

Artificial Intelligence in Adjudication and Administration:
A Status Report on Governmental Use of Technology in the United States

Cary Coglianese* & Lavi Ben Dor**

Artificial intelligence has begun to permeate many aspects of U.S. society.¹ In settings as varied as medicine, transportation, financial services, and entertainment, digital technologies continue to emerge that rely on machine-learning algorithms to process vast quantities of data and make highly accurate predictions that can often outperform human ability to perform similar tasks.² As a result, the potential utility of artificial intelligence in the legal field has not gone

* Edward B. Shils Professor of Law and Political Science, Director, Penn Program on Regulation, University of Pennsylvania Law School.

** J.D. Candidate, University of Pennsylvania Law School.

© 2019 All rights reserved. This report has been prepared as a contribution to the “Government by Algorithms” project based at Leiden University in the Netherlands. The authors gratefully acknowledge assistance from David Lehr with portions of the research underlying this report. The report should be considered a draft, subject to updating and correction. It reflects the authors’ understanding of the state of the art of machine-learning technology and its digital building blocks in the United States as of October 2019. Comments are welcome and should be directed to cary_coglianese@law.upenn.edu.

¹ While a variety of definitions for the term “artificial intelligence” exist, a helpful one is “[t]he theory and development of computer systems able to perform tasks normally requiring human intelligence.” *Artificial Intelligence*, OXFORD DICTIONARIES, <https://en.oxforddictionaries.com/definition/us/artificial-intelligence>. The terms “machine learning” and “artificial intelligence” are to some extent interchangeable and are used as such throughout this Essay. Cf. Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1, 2 n.2 (2019) (“By “artificial intelligence” and “machine learning,” we refer . . . to a broad approach to predictive analytics captured under various umbrella terms For our purposes, we need not parse differences in the meaning of these terms, nor will we delve deeply into specific techniques within machine learning.”).

² See, e.g., Cary Coglianese, *Using Machine Learning to Improve the U.S. Government*, THE REG. REV. (Aug. 12, 2019), <https://www.theregreview.org/2019/08/12/coglianese-using-machine-learning-to-improve-us-government/>; Peter Dizikes, *AI, the Law, and Our Future*, MIT NEWS OFFICE (Jan. 18, 2019), <http://news.mit.edu/2019/first-ai-policy-congress-0118>; Jillian D’Onfro, *AI 50: America’s Most Promising Artificial Intelligence Companies*, FORBES (Sept. 17, 2019), <https://www.forbes.com/sites/jilliandonfro/2019/09/17/ai-50-americas-most-promising-artificial-intelligence-companies/#54bfb84c565c>; Chris Weller, *A California Police Department is Using Software to Decide if You’re About to Commit a Crime* (Jan. 12, 2016), <https://www.businessinsider.com/intrado-beware-system-tracks-threat-levels-2016-1> (“A new piece of software in place at the Fresno Police Department in central California uses huge batches of data, ranging from criminal history to Twitter feeds, to assess how likely someone is to commit a crime and whether the police ought to keep tabs on them.”).

unnoticed, with scholars, attorneys, and judges beginning to examine the implications it could have for the United States legal system.³

This paper seeks to capture the state of the art of the current uses of digitization, algorithmic tools, and machine learning in domestic governance in the United States. It serves, in effect, as a status report on non-military governmental use—that is, functions by courts and administrative agencies—of artificial intelligence and its building blocks throughout the United States.⁴ That said, probably no single report in such a fast-changing area of technology can claim to be entirely complete. Moreover, the domestic governmental system in the United States is vast and fragmented. The U.S. Constitution created a federal framework for domestic governance. It authorizes the national (or federal) government to engage in specific functions and reserves for the states all other governing powers.⁵ As a result, the responsibility for domestic governance is divided across 51 governments—those of the 50 states, plus the national government—through a process at times described (with varying degrees of accuracy) as “cooperative” federalism.

Mere recognition of the U.S. government’s federal structure only scratches the surface of the obstacles confronting any effort to survey uses of specific technologies, as the development and implementation of technology is not necessarily determined at a central level even within any of the nation’s 51 main governments. Technology decisions are made at the federal level in as many as several hundreds of separate administrative agencies.⁶ The number of comparable

³ See, e.g., Daniel L. Chen, *Machine Learning and the Rule of Law* 4-7 (Working Paper, 2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3302507 (noting that machine learning may be useful in detecting and adjusting for bias in judicial decision-making on asylum requests).

⁴ Military and security intelligence-gathering uses would obviously be subject to security classification. For the most in-depth treatment of artificial intelligence in U.S. military applications, see PAUL SCHARRE, *ARMY OF NONE: AUTONOMOUS WEAPONS AND THE FUTURE OF WAR* (2018).

⁵ U.S. CONST. amend. X.

⁶ Indeed, just getting a count of the number of federal agencies is difficult. One scholarly report published by a governmental agency noted that “there is no authoritative list of government agencies. Every list of federal agencies

agencies at the state and local level surely runs into the tens of thousands, and apparently no one has ever tried to count them all. With respect simply to law enforcement agencies, it has been noted that “the decentralized, fragmented, and local nature of law enforcement in the United States makes it challenging to accurately count the number of agencies.”⁷ But, in 2008, approximately 18,000 different police departments and other law enforcement agencies responded to a federally sponsored Census of State and Local Law Enforcement Agencies.⁸

Similar numbers describe the judiciary in the United States. The federal court system comprises, in addition to one Supreme Court, a total of 13 “circuits” in the federal appellate court system and 94 trial court “districts” (each with dozens of trial judges that in total number over 650 courtrooms).⁹ At the state level, the number of different courts proliferates still further—especially given that state governments further delegate domestic authority to county and municipal governments. According to the National Center for State Courts, approximately 15,000-17,000 different state and municipal courts exist in the United States.¹⁰

Any one of these numerous judicial or administrative entities could in principle have its own policy with respect to electronic filing, digitization of documents, or the use of algorithms to support decision-making.¹¹ As a result, any generalizations about the use of artificial intelligence

in government publications is different.” DAVID E. LEWIS & JENNIFER L. SELIN, SOURCEBOOK OF UNITED STATES EXECUTIVE AGENCIES (1st ed. 2012), https://www.acus.gov/sites/default/files/documents/Sourcebook-2012-Final_12-Dec_Online.pdf (reporting estimates of the number of federal administrative agencies that range from 252 to 405).

⁷ U.S. Department of Justice, National Sources of Law Enforcement Employment Data 1 (Oct. 4, 2016), <https://www.bjs.gov/content/pub/pdf/nsleed.pdf>.

⁸ *Id.* at 6.

⁹ *See, e.g.*, U.S. Courts, Court Role and Structure, <https://www.uscourts.gov/about-federal-courts/court-role-and-structure>.

¹⁰ This estimate is based on a telephone and email exchange with NCSC staff, and it includes a vast number of municipal courts. Indeed, the uncertainty reflected in the range (rather than a point estimate) is apparently due to fairly regular changes in the size and organization of municipal courts.

¹¹ *See, e.g.*, U.S. SUPREME COURT, RULES OF THE SUPREME COURT (adopted Sept. 27, 2017), <https://www.supremecourt.gov/filingandrules/2017RulesoftheCourt.pdf>; FED. R. CIV. P. 83(a) (“[A] district court, acting by a majority of its district judges, may adopt and amend rules governing its practice.”); FED. R. APP. P.

or its building blocks must be made with appropriate caution. Neither time nor resources permitted a comprehensive survey of all of these governmental entities, at any level of government. Instead, this paper is based primarily on extensive searches of academic literature and media publications. We also spoke with court and agency officials and others who would be in a position to know about the current uses of artificial intelligence and its building blocks by governmental entities, including leading consultants and academic experts developing and studying such possible uses. This effort generated as thorough a survey of judicial and administrative uses of machine learning across federal and state governments as any of which we know.

The results of our research lead us to be quite confident in two overarching conclusions. First, no judicial or administrative body in the United States has instituted a system that provides for total decision-making by algorithm, such that a digital system makes a fully independent determination (that is, human “out of the loop” decision).¹² Second, we are also aware of no *court* that is currently relying in any way, even on a human-in-the-loop basis, on machine-learning algorithms for *judicial* decisions. That said, one state has a parole board using a system based on a machine-learning algorithm to support prisoner release decisions and numerous federal and state agencies have deployed or are researching the use of machine learning in support of various *administrative* functions.

Such machine-learning algorithms—which we treat here as defining “artificial intelligence”—are distinguished here from two building blocks that might help lead to the

47(a)(1) (“Each court of appeals acting by a majority of its judges in regular active service may . . . make and amend rules governing its practice.”)

