

University of Kentucky

UKnowledge

Law Faculty Books and Chapters

Law Faculty Publications

2020

Benefits, drawbacks, and risks of AI

James M. Donovan

University of Kentucky College of Law, james.donovan@uky.edu

Follow this and additional works at: https://uknowledge.uky.edu/lawfac_book



Part of the [Digital Communications and Networking Commons](#), and the [Law Librarianship Commons](#)

10

Benefits, Drawbacks, and Risks of AI

The specter of the impact of artificial intelligence (AI) on law and legal education casts a long and uncertain shadow. Its mix of enthusiasm and trepidation can arise from a thin idea of what the label means. Four components differentiate AI from even high-end automation: big-data and predictive analytics, deep-learning software, cloud computing, and natural language processing (NLP).¹ The perspective of the person on the street, though, is captured as “the art of creating machines that perform functions that require intelligence when performed by people,” centering on the ability to make independent choices.² While that gloss might fail as a technical specification, it appears to capture the ordinary meaning and will suffice for the purposes of this chapter, with only one additional clarification.

In the following discussion, we focus on what exists today: narrow (or weak) AI. Narrow AI refers to algorithms that are application-specific, as compared to the generalized AI that mimics the human

JAMES M. DONOVAN is the director of the law library and James and Mary Lassiter associate professor of law at the University of Kentucky College of Law Library.

ability to learn and perform on any topic. Generalized (or strong) AI does not currently exist, and some experts are skeptical that it ever will.³ If society decides that even this limited version of AI is undesirable, all things considered, then that determination should preclude any need to separately evaluate the general variety. For the time being, therefore, we are justified in considering only the narrow version.

By that definition, we are already immersed in an AI-rich environment. The appeal is obvious: A dispassionate computer algorithm returns lower error rates when compared to humans, in part because it reasons from data without being swayed by emotions or preferred outcomes. A common example of everyday AI application occurs whenever Amazon shoppers find book recommendations based on their past views and purchases. AI's ability to find hidden patterns received a powerful demonstration when Target identified a young customer whose shopping habits signaled a pregnancy that was unknown to her parents, who had complained when the store started sending baby-related coupons.⁴ AI helps spam filters learn to exclude emails of a certain kind and not just those that contain certain words. Our phones work to understand our individual dialects. With the development of self-driving cars, intelligent refrigerators, and medical diagnostics, it will be a rare individual who can claim that they have insulated themselves from this new world of nonhuman thinking and learning. Such a ubiquitous impact cannot help but reach into the practice of law.

Benefits of AI

Always seeking to gain a competitive advantage, law firms have a tradition of being early adopters of new technology (perhaps second only to health care). Legal professionals had earlier demonstrated their openness to cutting-edge methods with their quick acceptance in the seventies of one of the first full-text data services developed by LEXIS. In keeping with that pragmatism, the field has implemented a range of useful AI applications to assist in the efficient performance of routine tasks. The following are representative examples of these valuable benefits.

Legal Research

The major commercial platforms have incrementally incorporated AI into their existing tools. Most apparent to users will be the improved ability for the system to comprehend natural language, both in the queries and in the documents. This achievement expands legal research results by the identification of patterns of conceptual use and other

variables beyond the strict keyword or headnote analysis that was the standard only a short time ago. Further functionality analyzes past searches and retrievals to identify the kinds of materials of interest to the patron, and then suggests additional items the algorithm concludes as being helpful. These improvements have resulted in more efficient searching, capable of identifying relevant authorities in a fraction of the time when applying traditional Boolean techniques.

Jury Selection

The hallmark of AI is the use of “Big Data”—the manipulation of large, complex datasets, too unwieldy for even simple processing by ordinary applications—to identify hidden patterns as described in the Target anecdote. This represents a qualitative change: We don’t just know the usual things better and faster, but we have access to new and different kinds of relationships than were discernible just a few years ago.

