 (“[P]rivate, online intermediaries . . . often use robots to handle the immense traffic of online content.”); Yu, *Fair Use*, *supra* note 18, at 347–48 (“With the creation and dissemination of hundreds of exabytes of data and digital content every day, it is almost impossible for technology platforms to not rely on algorithms to determine whether a specific use of a copyrighted work has complied with copyright law.” (footnote omitted)).

161. Although Section 512(f) of the Copyright Act penalizes those who “knowingly materially misrepresent[]” information, 17 U.S.C. § 512(f), “copyright’s ambiguity assures that many statements of infringement can be made in good faith, even though a court may find that no infringement actually exists.” Alfred C. Yen, *Internet Service Provider Liability for Subscriber Copyright Infringement, Enterprise Liability, and the First Amendment*, 88 GEO. L.J. 1833, 1888 n.278 (2000); *see also* Karaganis & Urban, *supra* note 159, at 30 (“Stronger liability for reckless or malicious notice use might be a good step in curbing the worst notice practices, which can include deceptive or predatory behavior. But such changes are currently a dead letter in U.S. copyright politics.”).

162. *See* Perel & Elkin-Koren, *Black Box Tinkering*, *supra* note 39, at 204 (“[R]ecent studies prove that prominent [online service providers], facing a flood of robo-takedown notices sent automatically by right-holders, substitute human review of the vast majority of these notices with their own privately designed automated systems.”); Sag, *Internet Safe Harbors*, *supra* note 40, at 543 (“[I]n spite of the DMCA’s requirement that takedown notices attest to the complaining party’s ‘good faith belief’ in infringement, massive volumes of such notices are clearly sent, and often acted upon, without meaningful human review.” (footnote omitted)).

the change in technological environment and social preferences has caused society to give machine-made decisions more deference than they once had.

Similar changes can be found in the fair use context. Although fair use decisions could be made with built-in human oversight, the large volumes of online content that are being evaluated for fair use purposes will likely require the development of automated systems.¹⁶³ If so, humans will have to provide oversight after the fact.

One possibility for providing such oversight *ex post* is to allow machine-made decisions to be challenged in a court of law.¹⁶⁴ Upon such a challenge, a judge will be able to intervene should the automated system reach a wrong or undesirable decision. The allowance for judicial intervention precipitates the need to think more deeply about the hierarchy of decisions—Should judges always trump machines? From a rule-of-law or constitutional standpoint, there are considerable benefits to reserving final decisions to human judges.¹⁶⁵ As Tim Wu reminded us, a key advantage of retaining the use of human courts is procedural fairness.¹⁶⁶ As he observed, “There are . . . advantages to adjudication as a form of social ordering that are difficult to replicate by any known means.”¹⁶⁷

One caveat that is worth noting in this area concerns the challenge of deciding when to undertake human intervention. Just because the automated systems have made decisions that differ significantly from what human judges would have rendered does not mean that those machine-made decisions are wrong or undesirable.¹⁶⁸ When these decisions are challenged before courts,

163. See AGRAWAL ET AL., *supra* note 17, at 67 (“One major benefit of prediction machines is that they can scale in a way that humans cannot.”); TARLETON GILLESPIE, *CUSTODIANS OF THE INTERNET: PLATFORMS, CONTENT MODERATION, AND THE HIDDEN DECISIONS THAT SHAPE SOCIAL MEDIA* 97 (2018) (“Artificial intelligence techniques offer . . . to solve the problem of scale.”); Wu, *supra* note 8, at 2002 (“Compared with the legal system, software has enormous advantages of scale and efficacy of enforcement. It might tirelessly handle billions if not trillions of decisions in the time it takes a human court to decide a single case.”); Yu, *Fair Use*, *supra* note 18, at 347–49 (discussing the scalability of automated fair use systems).

164. See sources cited *supra* note 84.

165. See Michaels, *supra* note 8 (discussing the negative impact of automated adjudication on legal change, separation of powers, and the rule of law); Re & Solow-Niederman, *supra* note 8, at 262–78 (noting the concerns that artificial intelligence-based adjudication will make the legal system more incomprehensible, data-based, alienating, and disillusioning).

166. See Wu, *supra* note 8, at 2002 (“One set of advantages [of human courts] . . . is related to procedural fairness.”); see also Margot E. Kaminski, *Binary Governance: Lessons from the GDPR’s Approach to Algorithmic Accountability*, 92 S. CAL. L. REV. 1529, 1554–57 (2019) (recapitulating the literature on algorithmic due process); Re & Solow-Niederman, *supra* note 8, at 284 (“The idea of mechanized verdicts, especially criminal verdicts, . . . seems to cut at the heart of democratic self-government, as well as due process.”); Olivier Sylvain, *Recovering Tech’s Humanity*, 119 COLUM. L. REV. F. 252, 261 (2019) (“Human review is essential today because it confers a degree of legitimacy on the platforms’ moderation choices.”). See generally Tom R. Tyler, *Procedural Justice and the Courts*, 44 CT. REV. 26 (2007) (discussing the importance of procedural justice).

167. Wu, *supra* note 8, at 2002.

168. See *supra* text accompanying notes 119–120. The converse is also true. Just because the automated systems have made decisions that coincide with what human judges would have rendered does not mean that those decisions are necessarily more correct or desirable. Cf. Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 33 (2014) (“Scoring systems have a

judges will have to use their own professional judgment to determine whether to retain those seemingly incorrect or counterintuitive machine-made decisions. Indeed, making such determinations will remain a key exercise of judicial discretion.

Finally, should society decide to let machines make at least some autonomous decisions, we can still have a hierarchy of decisions favoring humans—for instance, by providing an opportunity to have an override.¹⁶⁹ A good example in this area concerns those algorithms that have been deployed in cars to facilitate automatic lane correction.¹⁷⁰ When a car veers into another lane, those built-in algorithms will quickly help the driver steer the car back to its original lane. Should the driver disagree with the computer-made decision, the human decisionmaker can hold on to the steering wheel or turn it in the opposite direction to initiate an override. By providing this override, the algorithms involved preserve a hierarchy of decisions that favors human decisions. Such an arrangement contrasts significantly with the arrangement for the automatic application of emergency brakes, in which machine-made decisions will trump human decisions.

C. *Legal Effects of Machine-Made Decisions*

Once we have figured out the hierarchy of decisions, there remains the final design question concerning what legal effects machine-made decisions will have.¹⁷¹ For instance, in an environment in which humans can intervene by making decisions that trump machine-made decisions, there will always be questions concerning what legal effects machine-made decisions will have should no human decisionmaker intervene.¹⁷² In an environment in which both

powerful allure—their simplicity gives the illusion of precision and reliability. But predictive algorithms can be anything but accurate and fair.”); Ric Simmons, *Big Data, Machine Judges, and the Legitimacy of the Criminal Justice System*, 52 U.C. DAVIS L. REV. 1067, 1075 (2018) (“[P]redictive algorithms . . . create an illusory ‘technocratic framing’ of who is dangerous and who deserves greater punishment, even though the algorithms’ conclusions are based on the same flawed data.”); Megan Stevenson, *Assessing Risk Assessment in Action*, 103 MINN. L. REV. 304, 375 (2018) (“Risk assessment tools wear the clothes of an evidence-based practice—they are developed with the use of large data sets and sophisticated techniques and endorsed by social scientists running policy simulations—but risk assessments should not be considered evidence-based until they have shown to be effective.”).