¹² For a discussion of the difference between using algorithms on a supportive versus determinative basis, see Coglianesse & Lehr, *supra* note 1, at 31; Cary Coglianesse & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147, 1167-1170 (2017).

eventual governmental use of artificial intelligence: digitization and algorithmic tools. Indeed, machine learning resides on the far end of a spectrum of digital technologies available to governments. The closest point on that spectrum begins with simple *digitization*—or the use of electronic filing or other data systems to manage information in electronic format. Digitization is a building block toward artificial intelligence because it can facilitate the availability of the “Big Data” on which machine learning is based. Next on the spectrum would be for governments to rely on we will call here *algorithmic tools*—that is, traditional, human-created statistical models, indices, or scoring systems that are then used as decision tools. These traditional algorithmic or statistical tools rely on humans to select the specific variables to be included in a decision aid and the precise mathematical relationships between these variables. Only the final step on the spectrum—*machine learning*—constitutes what we will consider artificial intelligence because learning algorithms essentially work “on their own” to process data and discover optimal mathematical relationships between them. This autonomous self-discovery is what gives machine-learning algorithms not only their name but also their frequent superior performance in terms of accuracy over traditional algorithmic tools. Of course, even with machine learning, humans must specify the objective that the learning algorithm is supposed to forecast or optimize, and humans must undertake a number of steps to “train” the algorithm and refine its operation.¹³ Yet these learning algorithms are different than traditional statistical tools because the precise ways that data are combined and analyzed are neither determined in advance by a human analyst nor easily explainable after the fact. For this reason, machine-learning algorithms are often described as “black-box” algorithms because they do not afford a ready way of

¹³ For an excellent primer on machine learning, see David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 669 (2017).

characterizing how they work—other than that they can be quite accurate in achieving the objectives they have been designed to achieve.

In the rest of this paper, we first take up the status of artificial intelligence in the federal and state judiciaries. More precisely, we report on three building blocks that might eventually lead to the use of artificial intelligence in the courts: the increased digitization of court records, the use of algorithmic tools for risk assessment in aspects of the criminal justice process, and the growth of online dispute resolution outside of the courts. The most widespread innovation in the courts has occurred in various forms of digitization (such as electronic filing and case management), while some courts have relied on algorithmic tools to support pretrial, sentencing, or parole decisions. Some courts recognize a role for online dispute resolution systems developed by the private sector.

We turn, in the second part of this paper, to a review of administrative agencies' use of artificial intelligence. Many administrative systems have been digitized for some time, and administrative agencies have also long relied on traditional statistical analysis or algorithmic tools. But most relevant to the purposes of this paper, some administrative agencies at the local, state, and federal level are also starting to use machine learning for certain analytical and decision support purposes. We thus devote our attention in the second part of the paper to these latter uses of machine learning in the administrative context.

In both parts of this paper, we also highlight some of the legal issues, and at times the litigation and public controversy, that has surrounded certain applications of algorithmic tools or machine learning. Given the increased use of artificial intelligence in other facets of society, as well as in many other parts of the world, the path toward greater governmental reliance on machine learning in the United States will likely continue to move forward. At some point in the

not-too-distant future, autonomous decision-making systems based on machine learning may well begin to take the place of a government singularly and literally “of the people” and “by the people” in the United States.¹⁴

I. ARTIFICIAL INTELLIGENCE BUILDING BLOCKS IN THE COURTS

As of today, of course, we know of no machine-learning tool has been adopted in any court in the United States to make an ultimate, fully automated determination on a legal or factual question.¹⁵ However, several trends in recent years have emerged that could signal movement towards the eventual use of such automated adjudication via artificial intelligence. To date, the principal building blocks of artificial intelligence in the courts comprise the digitization of court filings and processes, the introduction of algorithmic tools for certain criminal court decisions, and the emergence of online dispute resolution as an alternative to traditional court proceedings for small claims.

A. Digitization of Court Records

Artificial intelligence depends on data.¹⁶ Increasingly, court systems in the United States have made data more easily accessible through the growing digitization of court documents.¹⁷

¹⁴ Abraham Lincoln, The Gettysburg Address (Nov. 19, 1863).

¹⁵ See Richard C. Kraus, *Artificial Intelligence Invades Appellate Practice: The Here, The Near, and The Oh My Dear*, AM. BAR ASS'N (Feb. 5, 2019), https://www.americanbar.org/groups/judicial/publications/appellate_issues/2019/winter/artificial-intelligence-invades-appellate-practice-the-here-the-near-and-the-oh-my-dear/ (noting that in the United States, “the more fantastic ideas such as using AI to objectively decide cases by analyzing facts and applying law . . . are still figments of creative imaginations”).

¹⁶ See Willem Sundblad, *Data Is the Foundation for Artificial Intelligence and Machine Learning*, FORBES (Oct. 18, 2018), <https://www.forbes.com/sites/willemsundbladeurope/2018/10/18/data-is-the-foundation-for-artificial-intelligence-and-machine-learning/#4bd8c64051b4> (“[D]ata is both the most underutilized asset of manufacturers and the foundational element that makes AI so powerful.”).

¹⁷ See, e.g., Jenni Bergal, *Courts Plunge into the Digital Age*, PEW CHARITABLE TRUSTS: STATELINE (Dec. 8, 2014), <https://www.pewtrusts.org/en/research-and-analysis/blogs/stateline/2014/12/8/courts-plunge-into-the-digital-age> (noting that the status of courthouses’ digital use “has been changing dramatically in many courthouses across the country. States are moving to systems in which documents are submitted electronically, file rooms are disappearing and the judicial system is going paperless”); *Records/Document Management Resource Guide*, NAT’L CTR. FOR STATE COURTS, <https://www.ncsc.org/Topics/Technology/Records-Documents-Management/Resource-Guide.aspx> (last updated Sept. 26, 2018) (“Records and document management are at the core of most courts’ business

This digitization has in large part been internally driven by the courts. Courts at both the state and federal level, including the Supreme Court itself, have required electronic filing as one of several ways a party can submit motions or arguments to a court, or as the only method of doing so.¹⁸ In addition, virtually every state and the federal government post free forms online that can be downloaded and used by litigants.¹⁹ Some courts have created “dedicated computer kiosks” specifically designed to help litigants who lack legal representation.²⁰ In California, for example, an “‘Online Self-Help Center’ offers PDFs that can be filled in online and used for evictions, divorces, orders of protection, collection matters, small claims, and other issues.”²¹

The federal judiciary has instituted a “comprehensive case management system” known as the Case Management/Electronic Case Files (CM/ECF) system that allows for convenient filing and organization of court documents, party pleadings, and other relevant materials.²² In 2002, Congress directed the federal courts to ensure that, with exceptions for certain documents filed under seal, “any document that is filed electronically [is also] publicly available online.”²³ State and local courts have increasingly rolled out various electronic filing (or “e-filing”)

processes. . . . [M]any state courts have implemented electronic court records (ECR) and electronic data management systems (EDMS) in an effort to improve court operations and manage unruly paperwork.”)

¹⁸ See, e.g., SUP. CT. R. 29 (requiring that in addition to filing documents with the Court Clerk, “all filers who are represented by counsel must submit documents to the [Supreme] Court’s electronic filing system”); 7TH CIR. R. 25 (“All documents must be filed and served electronically.”); E.D. PA. LOCAL R. 5.1.2 (“All civil and criminal cases filed in this court are required to be entered into the court’s Electronic Case Filing (“ECF”) System”); CA R. OF CT. 2.253 (empowering state courts in California to either permit or require parties to file electronically).

¹⁹ National Center for State Courts, Self Representation, <https://www.ncsc.org/Topics/Access-and-Fairness/Self-Representation/State-Links.aspx?cat=Court%20Forms>.

²⁰ BARTON & BIBAS, *supra* note __, at 123.

²¹ BENJAMIN H. BARTON & STEPHANOS BIBAS, REBOOTING JUSTICE 119 (2017). Barton and Bibas report that in a single year more than 4 million people visited the California self-help portal. They also report successful experiences with other systems for “DIY” lawyering, such as a system in New York State. *Id.* at 119-123.

²² *Case Management/Electronic Case Files*, PACER, <https://www.pacer.gov/cmecf/>. Public access to PACER data is not free, which has generated some controversy. See, e.g., David Post, *Yes, PACER Stinks... But Is It Also Overcharging Its Customers?*, WASH. POST (Jan. 9, 2016), https://www.washingtonpost.com/news/volokh-conspiracy/wp/2016/01/09/yes-pacer-stinks-but-is-it-also-overcharging-its-customers/?utm_term=.49cf19383d86. Similar concerns have been expressed related to the video recording of judicial proceedings. See, e.g., Jonathan Sherman, *End the Supreme Court’s Ban on Cameras*, N.Y. TIMES (Apr. 24, 2015), <https://www.nytimes.com/2015/04/24/opinion/open-the-supreme-court-to-cameras.html>.

²³ E-Government Act of 2002, 107 P.L. 347, 116 Stat. 2899, §205(c)(1).

software to replace paper submissions and docketing, many in the past decade.²⁴ In Florida alone, individuals filed roughly 23.5 million documents totaling about 110 million pages from mid-2018 to mid-2019.²⁵ These systems have created massive repositories of filings from litigants and judicial decisions and orders, all held in centralized databases.