Voltaire (<https://voltaireapp.com/>) is a legal AI company that leverages this potential to give litigators deep insights in real time into the psyches of potential jurors. As described by the company, the service “explores all public data related to the potential juror, correlates the data against known patterns in human behavior, and then produces a detailed profile, with indications of the type of person they are and how their views and biases may be a positive or negative factor as part of a jury.”⁵

Outcome Prediction

For several years, attorneys have had the ability to examine how certain kinds of cases typically fare in specific jurisdictions or before certain judges, although without offering a determination on how one’s current case might fit into this pattern. Because it can access more data and identify hidden patterns not readily discernible to the human expert, AI makes possible that critical last step. For example, the London firm of Hodge Jones & Allen invested in a tool to predict likely outcomes of personal injury cases in order to help decide whether to settle.⁶ Additional work is begin pursued that will allow similar determinations, but at the level of the deciding judge.⁷

E-Discovery

Empirical research has demonstrated that technology-assisted review of large document sets to find relevant documents is both more effective at finding desirable items and accomplishes this superior result

more efficiently—and more cheaply—than traditional manual review.⁸ Beyond the search for mere keywords, AI identifies context, concepts, and tone through the application of predictive coding. Predictive coding—the marking by experts of sample documents that the program employs to identify similar items—enjoys official judicial approval after it was declared that “computer-assisted review is an available tool and should be seriously considered for use in large-data-volume cases where it may save the producing party (or both parties) significant amounts of legal fees in document review.”⁹

Contracts

The first successful use of AI in the law firm was LONald in 2015 by the firm of Berwin Leighton Paisner for the purpose of contract due diligence.¹⁰ Because AI is able to manipulate concepts as well as words, it has proven to be a valuable tool to generate consistent agreements. Although the greater speed and accuracy of AI systems as compared to human performance is often taken as given, contract analysis is one of the few areas in which the two sides have directly competed. When evaluating five Non-Disclosure Agreements for problematic issues, the AI “achieved an average 94 percent accuracy rate, ahead of the lawyers who achieved an average rate of 85 percent.”¹¹ The time required to reach these results was even more dramatically different: “On average, it took 92 minutes for the lawyer participants to complete all five NDAs. . . . In contrast, the AI engine completed the task of issue-spotting in 26 seconds.”

Personnel

As the results of the head-to-head comparison suggest, a key benefit of the new technology is that the proven superiority in performing the routine tasks of legal practice can help firms to control personnel costs. Although several authors have been predicting that AI will result in major disruptions in the legal market—and despite clients pushing firms to adopt these tools to realize lower charges on their service bills—thus far, the rise in AI applications has not directly impacted the numbers finding legal employment.¹²

Still, fear of displacement could lead practitioners to dismiss the benefits of AI. One of the authors of the e-discovery effectiveness study, Maura Grossman, has acknowledged that some attorneys harbor a resistance to admitting that a machine can do a task better than humans. She described an experiment showing that, even after a demonstration that the algorithm can better predict outcomes, subjects bet against it

7.⁸
is,
d-
m
er
id
re
of

in favor of their own judgments. Grossman interprets this pattern as illustrating a tendency of some lawyers to be more critical of machine failures but highly tolerant of human error, perhaps feeling that “there but for the grace of God. . . .”¹³

Instead of real numbers declining, what we have seen are changes in the kinds of work that attorneys now are able to do. Freed from rote tasks, attorneys will have more opportunity to undertake creative tasks that provide an intellectual and professionally stimulating challenge. The avoidance of unnecessary stress and frustration can lead to fewer attorneys experiencing burnout.¹⁴

Drawbacks of AI

Benefits of the thoughtful development and deployment of AI applications accrue to almost every fact of legal practice. However, we must also recognize that even when AI performs as intended, it can generate negative repercussions. These drawbacks require that AI not merely be adopted solely for the upside; a cost-benefit analysis should be performed in each case to determine if benefits outweigh the possible costs.

Potential problems associated with the introduction of AI into daily legal practice arise at two discrete levels of concern. The first considers how AI can create challenges within law, while the second acknowledges that society at large has its own interests in such developments. Whether new applications that offer clear benefits to the profession should be shelved because of the broader ramifications remains a legitimate question. The following are examples of the potential drawbacks.