169. See Re & Solow-Niederman, *supra* note 8, at 282 (“[H]uman and AI judges might collaborate by operating in tandem at specified stages of the judicial process, either by functioning with a human in-the-loop or by preserving an extra measure of human oversight and involvement at particular points.”).

170. See PREPARING FOR THE FUTURE, *supra* note 1, at 18 (noting the “advanced cruise controls that keep a car in its lane”).

171. Professor Hildebrandt defined legal effect as follows: “Legal effect denotes the consequences that legal norms attach to specific actions or states; legal effect changes the legal status of a person or other entity and attributes the ensuing rights and obligations to legal subjects.” HILDEBRANDT, *supra* note 10, at 168.

172. Cf. Mireille Hildebrandt, *A Vision of Ambient Law*, in REGULATING TECHNOLOGIES: LEGAL FUTURES, REGULATORY FRAMES AND TECHNOLOGICAL FIXES 175, 177–80 (Roger Brownsword & Karen Yeung eds., 2008) [hereinafter REGULATING TECHNOLOGIES] (discussing the distinction between technological and legal normativity); Bert-Jaap Koops, *Criteria for Normative Technology: The Acceptability of ‘Code*

humans and machines will make different legal decisions at the same time, or in which human decisions do not always trump machine-made decisions, this type of question will be raised even more frequently and will take on even greater significance.

Consider once again the case study of fair use automation. Should a user receive a machine-made fair use determination, would that determination have any legal effect in the sense that it will protect the user from future legal liability? This question will be important even if human judges can always intervene by overturning machine-made decisions. After all, if the machine-made decision has the force of law, and the user has in fact relied on that decision to disseminate the allegedly infringing content, that user will not infringe on the protected work until the court overturns the decision. If the machine-made determination is recognized by multiple platforms, including platforms that are available overseas, giving legal effects to machine-made decisions can help facilitate content distribution across these platforms, both domestically and globally.

By contrast, if the machine-made determination has no legal effect, the infringement can be traced back to the time before the court makes its fair use decision, even though a judge could reduce the damage award based on evidence of good-faith reliance on the machine-made determination. To be sure, users are unlikely to seek machine-made fair use determinations if they know in advance that such determinations will have no legal effects. However, because fair use determinations are often made at gateways when platform users upload content for dissemination, these users will still have strong incentives to seek those determinations or will have no choice but to go through with such determinations. For example, YouTube users seek machine-made determinations not because they rely on the legal effects of those determinations, but because such determinations are part of the content uploading process.¹⁷³

Finally, in determining the legal effects of machine-made decisions, one could take a middle approach by giving those decisions some deference while retaining some legal liability.¹⁷⁴ For instance, with respect to copyright infringement, society could introduce laws to allow machine-made fair use

as Law' in *Light of Democratic and Constitutional Values*, in REGULATING TECHNOLOGIES, *supra*, at 157, 161 ("The way in which a legal norm is translated and inscribed in technology is a separate activity that should be assessed in its own right, because 'law in the books' is not and cannot be exactly the same as 'law in technology.'"); Burk, *Algorithmic Fair Use*, *supra* note 25, at 297 ("Patterns detected by a machine evaluating fair use-related data should not be confused with a legal institutional determination of fair use."); Elkin-Koren, *supra* note 43, at 1099 ("AI systems do not decide fair use, but simply generate a score that reflects the probability of fair use.").

173. See *supra* text accompanying note 40 (discussing YouTube's Content ID system).

174. Cf. *Losing Humanity: The Case Against Killer Robots*, HUM. RTS. WATCH (Nov. 19, 2012), <https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots> (advancing in the context of automated weapons the trichotomy of "Human-in-the-Loop," "Human-on-the-Loop," and "Human-out-of-the-Loop").

determinations to absolve the user from the legal liability for compensation beyond what he or she has received. However, those laws could state that such determination will not prevent the user from being subject to an accounting of profit. Such a middle approach will likely be important to noncommercial users, as many of them will have limited economic resources and will actively rely on low-cost, or no-cost, machine-made decisions to advance their creative projects.¹⁷⁵

IV. THE FUTURE

Commentators have widely discussed the impact of artificial intelligence on the legal system. As Erik Brynjolfsson and Andrew McAfee reminded us, the change brought about by artificial intelligence will take effect “[g]radually and then suddenly,”¹⁷⁶ recalling Ernest Hemingway’s famous description of how one goes bankrupt in *The Sun Also Rises*.¹⁷⁷ Noting the large-scale ramifications in what they have coined “the Second Machine Age,” Brynjolfsson and McAfee observed:

Progress on some of the oldest and toughest challenges associated with computers, robots, and other digital gear was gradual for a long time. Then in the past few years it became sudden; digital gear started racing ahead, accomplishing tasks it had always been lousy at and displaying skills it was not supposed to acquire anytime soon.¹⁷⁸

Likewise, Lee Kai-fu lamented that “time is one thing that the AI revolution is not inclined to grant us.”¹⁷⁹

In the past few decades, commentators have widely explored how artificial intelligence will affect the legal field.¹⁸⁰ In view of this burgeoning and ever-growing literature, this Part does not intend to rehash prior research. Instead, it focuses on the various lessons we can glean from the earlier discussion of the interplay of artificial intelligence and the law. Although these lessons were drawn from a close analysis of automated fair use systems, they can be easily generalized to inform other bodies of law or the larger legal system. Covering the legislature, the bench, the bar, and academe, this Part underscores

175. See *supra* text accompanying notes 67–68 (discussing the benefits of automated systems in providing low-cost fair use determinations).

176. ERIK BRYNJOLFSSON & ANDREW MCAFEE, *THE SECOND MACHINE AGE: WORK, PROGRESS, AND PROSPERITY IN A TIME OF BRILLIANT TECHNOLOGIES* 20 (2014).