In principle, artificial intelligence could take advantage of all of this data. At law firms, the increasing use of algorithmic tools, including those involving machine-learning algorithms, can be found to support the review of documents during the discovery process. This “e-discovery” practice has been shown to have a “strong impact” on reducing the need for human labor—plus it has spawned services that seek to analyze trends and make legal forecasts.²⁶ In addition, artificial intelligence has been used by outside researchers to attempt to predict courts’ decisions using data. In a 2017 study, a machine-learning statistical model correctly predicted the outcome of 70% of 28,000 U.S. Supreme Court decisions and 72% of individual justices’ votes from 1816 to 2015.²⁷ With a growing amount of data available from courts at all levels across the country, it is likely that such predictive efforts will only improve in quality in the future. In time, it may also be possible that artificial intelligence tools will have gained enough “experience” in

²⁴ *Electronic Filing*, NAT’L CTR. FOR STATE COURTS, <https://www.ncsc.org/Topics/Technology/Electronic-Filing/State-Links.aspx>. For examples of state court “e-filing” systems, see *eCourts*, N.J. COURTS, <https://njcourts.gov/attorneys/ecourts.html>; *Electronic Filing in the Delaware Judiciary*, DEL. COURTS, <https://courts.delaware.gov/efiling/>; *Active Courts*, ODYSSEY EFILEGA, <http://www.odysseyefilega.com/active-courts.htm>; EFILETEXAS, <https://www.efiletexas.gov/>; *Superior Court Electronic Case Filing*, N.H. JUDICIAL BRANCH, <https://www.courts.state.nh.us/nh-e-court-project/superior-attorneys.htm> (noting that e-filing in the New Hampshire Superior Court became mandatory in September 2018).

²⁵ *2018-2019 Annual Statistics*, FL. COURTS E-FILING PORTAL, http://archive.flclerks.com/e-Filing_Authority/Resources/2018-2019_Board_Meetings/June_25_2019_Board_Meeting/Annual_Statistics_2018-2019May.pdf.

²⁶ See, e.g., Dana Remus & Frank Levy, *Can Robots Be Lawyers: Computers, Lawyers, and the Practice of Law*, 30 GEO. J. LEGAL ETHICS 501, 515-16 (2017). Various private sector efforts are underway to make use of court data for predictive analytic purposes. One major service is Lex Machina, which is used by law firms. <https://lexmachina.com/> Another service, Docket Navigator, performs some basic analytics (albeit not with machine learning) for intellectual property cases. <http://brochure.docketnavigator.com/>.

²⁷ Matthew Hutson, *Artificial Intelligence Prevails at Predicting Supreme Court Decisions*, SCIENCE (May 2, 2017), <https://www.sciencemag.org/news/2017/05/artificial-intelligence-prevails-predicting-supreme-court-decisions>.

document review to step into the role of the judges and, rather than just predicting their behavior, use the large troves of data available in electronic filing systems to help in making actual judicial determinations.

B. Risk Assessment Algorithms

Algorithmic tools have taken root in some court systems as an aid to judicial decision-making in criminal cases on questions of bail, sentencing, and parole. An algorithmic tool for bail (pre-trial) decisions, developed by the Arnold Foundation, has been adopted by four states (Arizona, Kentucky, New Jersey, and Utah) and about a dozen municipal courts, largely in major metropolitan areas.²⁸ About 20 states have reportedly adopted some kind of algorithmic risk assessment tool in sentencing decisions.²⁹ More than half of the states use some form of algorithmic tool for purposes of parole decision-making.³⁰ The federal government has recently announced an algorithmic tool for parole decisions: Prisoner Assessment Tool Targeting Estimated Risk and Needs (PATTERN).³¹ The PATTERN system was developed in response to federal legislation calling for the use of risk assessment in federal parole decisions. Similarly, some state statutes encourage or require the use of these algorithmic tools,³² while others are

²⁸ Arnold Ventures, Public Safety Assessment FAQs (“PSA 101”) (March 2019),

https://craftmediabucket.s3.amazonaws.com/uploads/Public-Safety-Assessment-101_190319_140124.pdf.

²⁹ Sonja Starr, *Evidence-Based Sentencing and the Scientific Rationalization of Discrimination*, 66 STAN. L. REV. 803, 809 (2014). Federal courts also must consider the Sentencing Guidelines, which set out suggested sentence ranges for federal offenses depending on a variety of factors. *See* U.S. SENTENCING GUIDELINES MANUAL (U.S. SENTENCING COMM’N 2018); *see also* *Kimbrough v. United States*, 552 U.S. 85, 108-09 (2007) (noting that the Sentencing Guidelines are advisory but nonetheless should play a “key role” in judges’ considerations).

³⁰ BERNARD E. HARCOURT, *AGAINST PREDICTION: PROFILING, POLICING, AND PUNISHING IN AN ACTUARIAL AGE* 77 (2007) (noting twenty-eight states were using an algorithmic risk assessment tool for parole decision-making as of 2004).

³¹ U.S. Department of Justice, *The First Step Act of 2018: Risk and Needs Assessment System* (2019), https://nij.ojp.gov/sites/g/files/xyckuh171/files/media/document/the-first-step-act-of-2018-risk-and-needs-assessment-system_1.pdf.

³² *See, e.g.*, ALA. STAT. § 12-25-33(6) (instructing the Alabama Sentencing Commission to develop a risk assessment instrument to be “predictive of the relative risk that a felon will become a threat to public safety”); 42 PA. CONS. STAT. § 2154.7 (2019) (requiring the Pennsylvania Commission on Sentencing to adopt a risk assessment tool to “be used as an aide in evaluating the relative risk that an offender will reoffend and be a threat to public safety”); OKLA. STAT. tit. 22, § 988.18 (requiring courts to use a risk assessment tool in determining an offender’s

selected at the discretion of state or local officials.³³ As best we can determine, only one jurisdiction (Pennsylvania) has implemented any risk assessment tool in criminal justice that is based on machine learning.³⁴ Despite somewhat frequent claims to the contrary in the popular media,³⁵ the remaining algorithmic tools are based on standard indices or conventional logistic regression models—not machine-learning algorithms.

For example, one of the more popular non-learning algorithmic tools for bail decisions, the Arnold Foundation’s Public Safety Assessment, considers nine factors: the defendant’s age, the current violent offense, pending charges at the time of the offense, prior misdemeanor, felony, and violent convictions, prior failure to appear in the past two years and prior to the past two years, and prior sentences to incarceration. It then weighs these factors in varying proportions to determine scores from one to six that predict the defendant’s likelihood of failure to appear in court, new criminal activity, and new violent criminal activity, which a court can use in determining whether to grant a defendant pretrial release.³⁶

eligibility for a sentence of community service); OHIO REV. CODE ANN. § 5120.114(A) (LexisNexis 2019) (allowing for the use of a risk assessment tool by a variety of adjudicatory bodies in the criminal justice system); KY. REV. STAT. ANN. § 532.007(3)(a) (LexisNexis 2019) (“Sentencing judges shall consider . . . the results of a defendant’s risk and needs assessment included in the presentence investigation”); W. VA. CODE § 62-12-6(a)(2) (West 2019); *see also* ARIZ. CODE OF JUDICIAL ADMIN. § 6-201.01(J)(3) (2016) (“For all probation eligible cases, presentence reports shall [] contain case information related to criminogenic risk and needs as documented by the standardized risk assessment and other file and collateral information.”).

³³ *See, e.g.*, BD. OF DIRS. OF THE JUDICIAL CONFERENCE OF IND., POLICY FOR INDIANA RISK ASSESSMENT SYSTEM (Apr. 23, 2010), <https://www.in.gov/judiciary/cadp/files/prob-risk-iras-2012.pdf> (last updated Sept. 14, 2012); R.I. DEP’T OF CORRS., LEVEL OF SERVICE INVENTORY-REVISED: A PORTRAIT OF RIDOC OFFENDERS (Apr. 2011), <http://www.doc.ri.gov/administration/planning/docs/LSINewsletterFINAL.pdf>.

³⁴ Richard Berk, *An Impact Assessment of Machine Learning Risk Forecasts on Parole Board Decisions and Recidivism*, 13 J. EXPT. CRIM. 193 (2017). And technically this use in Pennsylvania is by an agency, not a court: the Pennsylvania Board of Probation and Parole. Another state, Maryland, has apparently looked into using machine learning for parole, but does not appear to have implemented it.

³⁵ *See, e.g.*, Matt O’Brien & Dake Kang, *AI in the Court: When Algorithms Rule on Jail Time*, ASSOCIATED PRESS (Jan. 31, 2018), <https://www.apnews.com/ae7b23e20c874800aa5746b92210a2dc>.

³⁶ *Risk Factors and Formulas*, PUBLIC SAFETY ASSESSMENT, <https://www.psapretrial.org/about/>. *See also, e.g.*, N.J. COURTS, PUBLIC SAFETY ASSESSMENT: NEW JERSEY RISK FACTOR DEFINITIONS 1-4 (Dec. 2018), <https://njcourts.gov/courts/assets/criminal/psariskfactor.pdf>.

Another non-learning algorithmic tool, known as Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), has been adopted by several state court systems for pretrial decisions. It involves an extensive questionnaire examining issues such as the defendant's prior criminal history, compliance with probation, substance abuse, relationships with others who have been arrested or sent to jail, home and work environment, and personality.³⁷ The algorithm uses these data points to place the defendant along several "risk scales" purporting to predict the defendant's relative likelihood of pretrial failure (including failure to appear and new felony arrest after pretrial release) and recidivism.³⁸ Judges deciding whether to approve a defendant for pretrial release or analyzing the appropriate sentence to set can then take the values reached by these algorithms into account in their determinations.³⁹

A third basic algorithmic tool, LSI-R (Level of Service Inventory-Revised), also aims to predict a defendant's risk of recidivism by weighing a number of factors. These factors include criminal history, educational and employment background, financial, mental, and familial state, substance abuse, and other personal details.⁴⁰

³⁷ See generally NORTHPOINTE INC., PRACTITIONER'S GUIDE TO COMPAS CORE (Mar. 19, 2015), http://www.northpointeinc.com/files/technical_documents/Practitioners-Guide-COMPAS-Core-_031915.pdf. While the details of the algorithm are proprietary, investigative journalistic organization ProPublica uncovered the version of the questionnaire used in the Wisconsin state court system. See *Risk Assessment*, <https://www.documentcloud.org/documents/2702103-Sample-Risk-Assessment-COMPAS-CORE.html>. Other states' courts have adopted COMPAS as well, including Florida, Michigan, New Mexico, and Wyoming. See *Algorithms in the Criminal Justice System*, *supra*.