Costs

Most products must be developed inside each firm to take advantage of its specific expertise. Creation of AI utilities can be expensive in terms of both time and money. One vendor estimates that a typical AI application can cost up to \$300K to create and train.¹⁵ The firm’s expectation is that in the long run, it will be able to recoup this expenditure: Businesses making this investment are estimated to realize a 38 percent revenue boost.¹⁶ This view, however, ignores the opportunity costs that can be required. The machine learning (ML) that is at the core of successful AI development requires that experts who would otherwise spend their time billing for services to clients must instead spend their time training the system by coding documents and evaluating results. Some practices might not be in a financial position to make this investment, putting them at further disadvantage relative to those that are.

Unintended Consequences

A product intended to offer strategic advantage can lead to unexpected outcomes, some of them undesirable. As previously mentioned, many attorneys worry that if research results can now be obtained in minutes rather than hours, this could result in fewer personnel who would otherwise be hired to perform such tasks. Although firm staffing has not significantly seen reductions for this reason, that could change down the road.

The more immediate issue concerns not staffing, but billing. *Cass*, a Canadian case (discussed in chapter 7), denied a research bill as excessive, in light of the fact that “If artificial intelligence sources were employed, no doubt counsel’s preparation time would have been significantly reduced.”¹⁷ The more efficiently work is performed, the less likely clients are to pay the costs traditionally charged for that work product. Such a change would force a move away from the billable hour and the adoption of alternative arrangements—as yet unspecified—regarding how firms bill for services and, indirectly, how they pay their attorneys.¹⁸

Compulsory AI

Although few attorneys would mourn the loss of the billable hour, the movement toward the ordinary employment of AI tools points in another direction that they might find less attractive. At last count, 35 states have adopted the ABA Model Rule 1.1 (Duty of Competence) revised Comment 8, which requires technological knowledge for ethical legal practice.¹⁹ The rule encourages attorneys to have comparable relevant expertise as those with whom they frequently interact (e.g., C-level executives) and to ensure that client services are streamlined to maximize efficiency, privacy, and cost effectiveness.

When first proposed in 2012, the understanding of the standard focused on the lawyer’s understanding of matters such as how cloud computing can be ethically utilized to store client documents, proper management of social media, or how to strip inappropriate metadata from electronic documents before sharing.²⁰ Today, though, reasoning in cases like *Cass*, which underscores the efficiency of AI-powered researching, can, when paired with the duty of technological competence, lead to the further conclusion that the use of AI is not merely desirable and beneficial, but actually necessary to the competent practice of law.²¹ There have been hints that a similar progression is already

gaining ground in the practice of medicine.²² Should the adoption of AI tools become not a strategic option but either formally or de facto compulsory for the ethical and competent practice of law, the impact would fall heaviest on smaller practices that lack the resources to create their own tools or find the subscriptions to proprietary applications prohibitively expensive.

Lack of Creativity

Earlier it was mentioned that today's AI is of a narrow or weak variety. In other words, it is able to handle only a well-defined set of tasks. An AI trained in contract evaluations, for example, is unable to function on predicting litigation outcomes. Importantly, this is not simply due to a lack of training; the tool lacks the infrastructure to recognize and process the relevant data, even if exposed. Strong or general AI, on the other hand, would be able to process whatever the request, in much the same manner as a human can perform nearly any task required.

Although narrow AI is sufficient for many purposes, we should be aware of what its functional limitations include. AI (at least thus far) can only tell us, given past examples of applying current rules, what will be the likely legal solution to a query.²³ What it cannot tell us, though, is the *right* thing to do. For example, given the parameters under which the platform operates, almost by definition it can never counsel that the situation requires altering the system itself, perhaps by jury nullification or by arguing to reject the doctrine of stare decisis. Consequently, an AI-driven legal system will, on the whole, tend to further entrench the status quo.

One might respond to this observation by noting that AI only offers recommendations and that it is always up to the practitioner's professional judgment as to what should actually be done. In practical terms, as already suggested, professionals are often ill-advised to not contradict the AI system because it normally performs better than people. However, problem solving by humans frequently requires more than the mere application of rules, often needing creative connections and insights built upon cross-disciplinary knowledge that is inaccessible to the narrow AI. Although most legal tasks do not require such expansive resources, we might not be well served by growing too dependent on a system that excludes them.