177. ERNEST HEMINGWAY, *THE SUN ALSO RISES* 109 (Hemingway Library ed., 2014) (1926).

178. BRYNJOLFSSON & MCAFEE, *supra* note 176, at 20.

179. LEE, *supra* note 4, at 152.

180. See sources cited *supra* note 8.

our need to carefully analyze the potential impact of technological change on not only the law but also legal institutions.¹⁸¹

A. Law

In *Smart Technologies and the End(s) of Law*, Mireille Hildebrandt asked a highly provocative question concerning whether advances in artificial intelligence will spell the end of the law as we know it.¹⁸² As she observed:

If we do not learn how to uphold and extend the legality that protects individual persons against arbitrary or unfair state interventions, the law will lose its hold on our imagination. It may fold back into a tool to train, discipline or influence people whose behaviours are measured and calculated to be nudged into compliance, or, the law will be replaced by techno-regulation, whether or not that is labelled as law.¹⁸³

In the end, she concluded that whether the law as we know it will “end” “depends on how we design, construct and develop our information and communication infrastructures and how we engage with the mindless agents that will ‘people’ our onli/e world.”¹⁸⁴ To ensure the significantly more desirable outcome, she called on us to “build[] legal protection into our artefactual environment, reinventing recalcitrance . . . as well as the means to generate values and added value in a shared onli/e world that celebrates and affords both democracy and the Rule of Law.”¹⁸⁵

While Professor Hildebrandt was right that the law will still have important roles to play, the growing interplay of artificial intelligence and the law suggests that the role of law will change in at least three distinct ways. First, given the ever-growing algorithmic deployment to make legal decisions at the same time, the line between human and machine-made decisions will increasingly blur. While the law will initially leave most decisions to human decisionmakers, it is only a matter of time before people become more comfortable with machine-made decisions, especially on matters involving narrow or trivial areas. Moreover, if technology has improved to a state where machine-made decisions can closely approximate human decisions, it may be difficult to distinguish between these two types of decisions. Their indistinguishability immediately brings to mind the ongoing discussions in artificial intelligence literature relating

181. See COHEN, *supra* note 47, at 2 (underscoring the need to understand “how both information-economy disputes and new informational capabilities are reshaping the enterprise of law at the institutional level”).

182. HILDEBRANDT, *supra* note 10.

183. *Id.* at xiii.

184. *Id.*; see also LEE, *supra* note 4, at xi (“Our AI future will be created by us, and it will reflect the choices we make and the actions we take.”).

185. HILDEBRANDT, *supra* note 10, at xiii.

to passing the Turing test,¹⁸⁶ machine superintelligence,¹⁸⁷ and technological singularity.¹⁸⁸

Second, because intelligent machines will play increasingly important roles in the legal process, and computer code and algorithms are not as territorially tethered as the law, global and foreign norms will likely have a bigger impact on local decision-making processes than what we currently have in our legal system.¹⁸⁹ Just like how laws that have been transplanted abroad bring values

186. Developed by Alan Turing, this test determines whether one can distinguish between the intelligent behavior exhibited by a machine from that of a human. *See* A.M. Turing, *Computing Machinery and Intelligence*, 59 *MIND* 433 (1950) (advancing the Turing test). Interestingly, Turing believed that humans would be able to create a machine that can pass his test at the end of the twentieth century. *See id.* at 442 (“I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.”). By contrast, Ray Kurzweil set the date much later—at around 2029. *See* RAY KURZWEIL, *THE AGE OF SPIRITUAL MACHINES* 222 (1999); KURZWEIL, *THE SINGULARITY IS NEAR*, *supra* note 128, at 263.

187. *See* AGRAWAL ET AL., *supra* note 17, at 222 (discussing superintelligence in machines). *See generally* RAY KURZWEIL, *THE AGE OF INTELLIGENT MACHINES* (1990) (providing an overview of intelligent machines and exploring whether machines can be intelligent and what it means for them to be so).

188. As the Obama Administration observed in its white paper: People have long speculated on the implications of computers becoming more intelligent than humans. Some predict that a sufficiently intelligent AI could be tasked with developing even better, more intelligent systems, and that these in turn could be used to create systems with yet greater intelligence, and so on, leading in principle to an “intelligence explosion” or “singularity” in which machines quickly race far ahead of humans in intelligence. PREPARING FOR THE FUTURE, *supra* note 1, at 8; *see also* KURZWEIL, *THE SINGULARITY IS NEAR*, *supra* note 128, at 7 (defining “singularity” as “a future period during which the pace of technological change will be so rapid, its impact so deep, that human life will be irreversibly transformed”). *See generally* MURRAY SHANAHAN, *THE TECHNOLOGICAL SINGULARITY* (2015) (providing an overview of technological singularity). *But see* DOMINGOS, *supra* note 75, at 286–89 (challenging Kurzweil’s view on singularity). Benjamin Alarie extends the concept of technological singularity to the legal field. *See* Benjamin Alarie, *The Path of the Law: Towards Legal Singularity*, 66 *U. TORONTO L.J.* 443 (2016). As he explained:

The legal singularity contemplates the elimination of legal uncertainty and the emergence of a seamless legal order, which is universally accessible in real time. In the legal singularity, disputes over the legal significance of agreed facts will be rare. There may be disputes over facts, but, once found, the facts will map onto clear legal consequences. The law will be functionally complete.

Id. at 446.

189. *See* Woodrow Barfield & Ugo Pagallo, *Preface* to RESEARCH HANDBOOK, *supra* note 6, at xxiv, xxv–xxvi (“Artificial intelligence will not be ‘content’ to stay within the geographical boundaries of any particular jurisdiction, or nation state for that matter, therefore to be effective, the regulatory approach to AI will have to be international in scope.”); BRAD SMITH & CAROL ANN BROWNE, *TOOLS AND WEAPONS: THE PROMISE AND THE PERIL OF THE DIGITAL AGE* 300 (2019) (“[T]he inexorable course of technology is forcing more international collaboration. . . . [I]ssues like surveillance reform, privacy protection, and cybersecurity safeguards have all required governments to deal with each other in new ways.”).

with them,¹⁹⁰ technologies that are deployed overseas also export values.¹⁹¹ Langdon Winner rightly reminded us that technological artifacts embody the political, social, economic, and other conditions behind the development of these artifacts.¹⁹² Indeed, as technologies originating from developed and emerging countries are being rapidly and widely deployed throughout the world, one cannot help but wonder whether such deployment will lead to even greater convergence of legal norms, beyond what we have already seen through globalization and the efforts of international organizations and multilateral agreements.¹⁹³

Third, the increasing reliance on machine-based decision-making will have a direct impact on the future development of the legal community. In fact, commentators have already expressed concern that such reliance, and the increased allocation of decision-making power to machines, will undermine the effectiveness of that profession.¹⁹⁴ As Ajay Agrawal, Joshua Gans, and Avi Goldfarb observed in the artificial intelligence context, “If the machines get the experience, then the humans might not.”¹⁹⁵ Growing legal automation could therefore lead to the deskilling of the legal profession,¹⁹⁶ just like how our ability

190. As Alan Watson observed in his seminal work: “Transplanting frequently, perhaps always, involves legal transformation. Even when the transplanted rule remains unchanged, its impact in a new social setting may be different. The insertion of an alien rule into another complex system may cause it to operate in a fresh way.” ALAN WATSON, *LEGAL TRANSPLANTS: AN APPROACH TO COMPARATIVE LAW* 116 (2d ed. 1993); see also Otto Kahn-Freund, *On Uses and Misuses of Comparative Law*, 37 *MOD. L. REV.* 1, 24 (1974) (noting that, because transplanted laws often bring with them foreign values, they may upset longstanding traditions in the recipient countries while at the same time undermining institutions that are “closely linked with the structure and organisation of political and social power in their own environment”); Yu, *Digital Copyright Reform*, *supra* note 31, at 770 (“[If legal transplants] are hastily adopted without careful evaluation and adaptation, they may be both ineffective and insensitive to local conditions. They may also stifle local development while upsetting the existing local tradition.”).