³⁸ PRACTITIONER'S GUIDE TO COMPAS CORE, *supra*.

³⁹ See, e.g., Liptak, *supra*; O'Brien & Kang, *supra*; Ellora Thadaney Israni, *When an Algorithm Helps Send You to Prison*, N.Y. TIMES (Oct. 26, 2017), <https://www.nytimes.com/2017/10/26/opinion/algorithm-compas-sentencing-bias.html> ("Use of a computerized risk assessment tool somewhere in the criminal justice process is widespread across the United States . . . States trust that even if they cannot themselves unpack proprietary algorithms, computers will be less biased than even the most well-meaning humans.").

⁴⁰ For instance, the Rhode Island Department of Corrections has adopted this test. See R.I. DEP'T OF CORR., *supra*. Courts in other states have adopted some version of the LSI-R as well, including California, Colorado, Delaware, Hawaii, Iowa, Oklahoma, and Washington. See *Algorithms in the Criminal Justice System*, *supra*.

In addition to these three examples, some states have also adopted their own unique risk assessment algorithms.⁴¹ But yet again, simply to be clear, these are not artificial intelligence *per se*: these statistical models develop formulas based off of studying large data sets, and then apply those formulas to the inputs they are given for each defendant, rather than engaging in fully unguided “learning” to figure out what scores to give defendants.

Despite their increasing use by criminal courts, algorithmic risk assessment tools have not avoided scrutiny. Some scholars, lawyers, and concerned citizens challenge the lack of transparency behind some of these algorithms, as some of them are created by private consultants who claim commercial secrecy protection to avoid disclosure. Even when the parameters used in the analysis are publicly known, the owners of the risk assessment system will often decline to explain how exactly the factors that go into assessing an individual’s likelihood of recidivism or pretrial misbehavior are weighted.⁴² As Judge Noel Hillman of the District of New Jersey has put it, “[a] predictive recidivism score may emerge oracle-like from an often-proprietary black box. Many, if not most, defendants . . . will lack the resources, time,

⁴¹ See, e.g., BD. OF DIRS. OF THE JUDICIAL CONFERENCE OF IND., *supra*; Susan Turner et al., *Development of the California Static Risk Assessment (CSRA): Recidivism Risk Prediction in the California Department of Corrections and Rehabilitation 5-6* (U.C. Irvine Ctr. for Evidence-Based Corrections, Working Paper, Sept. 2013), <https://ucicorrections.seweb.uci.edu/files/2013/12/Development-of-the-CSRA-Recidivism-Risk-Prediction-in-the-CDCR.pdf> (listing factors considered by the California Static Risk Assessment tool); LA. SENTENCING COMM’N, RECOMMENDATIONS OF THE LOUISIANA SENTENCING COMMISSION FOR THE 2010 AND 2011 TERMS 14 (Mar. 1, 2012), http://www.lcle.state.la.us/sentencing_commission/2012_biannual_report_lsc_final.pdf (noting that Louisiana utilizes risk assessment tools for “both inmate management and programming . . . for persons held in state adult correctional facilities and supervision planning . . . for persons under probation or parole supervision provided by the Department”); THE COUNCIL OF STATE GOV’TS JUSTICE CTR., MONTANA COMMISSION ON SENTENCING 41 (Mar. 1-2, 2016), https://csgjusticecenter.org/wp-content/uploads/2016/03/Montana_Commission_on_Sentencing_third_meeting.pdf (noting the use of the Montana Offender Reentry Risk Assessment).

⁴² See, e.g., Noel L. Hillman, *The Use of Artificial Intelligence in Gauging the Risk of Recidivism*, AM. BAR ASS’N (Jan. 1, 2019), https://www.americanbar.org/groups/judicial/publications/judges_journal/2019/winter/the-use-artificial-intelligence-gauging-risk-recidivism/ (“[P]redictive technology becomes another witness against the defendant without a concomitant opportunity to test the data, assumptions, and even prejudices that underlie the conclusion.”); Israni, *supra* (“No one knows exactly how COMPAS works; its manufacturer refuses to disclose the proprietary algorithm. We only know the final risk assessment score it spits out . . . Something about this story is fundamentally wrong: Why are we allowing a computer program, into which no one in the criminal justice system has any insight, to play a role in sending a man to prison?”).

and technical knowledge to understand, probe, and challenge” the use of these tools.⁴³ In addition, a widely discussed 2016 ProPublica investigation suggested that the COMPAS tool systematically found black defendants to be at a higher risk of recidivism than similarly situated white defendants,⁴⁴ raising significant questions about the wisdom of integrating algorithms into judicial decision-making.⁴⁵

To date, the courts have only started to grapple with the legal implications of these algorithmic tools. Most prominently, in *State v. Loomis*, a defendant in Wisconsin state court challenged the state’s use of the COMPAS algorithm in sentencing him after he plead guilty.⁴⁶ Loomis’s COMPAS risk scores indicated that he had a high risk of recidivism; at sentencing, the court relied in part on the fact that he had been “identified, through the COMPAS assessment, as an individual who is at high risk to the community.”⁴⁷

In a post-conviction challenge to his sentence, Loomis argued that using the risk assessment violated his due process rights (1) to be sentenced based upon accurate information; (2) to receive an individualized sentence; and (3) to not have gendered assessments be used in sentencing.⁴⁸ The trial court denied the motion, holding that it had “used the COMPAS risk assessment to corroborate its findings and that it would have imposed the same sentence

⁴³ Hillman, *supra*.

⁴⁴ See Julia Angwin et al., *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>; see also Ed Yong, *A Popular Algorithm Is No Better at Predicting Crimes Than Random People*, THE ATLANTIC (Jan, 17, 2018), <https://www.theatlantic.com/technology/archive/2018/01/equivant-compas-algorithm/550646/>; Israni, *supra*; Liptak, *supra*.

⁴⁵ See, e.g., Cynthia Rudin et al., *The Age of Secrecy and Unfairness in Recidivism Prediction* (Duke University, Working Paper, 2019), at 1, <https://arxiv.org/pdf/1811.00731.pdf> (“[W]e show that COMPAS does not necessarily depend on race, contradicting ProPublica’s analysis”); Anne L. Washington, *How to Argue with an Algorithm: Lessons from the COMPAS-ProPublica Debate* (New York University, Working Paper, 2019), at 2-3, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3357874 (providing an alternative framework for evaluating the integrity of predictive algorithms).

⁴⁶ 881 N.W.2d 749 (Wis. 2016).

⁴⁷ *Id.* at 755.

⁴⁸ *Id.* at 757.

regardless of whether it considered the COMPAS risk scores,” and the Wisconsin Supreme Court affirmed.⁴⁹

The Supreme Court rejected Loomis’s due process challenges, noting that the variables that the COMPAS algorithms used were publicly available and that the risk assessment’s outcome was based fully on either the defendant’s answers to the questions or on publicly available information about his criminal history.⁵⁰ As a result, the use of COMPAS complied with due process, since the defendant had the “opportunity to verify that the questions and answers listed on the report were accurate.”⁵¹ The court further held that, although the use of risk assessment tools did involve group data, its inclusion among a mix of factors still achieved an individualized sentence for the defendant.⁵² Finally, the inclusion of gender in the COMPAS algorithm’s analysis did not violate any due process rights absent any proof that the court actually relied on gender as a factor in sentencing, since the algorithm simply accounted for differences in recidivism rates between men and women.⁵³

Loomis appealed to the United States Supreme Court.⁵⁴ The Court invited the Solicitor General⁵⁵ to weigh in, often a sign that the Court recognizes the significant potential value of the case.⁵⁶ The Solicitor General’s Office argued that the Court should not grant the petition, noting

⁴⁹ *Id.*

⁵⁰ *Id.* at 761.

⁵¹ *Id.*

⁵² *Id.* at 764-65. However, this state Supreme Court warned lower courts to be careful given the group-based nature of the COMPAS assessment.

⁵³ *Id.* at 765-67.

⁵⁴ Petition for Writ of Certiorari, *Loomis v. Wisconsin*, 2017 U.S. LEXIS 4204 (June 26, 2017) (No. 16-6387).

⁵⁵ The Solicitor General handles all litigation on behalf of the United States in the U.S. Supreme Court. See *Office of the Solicitor General*, U.S. DEP’T OF JUSTICE, <https://www.justice.gov/osg>.