Not all of the potential drawbacks of an AI-dominated legal industry will fall on the profession itself. Society at large risks exposure to more generalized ramifications.

Overdependence on Automation

A repeated complaint among law librarians is that the “Googlization” of Westlaw and Lexis has resulted in law students increasingly relying on the machine to do work that lawyers should perform. Perhaps much as the use of calculators in math class allowed students to externalize their understanding of the underlying principles, law students, because the discovery of results no longer depends on knowing beforehand exactly where the information sought is likely to reside, tend to come comparatively later to disciplinary defaults (such as that public laws will be found in the *Statutes at Large*). AI functionality promises to exacerbate this reliance. Instead of using the tool as an assist, it threatens to become a substitute for the professional expertise of the attorney.

Recall that algorithms cannot make strategic decisions for the case, but only assign a statistical likelihood to the differing options that satisfy the search criteria. Final decision falls to human consumers, who, if they are to remain sufficiently skilled to make that call, must retain the same level of material and methodological expertise as existed in the pre-AI environment. Yet the effort to learn something that a machine will normally do better will prove a poor enticement, with perhaps little reward, not only in law, but in all manners of professional and technical expertise. Society must be aware of the altering incentives for training and education, and shape its curricula and incentives accordingly.

Professional Status

While the unsettling impact of AI on the legal profession is frequently discussed, the effects of this disruption are not a purely internal concern. Law must maintain not only an actual level of efficiency and expertise, but be perceived as doing so. Otherwise the social contract that underlies its distinctive status as a profession might be deemed void. As already mentioned, overreliance on automated legal expertise could advance far enough that it impacts what attorneys are expected to know for themselves. The resulting changes in perceptions as to what it means to be a “lawyer” could lead to a decreased respect and a grant of self-regulation, should practitioners come to be viewed as mere technicians, while the machines do most of the serious work. Without prudent integration of these powerful knowledge tools into the practice of law, the field could risk losing recognition as one of the archetypical professions, with all that flows from that status.

Legal Diversity

One of the lessons from legal anthropology concerns the unintended consequences of a well-intentioned change in the accessibility of legal

information. For example, when Botswana was a part of the British Empire, the administrators asked Isaac Schapera to collect the oral legal traditions of the local peoples. To his disappointment, the resulting descriptive ethnography came to be applied as a prescriptive code of what the law should be, thereby freezing in place what to that point had been a fluid and flexible system of law.²⁴

Because law is fundamentally a set of rules and procedures, in theory the entire judicial process should be ideal for intelligent automation. In practice, however, laws exist to govern human behavior and form a framework for society. While consumers can hope that AI would be capable of recognizing and reflecting local variation on legal norms, the more likely outcome is that the algorithms will be trained on data from a specific jurisdiction. Exposing them to all regional variations, with a corresponding ability to limit analysis to a subset thereof, might not always be practical. Increased legal automation could therefore result in an unplanned legal homogenization. To the extent that we judge as desirable variations around what society considers acceptable or unacceptable, we can debate whether automating justice should be allowed to outpace our ability to program that diversity into the algorithm.

Even functioning as designed, AI can create problems while solving others. Now that a sizeable record of AI performance is available, the discussion has occasionally turned to risks arising when AI operations deviate from their designed parameters. They will still be following their programming, but not always performing in the manner the creators intended.

Bias

The marquee benefit of AI algorithms—the ability to predict future outcomes based on past events—has a dark side. Given the way in which AI self-programs, the means by which it arrives at its decisions cannot be fully known. In the literature, this is known as “the black box” problem, which became an issue in *State v. Loomis*.²⁵ In that case, a defendant who was sentenced based on a recidivism assessment derived from a commercial algorithm called “COMPAS” (an acronym for Correctional Offender Management Profiling for Alternative Sanctions). The court held that the use of an automated assessment was not a violation of his due-process rights, even though neither the court nor the defendant had access to the algorithm to better understand how it had arrived at its determination that the defendant was at risk to recommit.

It is worth noting that the EU has chosen a different route than the USA. Its newly effective General Data Protection Regulation explicitly limits the extent to which a person can be subject to automated judgments, going further to incorporate a right to explanation concerning

any such automated decision when it occurs.²⁶ How this second rule will be applied is yet unknown given that, due to the black box problem, it might not be physically possible to offer the explanation that the law requires and which the subject needs to mount a meaningful rebuttal.