191. See LEE, *supra* note 4, at 18 (noting that “American technology companies . . . were pushing their products and their values on users around the globe”); ERIC SCHMIDT & JARED COHEN, *THE NEW DIGITAL AGE: TRANSFORMING NATIONS, BUSINESSES, AND OUR LIVES* 111–12 (2013) (“Technology companies export their values along with their products . . .”).

192. See LANGDON WINNER, *THE WHALE AND THE REACTOR: A SEARCH FOR LIMITS IN AN AGE OF HIGH TECHNOLOGY* 19–39 (1986) (noting that technological artifacts can embody specific forms of power and authority); see also John Naughton, *Here Is the News—but Only if Facebook Thinks You Need to Know*, *GUARDIAN* (May 15, 2016, 4:00 EDT), <https://www.theguardian.com/commentisfree/2016/may/15/facebook-instant-articles-news-publishers-feeding-the-beast> (“Any algorithm that has to make choices has criteria that are specified by its designers. And those criteria are expressions of human values. Engineers may think they are ‘neutral’, but long experience has shown us they are babes in the woods of politics, economics and ideology.”).

193. See generally Peter K. Yu, *Currents and Crosscurrents in the International Intellectual Property Regime*, 38 *LOY. L.A. L. REV.* 323, 429–35 (2004) (discussing the international harmonization of intellectual property standards).

194. See Michaels, *supra* note 8, at 1096–98 (discussing how the switch from human judges to robot judges would weaken the legal community); Re & Solow-Niederman, *supra* note 8, at 247 (“Increasing use of AI will . . . foster lay and even professional alienation from law as adjudication increasingly moves within the exclusive dominion of technical specialists.”).

195. AGRAWAL ET AL., *supra* note 17, at 192; see also *id.* at 193 (“[E]xperience is a scarce resource, some of which you need to allocate to humans to avoid deskilling.”).

196. See *id.* at 192 (noting the concern that “automation could result in the deskilling of humans”).

to use maps will greatly decrease with our increasing reliance on apps or software utilizing the Global Positioning System.¹⁹⁷ Moreover, the increased use of artificial intelligence may reduce the participation of the existing legal community. As Richard Re and Alicia Solow-Niederman observed, “As AI adjudicators play a larger role in the legal system, human participation will change and, in some respects, decrease. Those developments raise the prospect of alienation, or the tendency for some or all people to cease participating in the legal system and even lose interest in its operations.”¹⁹⁸

B. *Legislature*

As far as the interaction between artificial intelligence and the legislative process is concerned, commentators have explored three broad sets of legislative roles that will help facilitate legal automation. First, the legislature will determine what type of decision can be automated. Second, it will provide assistance to ensure the successful automation of those decisions, including the provision of funding support and the introduction of laws to limit the liability for faulty machine-made decisions¹⁹⁹ and to prevent security breaches and malicious interferences.²⁰⁰ Third, the legislature will provide legal remedies, including institutional mechanisms, to address problems that will arise from the automation of these decisions.²⁰¹ To fashion these remedies, both the government and the legislature will have important roles to play.

One area that has received only limited attention concerns the legislature’s role in determining what type of algorithms could be deemed suitable for automating laws and legal decisions. In a recent article, Professor Elkin-Koren suggested that courts should play some role in making this type of decision.²⁰² As she observed, with the growing use of artificial intelligence and machine learning, they may have to “determin[e] acceptable error rates when testing the outcome of such a system compared to determination by the court.”²⁰³

While I agree with her on the need for determining acceptable error rates, the legislature’s greater fact-finding capacity and its ability to bring in technologists for testimonies will likely make the branch superior for making

197. See Joseph Stromberg, *Is GPS Ruining Our Ability to Navigate for Ourselves?*, VOX (Sept. 2, 2015, 11:31 AM), <https://www.vox.com/2015/9/2/9242049/gps-maps-navigation> (exploring whether the use of the Global Positioning System has undermined our navigation skills).

198. Re & Solow-Niederman, *supra* note 8, at 275.

199. See *infra* note 240 (providing sources examining the legal liability raised by autonomous vehicles).

200. See Crootof, *supra* note 8, at 240 (“Unintended glitches and intended interference from malicious actors create other potential sources of error.”); Volokh, *supra* note 8, at 1171–77 (discussing the potential hacking of the artificial intelligence judge programs and the exploitation of unexpected glitches in those programs).

201. See Mark A. Lemley & Bryan Casey, *Remedies for Robots*, 86 U. CHI. L. REV. 1311 (2019) (discussing the challenges in designing a remedies regime for robots).

202. Elkin-Koren, *supra* note 43, at 1099.

203. *Id.*

this type of determination. If the legislature chooses, it could also create a certification process or an institutional mechanism to help determine what type of algorithm could be deemed suitable for making those determinations.²⁰⁴

The development of this process or mechanism is important for two reasons. First, such development will be needed to address the likely existence of a wide variety of algorithms that could make satisfactory automated fair use determinations.²⁰⁵ Indeed, the diverging algorithms that are being developed will likely involve different trade-offs, such as “more speed, less accuracy; more autonomy, less control; more data, less privacy.”²⁰⁶ Allowing for the existence of multiple algorithms will therefore help increase consumer choices while promoting competition in algorithmic quality.²⁰⁷

Second, past experience has shown that for-profit entities are unlikely to develop a satisfactory arrangement that is in the best interest of the public. As Olivier Sylvain observed:

The ambition to foster “healthy” online engagement, while more than an afterthought, is hardly the Big Tech companies’ main priority. These companies are not (and do not see themselves as) chiefly in the business of calibrating the right balance between human moderators and screening algorithms. Rather, their aim is to hold and expand their dominion over networked information flows.²⁰⁸

204. Cf. Susan Saab Fortney, *Online Legal Document Providers and the Public Interest: Using a Certification Approach to Balance Access to Justice and Public Protection*, 72 OKLA. L. REV. 91, 117–22 (2019) (discussing the benefits of using certification to enhance consumer protection and to promote competition in the market of online providers of automated legal documentation); Yu, *Fair Use*, *supra* note 18, at 358 (noting the need to “set up a neutral and representative body that would supervise the development of fair use algorithms” and a process for certifying algorithms “that are . . . capable of making high-quality decisions”).