⁵⁶ See *Loomis v. Wisconsin*, 137 S. Ct. 1240 (2017). For discussions of the role of the Solicitor General in influencing the Court’s docket and merits decision, see, for example, Ryan C. Black & Ryan J. Owens, *Solicitor General Influence and Agenda Setting on the U.S. Supreme Court*, 64 POL. RES. Q. 765, 766 (2011) (“[W]e find strong support for SG influence. Justices who completely disagree with the SG nevertheless follow her recommendations 35 percent of the time, a number we take to be powerful evidence of influence.”); Margaret Meriwether Cordray & Richard Cordray, *The Solicitor General’s Changing Role in Supreme Court Litigation*, 51 B.C. L. REV. 1323, 1324 (2010) (“The U.S. Solicitor General, as the U.S. Supreme Court’s premier advocate, has

that no division of authority yet existed on the validity of the use of these algorithms and asserting that “[t]he issues that this petition raises . . . would benefit from further percolation. Most of the developments related to the use of actuarial risk assessments at sentencing have emerged within the last several years.”⁵⁷ Ultimately, the Court declined to take up the case, leaving the issue unresolved.⁵⁸

In another case, *Malenchik v. State*, the defendant pled guilty to a felony and admitted to being a habitual offender.⁵⁹ On appeal, the defendant challenged the use of the results of two risk assessment tests (one of which was the LSI-R) in determining his sentence; the tests’ results indicated that he was at high risk of recidivism.⁶⁰ The Indiana Supreme Court emphasized that the sentence had been based on factors other than the risk assessments, since the trial court had also relied on the defendant’s prior criminal history and refusal to accept responsibility for his actions and change his behavior and had not used the algorithm’s output as an independent aggravating factor.⁶¹ The court noted that such tools are neither “intended nor recommended to substitute for the judicial function of determining the length of sentence,” but are instead “significant sources of valuable information for judicial consideration in deciding whether to suspend all or part of a sentence, how to design a probation program for the offender, whether to assign an offender to alternative treatment facilities or programs, and other such corollary sentencing matters.”⁶² As a result, the Indiana Supreme Court held that a trial court can properly “supplement and enhance” its evaluation of the evidence before it at sentencing by considering

long exerted significant influence over both the Court’s case selection decisions and its substantive decisions on the merits.”).

⁵⁷ Brief for the United States as Amicus Curiae at 21-22, *Loomis v. Wisconsin*, 2017 U.S. LEXIS 4204 (June 26, 2017) (No. 16-6387).

⁵⁸ See *Loomis v. Wisconsin*, 2017 U.S. LEXIS 4204 (June 26, 2017).

⁵⁹ 928 N.E.2d 564, 566 (Ind. 2010).

⁶⁰ *Id.* at 566-67.

⁶¹ *Id.* at 568.

⁶² *Id.* at 573.

the results of a risk assessment, which can “provide usable information based on extensive penal and sociological research to assist the trial judge in crafting individualized sentencing schemes with a maximum potential for reformation.”⁶³

Meanwhile, a third case, *State v. Rogers*, involved a claim by a convicted defendant challenging the *lack of* the use of a risk assessment tool in his sentencing. The Supreme Court of Appeals of West Virginia rejected the claim because the defendant failed to enter a proper objection at the time of initial sentencing. But Justice Loughry, in a separate concurring opinion, argued that a risk assessment algorithm is “merely a tool that may be used by [trial court] judges during sentencing,” a process over which judges have broad discretion and that courts are under no obligation to use it.⁶⁴

Although it is still early in the judiciary’s assessment of legal issues surrounding courts’ use of algorithmic tools, it seems noteworthy that, in all three cases decided to date, courts reviewing algorithmic tools appear to have taken pains to emphasize that such tools only serve as one of multiple factors that a human judge takes into account in reaching a decision. Perhaps this suggests that, as long as humans remain in the loop, whether with standard algorithmic tools or even with machine learning, their use by courts will continue to receive approval.⁶⁵

C. Online Dispute Resolution

Online Dispute Resolution (ODR) has arisen in recent years as a tool for resolving disagreements among parties using technology, growing in part out of prior developments in the

⁶³ *Id.* at 573-75.

⁶⁴ No. 14-0373, 2015 WL 869323, at *2 (W. Va. Jan. 9, 2015); *id.* at *4-5 (Loughry, J., concurring).

⁶⁵ *See, e.g., Malenchik*, 928 N.E.2d at 568 (“[T]he trial court’s sentencing decision was clearly based on factors apart from the defendant’s LSI-R and SASSI results. . . . The trial judge did not rely on either the LSI-R or SASSI as an independent aggravating factor in deciding to impose more than the advisory sentence.”). *See generally* Melissa Hamilton, *Risk-Needs Assessment: Constitutional and Ethical Challenges*, AM. CRIM. L. REV. 231 (2015); Roger K. Warren, *Evidence-Based Sentencing: The Application of Principles of Evidence-Based Practice to State Sentencing Practice and Policy*, 43 U.S.F. L. REV. 585 (2009).

field of Alternative Dispute Resolution, which refers to a range of methods such as mediation and arbitration that aim to settle disputes without the use of litigation and the court system.⁶⁶

ODR mechanisms first mimicked ADR approaches to conflict resolution before evolving to their current forms, which harness the advantages of technology to aid their mission.⁶⁷

The initial growth of ODR was largely driven by the private sector.⁶⁸ Most notably, eBay and PayPal developed ODR systems to handle the millions of disputes that regularly arise on their platform from and among users.⁶⁹ Realizing that they could not afford to hire enough human mediators to resolve all of these disputes or arrange for parties to videoconference with each other, these companies leveraged the extensive amounts of data they had collected on consumer behavior and usage.⁷⁰ They parsed the dispute process into steps; developed automated software that analyzed the situation and suggested easy resolutions at the individual steps and facilitated interaction among the parties; and documented every aspect of the process for future analysis.⁷¹ As the success of these systems inspired other firms to develop similar and increasingly sophisticated programs, algorithms became a more prominent part of the solution, allowing companies to automate away many (if not all) of the steps of the dispute resolution and decision-making process.⁷² For example, Amazon has developed algorithms that can resolve a consumer complaint about a defective product without requiring any human intervention.⁷³

⁶⁶ See ETHAN KATSH & ORNA RABINOVICH-EINY, DIGITAL JUSTICE 33-34 (2017); *Online Dispute Resolution Moves From E-Commerce to the Courts*, PEW CHARITABLE TRUSTS (June 4, 2019), <https://www.pewtrusts.org/en/research-and-analysis/articles/2019/06/04/online-dispute-resolution-moves-from-e-commerce-to-the-courts> [hereinafter *Online Dispute Resolution I*]; *Alternative Dispute Resolution*, LEGAL INFORMATION INSTITUTE, https://www.law.cornell.edu/wex/alternative_dispute_resolution (last updated June 8, 2017).

⁶⁷ See KATSH & ORNA RABINOVICH-EINY, *supra*; *Online Dispute Resolution I*, *supra*.

⁶⁸ See *Online Dispute Resolution I*, *supra*.

⁶⁹ See BENJAMIN H. BARTON & STEPHANOS BIBAS, REBOOTING JUSTICE 111 (2017); KATSH & ORNA RABINOVICH-EINY, *supra*, at 34-35.

⁷⁰ BARTON & BIBAS, *supra*, at 111; KATSH & ORNA RABINOVICH-EINY, *supra*, at 34-35.

⁷¹ BARTON & BIBAS, *supra*, at 111-115; KATSH & ORNA RABINOVICH-EINY, *supra*, at 34-36.

⁷² KATSH & ORNA RABINOVICH-EINY, *supra*, at 46-48.

⁷³ *Id.* at 48.

Some courts have also begun experimenting with ODR as a mechanism to attempt to resolve lawsuits without requiring the use of judicial decision-making. Although much of the innovation in this area occurred in other parts of the world, dozens of state and local courts in the United States, including in Michigan, Ohio, California, and Utah, have adopted some form of “court ODR” in cases involving small claim civil matters, traffic violations, outstanding warrant cases, and low-conflict family court cases.⁷⁴ What counts as an ODR system can vary from a simple website that facilitates entering pleas for traffic tickets online to an online portal for engaging in asynchronous negotiations. These are not mandatory systems in any jurisdiction of which we are aware, but instead they are offered as an option to avoid appearing in court. In jurisdictions with these systems, parties are notified of the ODR option via mailings or websites.⁷⁵ Parties can access the ODR system at any time, and with the more interactive systems they can communicate and negotiate with each other, obtain legal information and suggested resolutions from the system, and easily manage electronic documents—all without having to see the inside of a courtroom.⁷⁶ These systems can usually reach resolution in a dispute faster and at lower cost to the parties and are far more accessible than traditional court-centered adjudication.⁷⁷

ODR provides an emerging avenue for litigants and courts to engage in dispute resolution outside of the presence of a courtroom and absent a human judge. Court ODR systems, as well as the private-sector iterations that inspired them, have increasingly adopted automated processes

⁷⁴ KATSH & ORNA RABINOVICH-EINY, *supra*, at 161-62; *Online Dispute Resolution Offers a New Way to Access Local Courts*, PEW CHARITABLE TRUSTS (Jan. 4, 2019), <https://www.pewtrusts.org/en/research-and-analysis/factsheets/2019/01/online-dispute-resolution-offers-a-new-way-to-access-local-courts> [hereinafter *Online Dispute Resolution II*].