In any event, the *Loomis* court's certainty that no rights were violated by the procedural reliance upon automated risk assessments should be questioned in light of ProPublica's finding that COMPAS reported that "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."²⁷ In other words, AI algorithms have proven capable of being as biased as the societies that created them, contrary to the expectations of their programmers.

Bias of one kind or another is a common result, certainly not limited to this one product. The factors that skewed the COMPAS predictions have impacted other systems as well, such as one used to predict child abuse.²⁸ Hewlett Packard computers equipped with facial-tracking software that functioned with white faces proved unable to recognize African Americans; similarly, an AI-judged beauty pageant appeared unfavorable to nonwhite contestants.²⁹ Google was shown to display fewer employment ads for higher paying jobs when the user was marked as female than as male.³⁰ The list of similarly disturbing AI examples of the irrelevant biases of humans could be greatly expanded.

These are all examples of inaccurate biases due to training on historical data that incorporates those biases. Those assumptions then become part of the pattern learned by the system to interpret future cases. Even when overt markers of race, sex, etc., are removed in an attempt to guide the algorithm away from considering those variables, the system can discover proxy indicators (e.g., ZIP codes for race) that yield the same outcomes.

The risk for the profession, as well as for society in general, is that these biases might not be recognized. If, as considered earlier, attorneys tend to defer to the system's suggestions, we could find these biases not only propagated but also routinized for the future.

To avoid such cases, it has been suggested that we should create an "FDA"-type watchdog for algorithms.³¹ In line with that idea, the previous Congress saw the introduction of the "FUTURE of Artificial Intelligence Act of 2017," which had as one purpose to study "how bias can be identified and eliminated in the development of artificial intelligence."³² Although national legislation on this problem remains a distant prospect, New York City has become the first jurisdiction to require that all automated decision systems be reviewed for "equity, fairness, and accountability."³³

Legal Responsibility

In 2015, a factory robot grabbed and killed a man.³⁴ A self-driving car killed a pedestrian in 2018.³⁵ Who bears legal responsibility for these acts? The machine? The programmer? The owner or user?

Jacob Turner's recent book, *Robot Rules*, considers these alternatives in some detail. The issues are complex and will require attorneys and policy makers to think more deliberately about assumptions they had previously taken as self-evident. Unlike other machines, AI is exceptional because it makes moral choices and can develop independently.³⁶ The second feature especially suggests that the existing product liability regimes cannot be extended to encompass the bad actions of AI.³⁷

Alternatively, we must consider whether the AI itself should be held responsible. Nonhuman entities can already have legal standing and responsibility. Corporations are perhaps the most obvious example, and some advocate for the legal personality of animals. Where will AI fall in this continuum, especially given developments such as Saudi Arabia's grant of "citizenship" to a humanoid robot?³⁸ To the extent that legal responsibility follows choice and understanding of consequences, AI seems quite close to what we require of humans when holding them liable for actions.

If those questions seem a ways off for legal practices, we can begin with a relatively simple problem closer to home, such as who is responsible for a false positive or negative produced by the AI engine upon which the attorney subsequently relies? If it veers toward malpractice to fail to use AI, as the direction we are headed in suggests, who is ultimately responsible for the bad decision?

Even in his text, Turner is unable to definitively answer such questions, but he reminds us that as we work to create tools whose primary quality is their approximation of human intelligence, they might also come to shoulder some of the burdens of that accomplishment. It will fall to lawyers using good reasons to draw the necessary distinctions and stake out the defensible boundaries.

NOTES

1. Joanna Goodman, *Robots in Law: How Artificial Intelligence Is Transforming Legal Services* (London: Ark, 2016), 7.
2. Ray Kurzweil, *The Age of Intelligent Machines* (Cambridge, MA: MIT Press, 1992), 14; Jacob Turner, *Robot Rules: Regulating Artificial Intelligence* (London: Palgrave Macmillan, 2019), 15.
3. Turner, *Robot Rules*, 6.