205. See AGRAWAL ET AL., *supra* note 17, at 5 (“There is often no single right answer to the question of which is the best AI strategy or the best set of AI tools . . .”).

206. *Id.*; see also PAUL R. DAUGHERTY & H. JAMES WILSON, HUMAN + MACHINE: REIMAGINING WORK IN THE AGE OF AI 126 (2018) (“A deep-learning system . . . provides a high level of prediction accuracy, but companies may have difficulty explaining how those results were derived. In contrast, a decision tree may not lead to results with high prediction accuracy but will enable a significantly greater explainability.”).

207. As I noted in a recent article:

Competition is imperative if society is to develop more efficient, more effective, and less biased algorithms. Such competition is particularly needed when algorithmic choices are increasingly difficult, or time consuming, to explain. Indeed, without competition, it would be hard to identify problems within an algorithm or to determine whether that algorithm has provided the best solution in light of the existing technological conditions and constraints.

Yu, *Algorithmic Divide*, *supra* note 69, at 382–83 (footnotes omitted); see also Peter K. Yu, *Data Producer’s Right and the Protection of Machine-Generated Data*, 93 TUL. L. REV. 859, 927 (2019) (noting that competition law is “a critical area relating to data governance”); Annie Lee, Note, *Algorithmic Auditing and Competition Under the CFAA: The Revocation Paradigm of Interpreting Access and Authorization*, 33 BERKELEY TECH. L.J. 1307, 1310 (2018) (“Online competitors . . . promote fair online practices by providing users with a choice between competitive products . . .”).

208. Sylvain, *supra* note 166, at 264; see also Kroll et al., *supra* note 98, at 682 (“A prejudiced decisionmaker could skew the training data or pick proxies for protected classes with the intent of generating discriminatory results.”).

If the legislature goes the certification route, it will have to take its role seriously, lest it allow justice to be privatized.²⁰⁹ To protect the public, Richard Re and Alicia Solow-Niederman proposed to “remove profit-seeking actors from the market for jurisprudential tools” while calling on the government to “produce a ‘public option’ jurisprudential tool for key purposes, such as criminal justice.”²¹⁰ In earlier articles, I also noted the need to “set up a neutral and representative body that would supervise the development of fair use algorithms.”²¹¹ The creation of this neutral and representative body will be of critical importance if we are to prevent industry lobbies and interest groups from capturing the algorithm design process the same way they would capture the legislative process.²¹²

C. Bench

When artificial intelligence is mentioned alongside judges, an oft-raised question concerns whether we are now ready for machine-generated decisions. In a recent article, Eugene Volokh advanced a highly provocative thought experiment concerning society’s readiness for robot judges.²¹³ His thought experiment went as follows: if an automated system can generate a set of opinions as persuasive as those written by an average human judge in an opinion-writing competition, and if that system can be adequately protected from hacking or other vulnerabilities, that system should be deemed to be “an adequate substitute for humans.”²¹⁴

209. See generally Eldar Haber, *Privatization of the Judiciary*, 40 SEATTLE U. L. REV. 115 (2016) (highlighting the danger of privatization of the judiciary to democratic society).

210. Re & Solow-Niederman, *supra* note 8, at 285.

211. Yu, *Fair Use*, *supra* note 18, at 358; accord Yu, *Anticircumvention and Anti-anticircumvention*, *supra* note 43, at 68 (“[W]e need to develop a process that brings together copyright holders, technology developers, consumer advocates, civil libertarians and other stakeholders.”); see also IAN BROWN & CHRISTOPHER T. MARSDEN, *REGULATING CODE: GOOD GOVERNANCE AND BETTER REGULATION IN THE INFORMATION AGE* 177 (2013) (“[G]reater multistakeholder involvement will improve the quality of regulatory design, including the technical understanding of code.”); COHEN, *supra* note 47, at 192 (“Mastering the processes by which technical standards are developed . . . requires . . . new public accountability mechanisms.”); SMITH & BROWNE, *supra* note 189, at 208 (“[A] global conversation about ethical principles for artificial intelligence will require . . . seats at the table not only for technologists, governments, NGOs, and educators, but for philosophers and representatives of the world’s many religions.”).

212. See sources cited *supra* note 58.

213. Volokh, *supra* note 8.

214. *Id.* at 1138–39. Specifically, Professor Volokh utilized what he described as the “Modified John Henry Test,” which runs as follows:

The way to practically evaluate results is the Modified John Henry Test, a competition in which a computer program is arrayed against, say, ten average performers in some field—medical diagnosis, translation, or what have you. All the performers would then be asked to execute, say, ten different tasks—for instance, the translation of ten different passages.

Sometimes this performance can be measured objectively. Often, it can’t be, so we would need a panel of, say, ten human judges who are known to be experts in the subject—for example, experienced doctors or fluent speakers of the two languages involved in a translation. Those

The aspiration of having robot judges is nothing new. In fact, the literature on the application of artificial intelligence to the law dates back to as early as the 1970s.²¹⁵ While the artificial intelligence we have today is very different from what we had at that time—with the latter featuring mostly mainframes, much more limited processing power, and no big data analytics²¹⁶—many of the legal and ethical questions have remained the same.

Thus far, commentators have widely debated over whether robots should be allowed to take the role of judges.²¹⁷ Even if one agrees with Professor Volokh that robots can eventually succeed in judicial roles and is willing to ignore the fact that our state of technology is still quite far away from that very scenario, judges will still be in a good position to contribute to the better development of the law–machine interface. First, judges can determine what type of technology can be satisfactorily deployed to assist with the adjudication process. As Part II.B has noted, machines can perform certain tasks better than humans. Allowing machines to focus on those specific tasks will provide what commentators have referred to as “intelligence augmentation.”²¹⁸ Such augmentation will free the judges “to focus on more complex legal questions,”²¹⁹ although commentators continue to debate the desirability of hybrid decision-making.²²⁰

judges should evaluate everyone’s performance without knowing which participant is a computer and which is human.

If the computer performs at least as well as the average performer, then the computer passes the Modified John Henry Test. We can call it “intelligent” enough in its field. Or, more to the point, we can say that it is an adequate substitute for humans.

Id. (footnotes omitted).

215. For this literature, see generally Bruce G. Buchanan & Thomas E. Headrick, *Some Speculation About Artificial Intelligence and Legal Reasoning*, 23 STAN. L. REV. 40 (1970); Anthony D’Amato, *Can/Should Computers Replace Judges?*, 11 GA. L. REV. 1277 (1977); L. Thorne McCarty, *Reflections on Taxman: An Experiment in Artificial Intelligence and Legal Reasoning*, 90 HARV. L. REV. 837 (1977). The literature cited here was collected in Volokh, *supra* note 8, at 1137 n.3.