⁷⁵ KATSH & ORNA RABINOVICH-EINY, *supra*, at 161-62; *Online Dispute Resolution II*, *supra*.

⁷⁶ *Online Dispute Resolution II*, *supra*.

⁷⁷ *Online Dispute Resolution I*, *supra*.

and rely on algorithmic tools to aid in reaching what some observers characterize as fair and low-cost solutions to the parties' disputes.⁷⁸ As some researchers have already begun to note, court systems could take these algorithms to the next "level" of autonomy by integrating artificial intelligence into ODR processes, allowing for increasingly automated forms of decision-making.⁷⁹

II. ARTIFICIAL INTELLIGENCE IN THE ADMINISTRATIVE STATE

Government agencies have long pursued the use of information technology to support vital services and programs. Even outside of the military, intelligence-gathering, and space exploration contexts, computers have been used for decades by government agencies to support administration and data management for tax collection and large national benefits programs such as Social Security and Medicare.⁸⁰ The technologies used by government have tended to lag behind those deployed in the private sector. Federal and state agencies relied on mainframe computers, for example, long after the personal computer revolution hit the private sector in the 1980s, and they continue to remain behind the innovation curve to this day.⁸¹ Many government computer systems have grown quite antiquated. As of 2016, auditors could report that three-quarters of annual federal spending on computer technology in the United States is devoted to

⁷⁸ KATSH & ORNA RABINOVICH-EINY, *supra*, at 163 ("The use of ODR in courts is also introducing algorithms into the judicial decision-making process."); Loïc E. Coutelier, *The New Frontier of Online Dispute Resolution: Online Divorce Mediation*, AM. BAR ASS'N (Apr. 1, 2016), https://www.americanbar.org/groups/young_lawyers/publications/tyl/topics/dispute-resolution/new-frontier-online-dispute-resolution-online-divorce-mediation/ (discussing a form of ODR used in divorce mediation that relies on "an innovative algorithm that uses game theory negotiation to maximize the return for divorcing couples who are dividing assets").

⁷⁹ See generally Arno R. Lodder & John Zeleznikow, *Artificial Intelligence and Online Dispute Resolution*, in ONLINE DISPUTE RESOLUTION: THEORY AND PRACTICE 73-94 (Mohamed S. Abdel Wahab et al. eds., 2012).

⁸⁰ Harold C. Relyea & Henry B. Hogue, *A Brief History of the Emergence of Digital Government in the United States*, in DIGITAL GOVERNMENT 16 (Alexei Pavlichev & G. David Garson eds., 2004).

⁸¹ Tod Newcombe, *The Complicated History of Government Technology*, GOVERNING (Oct. 2, 2017), <https://www.govtech.com/computing/The-Complicated-History-of-Government-Technology.html>; Jack Moore, *The Crisis in Federal IT that's Scarier than Y2K Ever Was*, NEXTGOV (Nov. 20, 2015), <http://www.nextgov.com/cio-briefing/2015/11/crisis-federal-it-rivals-y2k/123908/>.

“legacy systems” which “are becoming increasingly obsolete” due to “outdated software languages and hardware parts that are unsupported.”⁸²

Still, the internet revolution in the 1990s did prompt state and federal government agencies to begin to digitize many of their services and make greater use of the worldwide web. Initially, of course, the movement was slow. According to one survey, by the year 2000 states had websites containing an average each of only about four automated or online governmental services.⁸³ The most popular such digitized services at that time: applying for a state government job (32 states).⁸⁴ The second most popular was electronic filing of income taxes (24 states) and the third most popular was the online renewal of drivers licenses (17 states).⁸⁵ Today, all states have these basic services digitized—and many more services as well.

The federal government adopted the E-Government Act of 2002 “to develop and promote electronic Government services and processes” and “[t]o promote use of the Internet and other information technologies to provide increased opportunities for citizen participation in Government.” That law established a federal Office of Electronic Government, imposed a duty on all federal agencies to make vast quantities of government information available online, and generally to accept online submissions of public comments on proposed regulations. The federal government has created portals such as *Regulations.gov* and *Data.gov* to make available online massive amounts of information previously housed in paper records or internal government computers.⁸⁶

⁸² U.S. Government Accountability Office, *Federal Agencies Need to Address Aging Legacy Systems* (2016) (testimony of David A. Powner, Director, Information Technology Management Issues), <https://www.gao.gov/assets/680/677454.pdf>.

⁸³ Jane E. Fountain, *The Virtual State: Transforming American Government?*, 90 NATL. CIV. REV. 241, 242 (2001).

⁸⁴ *Id.*

⁸⁵ *Id.*

⁸⁶ For discussion of these and related efforts, see *The Obama Administration's Commitment to Open Government: A Status Report* (2011), https://obamawhitehouse.archives.gov/sites/default/files/opengov_report.pdf

Today, the United States is regarded as among the nations that has made the most advanced progress in implementing practices of e-government. According to the latest e-government ranking by the United Nations, the United States places second among all countries for “online service delivery” (tying with Singapore, and just behind Denmark).⁸⁷ It also ranks among the top nations for “e-participation.”⁸⁸

These rankings suggest that, even if administrative agencies in the United States may have been slower out of the starting gate than the private sector in their use of information technology, they appear ahead of many counterpart governments elsewhere in the world. They have also moved much earlier to digitize their operations and services than has the U.S. court system. In this respect, administrative agencies are well along a path that will support greater use of machine learning.

Some agencies have undertaken focused efforts to make data more easily accessible for use in machine-learning applications. For example, officials at the Federal Deposit Insurance Corporation have expressly focused on developing “the back-end disciplines of in-memory analytics, big data, and data quality.”⁸⁹ Staff at the Federal Communications Commission (FCC) established a Data Innovation Initiative with similar goals.⁹⁰ Financial regulators have worked to create a dedicated “legal entity identifier” to be able to link disparate transactional and other data to the corresponding business entities.⁹¹ The Environmental Protection Agency has built

⁸⁷ UNITED NATIONS, UNITED NATIONS E-GOVERNMENT SURVEY 2018 89, 91, 114 (2018).

⁸⁸ *Id.*

⁸⁹ U.S. FED. DEPOSIT INS. CORP., BUSINESS TECHNOLOGY STRATEGIC PLAN 2013–2017 (2013), https://www.fdic.gov/about/strategic/it_plan/BusinessTechnologyStrategicPlan2013-2017.pdf

⁹⁰ Michael Byrne, Big Data, FCC BLOG (Oct. 28, 2010, 1:06 PM), <https://www.fcc.gov/news-events/blog/2010/10/28/big-data>.

⁹¹ Matthew Reed, *Legal Entity Identifier System Turns a Corner*, FROM THE MANAGEMENT TEAM (July 3, 2014), <https://financialresearch.gov/from-the-management-team/2014/07/03/legal-entity-identifier-system-turns-a-corner/>

databases that can be used to train algorithms,⁹² while the Food and Drug Administration has tapped into cloud storage capacity to give the agency the ability to analyze Big Data.⁹³

Beyond these data-centered building blocks to artificial intelligence, U.S. administrative agencies are generally light years ahead of the U.S. judicial system in terms of employing algorithmic tools. After all, algorithmic tools of the statistical kind have long been a staple of administrative decision-making, especially when these agencies set policies and regulations.⁹⁴ Some government agencies, such as the U.S. Department of Commerce, even have data collection and analysis as among their principal responsibilities.⁹⁵ As a result, it is also not surprising that administrative agencies are also ahead of the courts in terms of their use of full-fledged machine-learning tools, something that the courts have yet to deploy. Admittedly, the use of machine learning within administrative agencies is not yet as extensive as it is in the private sector, but artificial intelligence is beginning to emerge to assist with important administrative functions—even though, again, we know of no example where artificial intelligence has fully replaced human decision-making.

We also know of no comprehensive survey of all uses of machine learning by administrative agencies at the state and federal levels. A Stanford University research team, however, is currently engaged in a multi-year effort, involving more than two dozen researchers with backgrounds in law and computer science, to survey the use of machine learning by the

⁹² See EPA's Cross-Agency Data Analytics and Visualization Program, <https://web.archive.org/web/20160414154548/https://www.epa.gov/toxics-release-inventory-tri-program/epas-cross-agency-data-analytics-and-visualization-program>

⁹³ Taha Kass-Hout, *FDA Leverages Big Data via Cloud Computing*, FDA VOICE (June 19, 2014), <http://blogs.fda.gov/fdavoices/index.php/2014/06/fda-leverages-big-data-via-cloud-computing>.

⁹⁴ For more recent discussions of the use of algorithmic analysis in public administration, see, e.g., DONALD F. KETTL, *LITTLE BITES OF BIG DATA FOR PUBLIC POLICY* (2018); JIM NUSSLE & PETER ORSZAG, EDs., *MONEYBALL FOR GOVERNMENT* (2014); ROBERT D. BEHN, *THE PERFORMANCESTAT POTENTIAL: A LEADERSHIP STRATEGY FOR PRODUCING RESULTS* (2014).