4. Charles Duhigg, "How Companies Learn Your Secrets," *New York Times*, February 16, 2012, <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html?pagewanted=1&r=1&hp>.
5. "Voltaire Uses AI and Big Data to Help Pick Your Jury," *Artificial Lawyer*, April 26, 2017, <https://www.artificiallawyer.com/2017/04/26/voltaire-uses-ai-and-big-data-to-help-pick-your-jury/>.
6. Michael Cross, "Role of Artificial Intelligence in Law," *Raconteur*, February 19, 2015, <https://www.raconteur.net/risk-management/time-for-technology-to-take-over>.
7. Daniel L. Chen, "Machine Learning and the Rule of Law," in *Computational Analysis of Law*, eds. M. Livermore and D. Rockmore (Santa Fe, NM: Santa Fe Institute Press, 2019).
8. Maura R. Grossman and Gordon V. Cormack, "Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review," *Richmond Journal of Law & Technology* 17, no. 3 (2011), 11.
9. *Da Silva Moore v. Publicis Groupe & MSL Grp.*, 287 F.R.D. 182 (2012).
10. Goodman, *Robots in Law*, 30.
11. Lawgeex, *Comparing the Performance of Artificial Intelligence to Human Lawyers in the Review of Standard Business Contracts* (New York: Lawgeex, 2018), <https://images.law.com/contrib/content/uploads/documents/397/5408/lawgeex.pdf>.
12. John O. McGinnis and Russell G. Pearce, "The Great Disruption: How Machine Intelligence Will Transform the Role of Lawyers into the Delivery of Legal Services," *Fordham Law Review* 82, no. 6 (2014), 3041. See also Richard Suskind, *Tomorrow's Lawyers: An Introduction to Your Future* (2d ed., Oxford: Oxford University Press, 2017); Richard Suskind and Daniel Suskind, *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (Oxford: Oxford University Press, 2016); Sam Skolnik, "Artificial Intelligence Creeps into Big Law, Endangers Some Jobs," *Bloomberg Law Big Law Business*, January 22, 2019, <https://biglawbusiness.com/artificial-intelligence-creeps-into-big-law-endangers-some-jobs>.
13. Maura Grossman, Ralph Losey, and Geoffrey Kingsporn, "E-Discovery," *Exterro EdTalks*, May 18, 2018, <https://www.exterro.com/resources/edtalks-presentations-5/>.
14. Leslie Gordon, "How Attorneys Can Avoid Burnout and Debilitating Anxiety," *ABA Journal* 101, no. 7 (July 2015): 59.
15. "How Much Does Artificial Intelligence (AI) Cost in 2019?," *Azati Software Blog*, January 23, 2019, <https://azati.com/how-much-does-it-cost-to-utilize-machine-learning-artificial-intelligence/>.
16. Ellyn Shook and Mark Knickrehm, *Reworking the Revolution* (New York: Accenture, 2018), 3, https://www.accenture.com/t20180613T062119Z_w_/nl-en/_acnmedia/PDF-69/Accenture-Reworking-the-Revolution-Jan-2018-POV.pdf#zoom=50.