216. See John O. McGinnis, *Accelerating AI*, in RESEARCH HANDBOOK, *supra* note 6, at 40, 42–45 (discussing the evolution of artificial intelligence in the past few decades).

217. On this debate, see generally Crootof, *supra* note 8; Michaels, *supra* note 8; Re & Solow-Niederman, *supra* note 8; Volokh, *supra* note 8; Wu, *supra* note 8.

218. See Pasquale, *A Rule of Persons*, *supra* note 8, at 54 (calling on the legal profession to pursue “a complementary vision of human-machine cooperation” and to focus more on intelligence augmentation); Liza Vertinsky & Todd M. Rice, *Thinking About Thinking Machines: Implications of Machine Inventors for Patent Law*, 8 B.U. J. SCI. & TECH. L. 574, 612 (2002) (“[I]ntelligence augmentation’ allows the effects of automatization to creep up the skill chain, providing for the substitution of white collar jobs by machines and allowing people with less formal training and education to perform more sophisticated tasks.”); Volokh, *supra* note 8, at 1147–52 (discussing the “AI Associate” and “AI Staff Attorney” models); Albert H. Yoon, *The Post-Modern Lawyer: Technology and the Democratization of Legal Representation*, 66 U. TORONTO L.J. 456, 466 (2016) (“Intelligence augmentation . . . reflects a symbiotic relationship between humans and technology. Humans continue to perform the task at hand, but they do so interactively with technology in order to do it better.”).

219. Yoon, *supra* note 218, at 468.

220. While the combined use of human and machine-made decisions has become increasingly common and can generate more desirable outcomes, such hybrid decision-making can also generate outcomes that are less desirable than those made solely by either humans or machines. See sources cited *supra* note 13.

Second, judges can determine how much of the decision-making power should be given to machines,²²¹ especially in so-called hard cases.²²² Even if society prefers to have lawmakers decide the proper allocation of decision-making power between humans and machines, the legislature could still leave some discretion to judges to fine-tune this allocation based on professional experiences and specific circumstances. Because fair use cases involve case-by-case balancing, judges will find it helpful to retain some ability to fine-tune such allocation.

Third, judges will have additional opportunities to influence the development of the law–machine interface. In addition to making individual case-by-case adjustments, they could exert influence as part of an epistemic community.²²³ Indeed, because many jurisdictions are now grappling with questions on the interplay of artificial intelligence and the law, there is an urgent need for an active cross-jurisdictional judicial dialogue. Such a dialogue will not only help achieve consensus at the national, regional, or international level, but will also enhance the judges’ ability to anticipate and address unforeseen challenges in this area.

Finally, judges can share their views with legislators and technologists. With respect to the former, they can weigh in on the key algorithmic design questions discussed in Part III, such as the allocation of decision-making power, the

221. As Richard Re and Alicia Solow-Niederman suggested:

[One] form of human/machine division of labor would apportion discrete types of judicial decision-making to human as opposed to mechanized actors. The resulting separation could be based on subject matter, such as a rule barring automated judging in criminal cases. Or it could derive from more fine-grained determinations about which parts of a legal decision raise concerns about equitable and codified justice. For example, some types of fact-finding could be well-suited for mechanization, without a commensurate cost in disillusionment and alienation, so long as there is a human judge who engages in the analytically severable task of applying the facts to the law. Even within appellate courts, a split in judicial function between human rule-generation and mechanized rule-application might be desirable. More broadly, codified justice already marks key aspects of many bureaucratic legal systems, and AI adjudicators might simply offer a better version of codified justice, limited to those contexts.

Re & Solow-Niederman, *supra* note 8, at 283 (footnotes omitted).

222. As Professor Wu observed:

[One] benefit of human courts over software is their advantages in hard cases, and the prevention of absurd errors, obviously unjust results, and other inequitable consequences of a blind adherence to rules. There are, on closer examination, several ways in which a case can be “hard.” Some cases might be hard only because the software lacks the ability to understand context or nuance, as in understanding that “I’m going to kill my husband” may be a figurative statement, not a death threat. And, others may be hard in the jurisprudential sense because they require the balancing of conflicting values or avoidance of absurd consequence. Finally, it may be that the stakes just seem large enough to merit human involvement, as in the decision to sentence someone to death.

Wu, *supra* note 8, at 2023.

223. See ANNE-MARIE SLAUGHTER, *A NEW WORLD ORDER* 65–103 (2004) (discussing the interactions of judges in a transnational network); see also MICHAEL P. RYAN, *KNOWLEDGE DIPLOMACY: GLOBAL COMPETITION AND THE POLITICS OF INTELLECTUAL PROPERTY* 15 (1998) (noting that epistemic communities “are valuable for their enormous pools of information and their capacities to acquire and generate more”).

hierarchy of decisions, and the legal effects of machine-made decisions.²²⁴ With respect to the latter, judges can educate technologists on how they make decisions and how to think like a lawyer.²²⁵ To the extent that we want to preserve the existing judicial system and to avoid undue disruption by machines, learning how judges make decisions will remain highly important. Such knowledge will be even more important when society has chosen the translation pathway over the other pathways to facilitate legal automation.

D. Bar

Similar to the question about judges, many commentators have questioned our readiness for robot lawyers,²²⁶ including prosecutors, defenders, and associates.²²⁷ Obviously, many questions still remain, ranging from the capacity of intelligent machines to provide legal advice²²⁸ to their ability to effectively handle ethical challenges.²²⁹ Instead of rehashing the answers to these questions, this Subpart turns to a new area that has not received sufficient policy and scholarly attention: the need for new legal personnel to play roles that did not exist before the age of artificial intelligence.

224. See discussion *supra* Part III.

225. See sources cited *supra* note 112.

226. See sources cited *supra* note 8.

227. See generally Kristen Thomasen, *Examining the Constitutionality of Robot-enhanced Interrogation*, in *ROBOT LAW*, *supra* note 16, at 306 (discussing how robot interrogators may engage the fundamental constitutional rights to privacy and silence); Volokh, *supra* note 8, at 1147–52 (discussing artificial intelligence-driven associates and staff attorneys).

228. As my colleague Milan Markovic aptly observed:

Regardless of their level of sophistication, clients often do not have clear objectives and require assistance in shaping them. Clients also sometimes misunderstand the legal system and do not view their situations, including any wrongs they may have suffered, in legalistic terms. A fully autonomous, composed, and decided client may not require the counseling of an attorney, but that is not the messy reality of the law as lived.

....

[Moreover, a]n intelligent machine may be able to determine if a course of conduct is unlawful; it may also be able to calculate the probability that any misconduct will be detected. What it cannot do is fulfill the other crucial “half” of a lawyer’s role: shaming and persuading clients and would-be clients “that they are damned fools and should stop.” As David Luban has explained, intelligent machines lack emotional intelligence and moral authority and cannot buttress legal and non-legal considerations to exhort clients to act in accordance with the law.