⁹⁵ See MICHAEL LEWIS, *THE FIFTH RISK* (2018).

federal government and develop a series of case studies.⁹⁶ Although this team has yet to issue a written report, its leaders did share some initial findings orally at a public meeting of the Administrative Conference of the United States in June 2019. Based on the government’s transcript of their remarks, the research team has “scoured public sources” for references to possible machine-learning uses at 100 federal agencies, yielding 117 “use cases” involving the reliance on algorithms—the “bulk” of which were “true AI in the modern learning sense as opposed to processed automation and conventional statistics.”⁹⁷ However, these examples were not distributed evenly across agencies: the Securities and Exchange Commission, for example, had “12 distinct use cases.”⁹⁸ As a result, only about “48 percent of agencies actively have invested in AI and machine learning tools.”⁹⁹ Furthermore, when the team members with computer science backgrounds looked closely at each use, they found only eight percent of these uses ranked “high in methodological sophistication”—revealing that there remains “still a fair bit of work to be done to close the private/public sector technology gap.”¹⁰⁰ Finally, the researchers appear not entirely confident that all of their use case actually entailed full machine learning systems, as they reported “some degree of puffery amongst agencies when they describe the adoption of machine learning and AI tools.”¹⁰¹

⁹⁶ Transcription of 71st Plenary Meeting of the Administrative Conference of the United States (June 13, 2019) (statement of David Engstrom). In the early part of this century, the federal General Accountability Office (GAO) conducted a survey of more than 125 federal agencies and reported that 52 relied on some form of “data mining,” which the GAO defined broadly “as the application of database technology and techniques—such as statistical analysis and modeling—to uncover hidden patterns and subtle relationships in data and to infer rules that allow for the prediction of future results.” U.S. GEN. ACCT. OFFICE, GAO-04-548, DATA MINING: FEDERAL EFFORTS COVER A WIDE RANGE OF USES 4 (2004). The GAO did not report whether any of these applications relied on machine learning rather than traditional analytic tools.

⁹⁷ *Id.* (statement of Dan Ho).

⁹⁸ *Id.*

⁹⁹ *Id.*

¹⁰⁰ *Id.*

¹⁰¹ *Id.*

It is also not entirely clear, based on what the Stanford team has reported to date, the precise stage of implementation at which all of the use cases find themselves. Some of these efforts appear to be prototypes, pilots, or even research studies. Still, the Stanford team's finding of approximately 100 use cases across the federal government at least suggests a plausible upper bound of the current extent of uses of machine learning at the federal level. Obviously still more uses exist at the state and local government level. We cannot purport to chronicle all instances of administrative machine learning in this paper, but instead we provide a range of examples to convey the variety of uses to which machine learning is being put by various agencies throughout the United States.

It is revealing that, among the use cases the Stanford team identified, the largest proportion (36 percent) were devoted to enforcement targeting—that is, helping identify cases of possible fraud or regulatory violations have human auditors or inspectors follow up to investigate. We thus first proceed in the next section to provide illustrative instances of machine learning used in the context of enforcement. We then proceed with examples in government services and program administration. Finally, we turn to a discussion of some of the merits, controversies, and legal issues surrounding the use of artificial intelligence in the administrative setting. Our discussion throughout both sections includes examples of machine learning at the federal, state, and local levels of government.

A. Enforcement

It is a common refrain that administrative agencies have more problems to deal with than they have resources to devote to solving them. Perhaps nowhere could this be more accurate than in the context of administrative enforcement. Agencies have limited number of auditors, inspectors, and other enforcement personnel who must oversee a vast number of the individuals

and businesses for their compliance with myriad pages of laws and regulations. The federal Occupational Safety and Health Administration, for instance, has no more than about 2,000 inspectors who oversee more than eight million workplaces employing about 130 million workers.¹⁰² To deploy these limited oversight resources smartly, agencies need to know which businesses or individuals are most likely to require oversight. Machine-learning algorithms can provide forecasts of the likelihood of violations, thus helping agencies allocate resources and decide which regulated entities to target.

The U.S. Internal Revenue Service (IRS), for example, began developing in 2001 machine learning risk tools to integrate data from prior tax records as well as data from other government agencies to help it predict cases of possible tax fraud and prioritize which taxpayers to target for auditing.¹⁰³ More recently, the IRS developed a machine-learning program that uses credit card information and other third-party data to forecast the probability of underreporting by businesses.¹⁰⁴ The Securities and Exchange Commission similarly uses machine learning and natural language processing to identify potential instances of insider trading or “bad apple” investment advisers or brokers.¹⁰⁵ The federal agency that oversees the major health care

¹⁰² Occupational Safety and Health Administration, Commonly Used Statistics, <https://www.osha.gov/oshstats/commonstats.html>.

¹⁰³ Jane Martin & Rick Stephenson, Risk-Based Collection Model Development and Testing, Presented at the Internal Revenue Service Research Conference (2005), <http://www.irs.gov/pub/irs-soi/05stephenson.pdf>; David DeBarr & Maury Harwood, Relational Mining for Compliance Risk, Presented at the Internal Revenue Service Research Conference (2004), <http://www.irs.gov/pub/irs-soi/04debarr.pdf>.

¹⁰⁴ See CHRIS WAGNER ET AL., TAXPAYER ADVOCACY SERV., *IRS Policy Implementation Through Systems Programming Lacks Transparency and Precludes Adequate Review*, in 2010 ANNUAL REPORT TO CONGRESS 71, 76, http://www.irs.gov/pub/irs-utl/2010arcmsp5_policythruprogramming.pdf; U.S. TREAS. INSPECTOR GEN. FOR TAX ADMIN., 2014-20-088, THE INFORMATION REPORTING AND DOCUMENT MATCHING CASE MANAGEMENT SYSTEM COULD NOT BE DEPLOYED (2014), <https://www.treasury.gov/tigta/auditreports/2014reports/201420088fr.pdf>.

¹⁰⁵ Pam Karlan & Joe Bankman, Artificial Intelligence and the Administrative State with Guests David Engstrom and Cristina Ceballos Stanford Legal (2019), <https://law.stanford.edu/stanford-legal-on-siriusxm/artificial-intelligence-and-the-administrative-state-with-guests-david-engstrom-and-cristina-ceballos/> (last visited Jun. 17, 2019).

benefits program, Medicare, relies in part on machine-learning algorithms to help it identify possible leads for its fraud investigators to pursue.¹⁰⁶

A number of state and local law enforcement authorities use algorithmic tools—some of which appear to be based on machine learning—in deciding where to send police patrols. Starting with a widely discussed Compstat initiative in New York City in the 1990s (which was not machine-learning based), many police departments across the United States have taken a more systematic approach to allocating law enforcement resources by using performance metrics and data analysis.¹⁰⁷ Today, this “moneyballing” effort includes both “place-based” and “person-based” predictive policing tools.¹⁰⁸ Place-based tools help police identify areas of a city that have a greater propensity for crime and may merit greater police patrols. For example, at least a dozen or more cities using a vendor-developed software called PredPol, which uses a proprietary algorithm to identify areas of a city which are more likely to be prone to criminal activity so that additional police resources can be allocated to those areas.¹⁰⁹ By contrast, the City of Los Angeles Police Department uses a Real-Time Analysis Critical Response (RACR) system and the New York City police department uses a machine-learning tool called Patternizr which are person-based: that is, they integrate information, detect patterns in crime incidents, and find

¹⁰⁶ Transcription, *supra* note 89 (statement of David Engstrom).

¹⁰⁷ By the end of the 1990s, a third of the largest police departments in the United States reported using a program like Compstat. David Weisburd et al., *The Growth of Compstat in American Policing*, Police Foundation Report (April 2004), <http://glenn.osu.edu/faculty/glenn-faculty/worley/June%2027%20-%20Reading%201%20-%20The%20Growth%20of%20Compstat%20in%20American%20Policing.pdf>.

¹⁰⁸ For the distinction between place-based and person-based prediction, see Ángel Díaz, *New York City Police Department Surveillance Technology* (Oct. 7, 2019), https://www.brennancenter.org/sites/default/files/2019-10/2019_NewYorkPolicyTechnology.pdf.

¹⁰⁹ Randy Rieland, *Artificial Intelligence is Now Used to Predict Crime. But is it Biased?*, SMITHSONIAN (Mar. 5, 2018), <https://www.smithsonianmag.com/innovation/artificial-intelligence-is-now-used-predict-crime-is-it-biased-180968337/>. See also Erica Goode, *Sending the Police Before There’s a Crime*, N.Y. TIMES (Aug. 15, 2011), <https://www.nytimes.com/2011/08/16/us/16police.html>; ANDREW GUTHRIE, *THE RISE OF BIG DATA POLICING SURVEILLANCE, RACE, AND THE FUTURE OF LAW ENFORCEMENT* (2017).

linkages between incidents in an effort to identify alleged perpetrators.¹¹⁰ Recent reports indicate that the Federal Bureau of Investigation and a number of state and local law enforcement agencies are using facial recognition tools marketed by private-sector firms such as Amazon in an effort to identify criminal suspects.¹¹¹ In May 2019, the city of San Francisco became the first major U.S. city to place restrictions on law enforcement's use of facial recognition and other surveillance tools.¹¹² Other cities are considering similar restrictions.