17. *Cass v. 1410088 Ontario Inc.*, 2018 ONSC 6959 (2018) (CanLII), 34 <http://canlii.ca/t/hw728>.
18. Geordie Clark, "Billable Hour 'Makes No Sense' in an AI World," *Raconteur*, November 27, 2018, <https://www.raconteur.net/risk-management/ai-analytics-billable-hour>.
19. Robert Ambrogi, "Tech Competence," *LawSites* (blog), <https://www.lawsitesblog.com/tech-competence>.
20. Jamie J. Baker, "Beyond the Information Age: The Duty of Technology Competence in the Algorithmic Age," *South Carolina Law Review* 69, no. 3 (2018): 557.
21. Jamie J. Baker, "Are Algorithms Required for Ethical Legal Research?," *The Ginger (Law) Librarian*, February 1, 2019, <http://www.gingerlawlibrarian.com/2019/02/are-algorithms-required-for-ethical.html>.
22. Matthew Turner, "Does a Doctor's Failure to Use AI to Diagnose a Patient Constitute Medical Malpractice?," *Sommers Schwartz Blog*, April 20, 2018, <https://www.sommerspc.com/blog/2018/04/does-a-doctors-failure-to-use-ai-to-diagnose-a-patient-constitute-medical-malpractice/>.
23. Legal philosophers may note that this observation suggests that law will increasingly resemble Justice Holmes's observation that law is whatever the courts are willing to enforce. Oliver Wendall Holmes, Jr., "The Path of Law," *Harvard Law Review* 10, no. 8 (1897): 457, 461.
24. Suzette Heald, "The Legacy of Isaac Schapera," *Anthropology Today* 19, no. 6 (Dec. 2003): 18.
25. Nizan Geslevich Packin and Yafit Lev Aretz, "Learning Algorithms and Discrimination," in *Research Handbook of Artificial Intelligence and Law*, eds. Woodrow Barfield and Ugo Pagallo (Northampton, MA: Edward Elgar, 2018), 88, 92. *State v. Loomis*, 881 N.W.2d 749 (Wis. 2016).
26. Art. 22: "The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her," <https://gdpr-info.eu/art-22-gdpr/>; Art 13(2)(f): "The controller shall, at the time when personal data are obtained, provide the data subject with the following further information necessary to ensure fair and transparent processing . . . the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject," <https://gdpr-info.eu/art-13-gdpr/>.
27. Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin, "How We Analyzed the COMPAS Recidivism Algorithm," *ProPublica*, May 23, 2016, <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>.
28. Virginia Eubanks, "A Child Abuse Prediction Model Fails Poor Families," *Wired*, January 15, 2018, <https://www.wired.com/story/excerpt-from-automating-inequality/>.

29. Brian X. Chen, "HP Investigates Claims of 'Racist' Computers," *Wired*, December 22, 2009, <https://www.wired.com/2009/12/hp-notebooks-racist/>; Sam Levin, "A Beauty Contest Was Judged by AI and the Robots Didn't Like Dark Skin," *The Guardian*, September 8, 2016, <https://www.theguardian.com/technology/2016/sep/08/artificial-intelligence-beauty-contest-doesnt-like-black-people>.
30. Amit Datta, Michael Carl Tschantz, and Anupam Datta, "Automated Experiments on Ad Privacy Settings," *Proceedings on Privacy Enhancing Technologies* 1 (2015), 92, <https://content.sciendo.com/view/journals/popets/2015/1/article-p92.xml>.
31. "New York Bulletin – AI Workshop + Do We Need an FDA for Algorithms," *Artificial Lawyer*, January 29, 2019, <https://www.artificiallawyer.com/2019/01/29/new-york-bulletin-ai-workshop-do-we-need-an-fda-for-algorithms/>.
32. H.R. 4625, 115th Cong. (2017).
33. New York City, Office of the Mayor, "Mayor de Blasio Announces First-In-Nation Task Force To Examine Automated Decision Systems Used by the City," News, May 16, 2018, <https://www1.nyc.gov/office-of-the-mayor/news/251-18/mayor-de-blasio-first-in-nation-task-force-examine-automated-decision-systems-used-by>.
34. Associated Press, "Robot Kills Worker at Volkswagen Plant in Germany," *The Guardian*, July 1, 2015, <https://www.theguardian.com/world/2015/jul/02/robot-kills-worker-at-volkswagen-plant-in-germany>; see also James Temperton, "When Robots Kill: Deaths by Machines Are Nothing New But AI Is About to Change Everything," *Wired*, March 17, 2017, <https://www.wired.co.uk/article/robot-death-wanda-holbrook-lawsuit>.
35. Troy Griggs and Daisuke Wakabayashi, "How a Self-Driving Uber Killed a Pedestrian in Arizona," *New York Times*, March 21, 2018, <https://www.nytimes.com/interactive/2018/03/20/us/self-driving-uber-pedestrian-killed.html?mtrref=www.google.com>.
36. Turner, *Robot Rules*, 64.
37. Turner, *Robot Rules*, 91.
38. Zara Stone, "Everything You Need to Know about Sophia, the World's First Robot Citizen," *Forbes*, November 7, 2017, <https://www.forbes.com/sites/zarastone/2017/11/07/everything-you-need-to-know-about-sophia-the-worlds-first-robot-citizen/#5da9179a46fa>.