Markovic, *supra* note 8, at 344–46 (footnotes omitted).

229. See Drew McDermott, *Why Ethics Is a High Hurdle for AI 2* (Feb. 29, 2008), <http://www.cs.yale.edu/homes/dvm/papers/ethical-machine.pdf> (“[E]thical behavior is an extremely difficult area to automate, both because it requires ‘solving all of AI’ and because even that might not be sufficient.”); see also DOMINGOS, *supra* note 75, at 280 (“[L]etting robots learn ethics by observing humans may not be such a good idea. The robot is liable to get seriously confused when it sees that humans’ actions often violate their ethical principles.”).

In their widely cited book on big data, Viktor Mayer-Schönberger and Kenneth Cukier discussed the future need for algorithmists.²³⁰ As they explained:

These new professionals would be experts in the areas of computer science, mathematics, and statistics; they would act as reviewers of big-data analyses and predictions. Algorithmists would take a vow of impartiality and confidentiality, much as accountants and certain other professionals do now. They would evaluate the selection of data sources, the choice of analytical and predictive tools, including algorithms and models, and the interpretation of results. In the event of a dispute, they would have access to the algorithms, statistical approaches, and datasets that produced a given decision.²³¹

Applying these insights to the present context, one cannot help but wonder whether two new types of legal professionals will emerge: algorithmically oriented lawyers and legal algorithmists.

Given the important and ever-growing roles of intelligent machines in the legal process and the growing importance of addressing issues at the law–machine interface, we will need to have lawyers that have a good grasp of artificial intelligence and what the latest technology can and cannot do.²³² The importance of algorithmic literacy²³³ has caused commentators and educators to emphasize the importance of computational thinking.²³⁴ In the future, those

230. See VIKTOR MAYER-SCHÖNBERGER & KENNETH CUKIER, *BIG DATA: A REVOLUTION THAT WILL TRANSFORM HOW WE LIVE, WORK, AND THINK* 180–82 (2013) (discussing the need for external and internal algorithmists).

231. *Id.* at 180.

232. See Crotofof, *supra* note 8, at 244 (“If we wish to elicit the benefits of human reasoning, teaming systems must be designed so that the human in the loop understands the AI program’s capabilities and limitations, has reason to exercise valued human skills, and is actively engaged in the decisionmaking process.” (footnote omitted)); Frank Pasquale, *Data-Informed Duties in AI Development*, 119 COLUM. L. REV. 1917, 1939 (2019) (“[T]here is a parallel duty for technology providers to have some basic understanding of the law as they serve their clients.”).

233. See INST. ELEC. & ELEC. ENG’RS, *supra* note 80, at 142 (“Improving digital literacy of citizens should be a high priority for the government and other organizations.”); RAINIE & ANDERSON, *supra* note 72, at 74–76 (surveying views on the need for algorithmic literacy); U.N. EDUC., SCI. & CULTURAL ORG. [UNESCO], *ARTIFICIAL INTELLIGENCE IN EDUCATION: CHALLENGES AND OPPORTUNITIES FOR SUSTAINABLE DEVELOPMENT* 6–7 (2019) (stating that “teachers must learn new digital skills to use AI in a pedagogical and meaningful way”); U.S. AGENCY FOR INT’L DEV. [USAID], *REFLECTING THE PAST, SHAPING THE FUTURE: MAKING AI WORK FOR INTERNATIONAL DEVELOPMENT* 74 (2018) (“Strengthening training programs for data science and machine learning in local development contexts can help create a pipeline of individuals who are ‘bilingual’ in the sense of understanding local context and having the technical skills to take an active role in developing [machine learning] tools.”); Yu, *Algorithmic Divide*, *supra* note 69, at 362–65 (discussing the need to increase algorithmic literacy).

234. The International Society for Technology in Education and the Computer Science Teachers Association provided the following operational definition of computational thinking:

Computational thinking (CT) is a problem-solving process that includes (but is not limited to) the following characteristics:

- Formulating problems in a way that enables us to use a computer and other tools to help solve them[]
- Logically organizing and analyzing data
- Representing data through abstractions such as models and simulations

lawyers who are equipped with a better understanding of the technological aspects of the legal decision-making process will likely be in better positions to serve their clients than those who do not or who rely solely, or mostly, on technology experts to provide gap-filling advice. The need for algorithmically oriented lawyers therefore arises.

The flip side is also true. Just as society needs to have algorithmically oriented lawyers, it also needs to have legal algorithmists. While internal algorithmists conduct audits inside the developers of automated systems,²³⁵ external algorithmists undertake evaluation from the outside and fulfill roles designated by the legislature or regulatory authorities.²³⁶ These algorithmists are legal algorithmists because they have a specialized focus on legal technology and on other technologies that have serious ramifications for the legal system.

E. *Academe*

As far as academic research is concerned, there is no shortage of materials on artificial intelligence and the law.²³⁷ In fact, law schools and legal

-
- Automating solutions through algorithmic thinking (a series of ordered steps)
 - Identifying, analyzing, and implementing possible solutions with the goal of achieving the most efficient and effective combination of steps and resources
 - Generalizing and transferring this problem solving process to a wide variety of problems

Int'l Soc'y for Tech. in Educ. & Comput. Sci. Teachers Ass'n, *Operational Definition of Computational Thinking for K–12 Education*, INT'L. SOC'Y TECH. EDUC., <https://id.iste.org/docs/ct-documents/computational-thinking-operational-definition-flyer.pdf> (last visited Aug. 19, 2019). *See generally* PETER J. DENNING & MATTI TEDRE, *COMPUTATIONAL THINKING* (2019) (providing an overview of computational thinking).

235. As Mayer-Schönberger and Cukier observed:

Internal algorithmists work inside an organization to monitor its big-data activities. They look out not just for the company's interests but also for the interests of people who are affected by its big-data analyses. They oversee big-data operations, and they're the first point of contact for anybody who feels harmed by their organization's big-data predictions. They also vet big-data analyses for integrity and accuracy before letting them go live. To perform the first of these two roles, algorithmists must have a certain level of freedom and impartiality within the organization they work for.

MAYER-SCHÖNBERGER & CUKIER, *supra* note 230, at 181–82.

236. As Mayer-Schönberger and Cukier elaborated:

We envision external algorithmists acting as impartial auditors to review the accuracy or validity of big-data predictions whenever the government requires it, such as under court order or regulation. They also can take on big-data companies as clients, performing audits for firms that want expert support. And they may certify the soundness of big-data applications like anti-fraud techniques or stock-trading systems. Finally, external algorithmists are prepared to consult with government agencies on how best to use big data in the public sector.

Id. at 181.

237. In only a few years, a vast literature has quickly built up on the question of whether creative works generated by intelligent machines are eligible for copyright protection. *See generally* Annemarie Bridy, *Coding Creativity: Copyright and the Artificially Intelligent Author*, 2012 STAN. TECH. L. REV. 5; Annemarie Bridy, *The Evolution of Authorship: Work Made by Code*, 39 COLUM. J.L. & ARTS 395 (2016); Daniel J. Gervais, *The Machine as Author*, 105 IOWA L. REV. 2053 (2020); Jane C. Ginsburg & Luke Ali Budiardjo, *Authors and Machines*, 34 BERKELEY TECH. L.J. 343 (2019); James Grimmelmann, *There's No Such Thing as a Computer-Authored Work—And It's a Good Thing, Too*, 39 COLUM. J.L. & ARTS 403 (2016); Lim, *supra* note 129, at 836–47; Carys J. Craig & Ian R. Kerr, *The Death of the AI Author* (Osgoode Hall Law Sch. Legal Studies Research Paper Series, 2019),

commentators have been actively organizing symposia and book projects to address questions arising at the intersection of artificial intelligence and the law.²³⁸

While many questions have been explored in regard to whether artificial intelligence will change the outcome of legal analysis—such as whether works created by artificial intelligence are eligible for copyright or patent protection²³⁹ or whether accidents caused by autonomous vehicles deserve the same type of legal liability²⁴⁰—it is time that academics explored whether some of these questions will have to be asked differently.

In the example concerning the development of automated fair use systems, the previous discussion has shown different pathways for legal automation: translation, approximation, and self-determination.²⁴¹ At the moment, we do not have enough evidence—empirical or otherwise—to inform whether one pathway will promote creativity better than the others. We also do not have sufficient research concerning the law–machine interface or how to facilitate a more optimal division of labor between humans and machines in the legal system.²⁴² Considering that the future of this system may be quite different from what we have today, it may be wise to start anticipating these potentially transformative changes and exploring what this process will become.

In addition to new thinking and research, academe needs to evaluate existing curricula and pedagogies to determine whether they are equipped to

<https://ssrn.com/abstract=3374951>. For a provocative discussion of the role of robots in copyright’s cosmology, see generally James Grimmelman, *Copyright for Literate Robots*, 101 IOWA L. REV. 657 (2016). For earlier discussions of copyright issues involving computer-generated works, see generally Ralph D. Clifford, *Intellectual Property in the Era of the Creative Computer Program: Will the True Creator Please Stand Up*, 71 TUL. L. REV. 1675 (1997); Arthur R. Miller, *Copyright Protection for Computer Programs, Databases, and Computer-Generated Works: Is Anything New Since CONTU?*, 106 HARV. L. REV. 977, 1042–72 (1993); Pamela Samuelson, *Allocating Ownership Rights in Computer-Generated Works*, 47 U. PITT. L. REV. 1185 (1986).

238. For symposia in the artificial intelligence area, see generally Symposium, *Artificial Intelligence and the Law*, 89 WASH. L. REV. 1 (2014); Symposium, *Artificial Intelligence, Big Data, and the Future of Law*, 66 U. TORONTO L.J. 423 (2016); Symposium, *Common Law for the Age of AI*, 119 COLUM. L. REV. 1773 (2019); Symposium, *Rise of the Machines: Artificial Intelligence, Robotics, and the Reprogramming of Law*, 88 FORDHAM L. REV. 381 (2019); *2019 Annual BCLT/BTLJ Symposium*, BERKELEY CTR. FOR LAW & TECH., <https://www.law.berkeley.edu/research/bclt/bcltevents/2019bcltblj-symposium> (last visited Sept. 13, 2020); “*Smart Law and Intelligent Machines*” Symposium, TEX. A&M UNIV. SCH. OF LAW, <http://law.tamu.edu/faculty-staff/news-events/conferences-and-symposia/smart-law-and-intelligent-machines-symposium> (last visited Sept. 13, 2020). For book projects and law review articles, see sources cited *supra* note 8.

239. See sources cited *supra* note 237.

240. For discussions of the legal liability raised by autonomous vehicles, see, for example, MARK CHINEN, *LAW AND AUTONOMOUS MACHINES: THE CO-EVOLUTION OF LEGAL RESPONSIBILITY AND TECHNOLOGY* 52–101 (2019); HANNAH YEEFEN LIM, *AUTONOMOUS VEHICLES AND THE LAW: TECHNOLOGY, ALGORITHMS AND ETHICS* 20–98 (2018); Bryant Walker Smith, *Automated Driving and Product Liability*, 2017 MICH. ST. L. REV. 1; Bryant Walker Smith, *Proximity-Driven Liability*, 102 GEO. L.J. 1777 (2014); David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117 (2014). See generally SAMIR CHOPRA & LAURENCE F. WHITE, *A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS* 119–51 (2011) (discussing tort liability for artificial agents).

241. See discussion *supra* Part II.

242. See discussion *supra* Part III.

train the next generation of lawyers.²⁴³ As the previous Subpart has noted, the need for algorithmically oriented lawyers will only continue to grow, and those lawyers who have high algorithmic literacy will be in better positions to help clients than those who do not.²⁴⁴ By introducing up-to-date curricula and pedagogies, law schools and other legal education providers will be able to train lawyers to take full advantage of the growing deployment of automated systems and artificial intelligence technologies in the legal field while at the same time responding effectively to the changes and challenges posed by these new technologies.

CONCLUSION

The age of artificial intelligence has brought to the legal field many thorny and complex questions. While some of them resemble questions that we are already asking in the legal discipline, or are extensions of those questions, others are novel and will require new legal, technological, or techno-legal insights.²⁴⁵ By utilizing the case study of fair use automation, this Article calls for greater attention not only to the impact of artificial intelligence on the law but also to the law–machine interface. If the impact of artificial intelligence in other areas of society is any guide,²⁴⁶ the technological advances in this area will likely precipitate profound changes to the legal system. The sooner we start thinking about these changes, the quicker we can harness these technological advances to improve the law, legal institutions, and the legal process, and the better off society will be.

243. See sources cited *supra* notes 232–234.

244. See discussion *supra* Part IV.D.

245. See Peter K. Yu, *Teaching International Intellectual Property Law*, 52 ST. LOUIS U. L.J. 923, 939 (2008) (“As [technological and legal protections] interact with each other, and improve over time, they result in a technolegal combination that is often greater than the sum of its parts. It is therefore important to understand not only law and technology, but also the interface between the two.”).

246. See generally BRYNJOLFSSON & MCAFEE, *supra* note 176 (examining the transformative impacts of emerging digital technologies on jobs and the economy); CARL BENEDIKT FREY, *THE TECHNOLOGY TRAP: CAPITAL, LABOR, AND POWER IN THE AGE OF AUTOMATION* (2019) (discussing the changing interplay of capital, labor, and power in the age of automation).