B. Services and Program Administration

Just as the police departments in major U.S. cities have deployed machine-learning tools to assist with law enforcement functions, cities are also using machine learning to support other key government functions. The New York City Fire Department, for example, follows its police counterparts in using machine learning algorithms to allocate and target its limited number of building inspectors that check for compliance with fire-related ordinances.¹¹³ In fact, New York City has established a central Office of Data Analytics which works to integrate data from across the city and develop “analytics tools to prioritize risk more strategically, deliver services more efficiently, enforce laws more effectively and increase transparency.”¹¹⁴ Other cities have

¹¹⁰ Caroline Haskins, *Dozens of Cities Have Secretly Experimented with Predictive Policing Software*, Motherboard (Feb. 6, 2019), https://www.vice.com/en_us/article/d3m7jq/dozens-of-cities-have-secretly-experimented-with-predictive-policing-software; Phil Goldstein, *How Pattern Recognition and Machine Learning Helps Public Safety Departments*, StateTech (May 3, 2019), <https://statetechmagazine.com/article/2019/05/how-pattern-recognition-and-machine-learning-helps-public-safety-departments-perfcon>.

¹¹¹ See, e.g., Jon Schuppe, *Facial Recognition Gives Police a Powerful New Tracking Tool. It's Also Raising Alarms*, NBC NEWS (July 30, 2018), <https://www.nbcnews.com/news/us-news/facial-recognition-gives-police-powerful-new-tracking-tool-it-s-n894936>; Jon Schuppe, *How Facial Recognition Became a Routine Policing Tool in America*, NBC NEWS (May 11, 2019), <https://www.nbcnews.com/news/us-news/how-facial-recognition-became-routine-policing-tool-america-n1004251>.

¹¹² San Francisco Ordinance on Acquisition of Surveillance Technology (May 6, 2019), <https://sfgov.legistar.com/View.ashx?M=F&ID=7206781&GUID=38D37061-4D87-4A94-9AB3-CB113656159A>

¹¹³ Brian Heaton, *New York City Fights Fire with Data*, GOV'T TECH. (May 15, 2015), <http://www.govtech.com/public-safety/New-York-City-Fights-Fire-with-Data.html>.

¹¹⁴ NYC ANALYTICS, <https://www1.nyc.gov/site/analytics/about/about-office-data-analytics.page>.

similarly created special offices or teams devoted to data analysis and prediction.¹¹⁵ Los Angeles has established a formal partnership, the Data Science Federation, with local colleges and universities to promote “predictive ... analysis that will help drive data driven decision making within the city.”¹¹⁶ The City of Chicago worked with a consortium of university partners to create a SmartData Platform which helps facilitate the use of machine learning in support of city services.¹¹⁷

In Chicago, some of these services supported by machine-learning tools include health inspections of restaurants, with inspectors assigned based on the algorithmic forecasts of the establishments posing the greatest risks.¹¹⁸ Both Chicago and Washington, D.C., are using machine learning to optimize rodent bait placement throughout their cities.¹¹⁹ In Flint, Michigan, following a major fiasco in the management of the city’s water supply, officials have benefited from machine-learning predictions to identify priorities for replacing pipes contributing to lead contamination.¹²⁰ In Los Angeles, traffic lights operate automatically based on a machine-learning system that optimizes for congestion avoidance using data fed by a network of sensors

¹¹⁵ See, e.g., City of Boston, Analytics Team, <https://www.boston.gov/departments/analytics-team>; City of Chicago, Data Science, https://www.chicago.gov/city/en/depts/doi/provdrs/data_sciences.html.

¹¹⁶ <https://datasciencefederation.lacity.org/about-us>.

¹¹⁷ Ash Center Mayor’s Challenge Research Team, *Chicago’s SmartData Platform: Pioneering Open Source Municipal Analytics*, DATA-SMART CITY SOLUTIONS (Jan. 8, 2014), <http://datasmart.ash.harvard.edu/news/article/chicago-mayors-challenge-367>.

¹¹⁸ Stephen Goldsmith, *Chicago’s Data-Powered Recipe for Food Safety* (May 21, 2015), <https://datasmart.ash.harvard.edu/news/article/chicagos-data-powered-recipe-for-food-safety-688>. Boston developed similar restaurant-inspection algorithms through a crowdsourcing project. See Edward Glaeser et al., *Crowdsourcing City Government: Using Tournaments to Improve Inspection Accuracy*, 106 AM. ECON. REV. 114 (2016).

¹¹⁹ Linda Poon, *Will Cities Ever Outsmart Rats?*, CityLab (Aug. 9, 2017), <https://www.citylab.com/solutions/2017/08/smart-cities-fight-rat-infestations-big-data/535407/>.

¹²⁰ Gabe Cherry, *Google, U-M to Build Digital Tools for Flint Water Crisis*, UNIV. OF MICH. NEWS (May 3, 2016), <http://ns.umich.edu/new/multimedia/videos/23780-google-u-m-to-build-digital-tools-for-flint-water-crisis>.

in the city's streets.¹²¹ Johnson County, Kansas, has used algorithmic determinations of risk to determine how to allocate its social service counselors and mental health professionals.¹²²

The Data-Smart City Solutions initiative at Harvard University's John F. Kennedy School of Government has cataloged 64 uses of data analytics by local governments, some but not all involving machine learning.¹²³ Its list includes tasks as varied as identifying children who could benefit from mentoring programs, prioritizing trees for trimming, and identifying businesses that might be underpaying taxes.

At the federal level too, predictive analytic tools, including machine learning, have been put to varied uses. One of the earliest uses of machine learning by the federal government actually helped spur innovations in the technique: the U.S. Postal Service's use of artificial intelligence to support automatic handwriting detection and mail sorting.¹²⁴ In addition, scientists at the National Oceanic and Atmospheric Administration have relied on machine learning for weather forecasting.¹²⁵ Risk analysts at the Environmental Protection Agency have used machine learning to forecast the likelihood that certain chemicals are toxic and need further study and management.¹²⁶ The Food and Drug Administration has employed machine learning to extract

¹²¹ Ian Lovett, *To Fight Gridlock, Los Angeles Synchronizes Every Red Light*, N.Y. TIMES (Apr. 1, 2013), <http://www.nytimes.com/2013/04/02/us/to-fight-gridlock-los-angeles-synchronizes-every-red-light.html>; David Z. Morris, *How Swarming Traffic Lights Could Save Drivers Billions of Dollars*, FORTUNE (July 13, 2015, 4:47 PM), <http://fortune.com/2015/07/13/swarming-traffic-lights>.

¹²² Robert Sullivan, *Innovations in Identifying People Who Frequently Use Criminal Justice and Healthcare Systems* (May 16, 2018), <https://www.prainc.com/innovations-identification-cj-healthcare/>.

¹²³ Data-Smart City Solutions, *A Catalog of Civic Data Use Cases*, <https://datasmart.ash.harvard.edu/news/article/how-can-data-and-analytics-be-used-to-enhance-city-operations-723>.

¹²⁴ One of the first automatic techniques for detecting handwriting emerged in the late 1980s in the context of mail sorting. See Ching-Huei Wang & Sargur Srihari, *A Framework for Object Recognition in a Visually Complex Environment and Its Application to Locating Address Blocks on Mail Pieces*, 2 INT'L J. COMP. VISION 125, 125 (1988); Ofer Matan et al., *Handwritten Character Recognition Using Neural Network Architectures*, presented at the 4th USPS Advanced Technology Conference (1990), <http://yann.lecun.com/exdb/publis/pdf/matan-90.pdf>.

¹²⁵ David John Gagne et al., *Day-Ahead Hail Prediction Integrating Machine Learning with Storm-Scale Numerical Weather Models*, Presented at the Twenty-Seventh Conference on Innovative Applications of Artificial Intelligence (2015), <https://www.aaai.org/ocs/index.php/IAAI/IAAI15/paper/viewFile/9724/9898>.

¹²⁶ Robert Kavlock et al., *Update on EPA's ToxCast Program: Providing High Throughput Decision Support Tools for Chemical Risk Management*, 25 CHEM. RES. TOXICOLOGY 1287, 1295 (2012); Matthew Martin et al.,

information from adverse event reports about drugs.¹²⁷ Similarly, the Bureau of Labor Statistics uses machine learning to code survey results about workplace injuries.¹²⁸ The U.S. Patent and Trademark Office is exploring how to use machine learning to identify existing literature that may be novelty-destroy “prior art” to patent applications.¹²⁹ The Customs and Border Protection uses facial-recognition algorithms to identify people when they arrive in the United States from international airplane flights.¹³⁰ The Social Security Administration uses a natural language processing tool based on machine learning that helps flag initial decisions adjudicating disability claims for further quality review.¹³¹

C. Impacts and Issues

The principal advantages of artificial intelligence in the administrative context are similar to those in the private sector: accuracy and efficiency. Machine-learning algorithms can make more accurate forecasts than can aid in governmental decision-making. For example, researchers have shown that if the U.S. EPA were to assign its water pollution inspectors using a machine-learning algorithm in instead of just identifying facilities at random to inspect, they could

Economic Benefits of Using Adaptive Predictive Models of Reproductive Toxicity in the Context of a Tiered Testing Program, 58 *SYS. BIOLOGY REPROD. MED.* 3, 4–6 (2012); Richard S. Judson et al., Estimating Toxicity-Related Biological Pathway Altering Doses for High-Throughput Chemical Risk Assessment, 24 *CHEM. RES. TOXICOLOGY* 451, 457–60 (2014).

¹²⁷ *Commissioner’s Fellowship Program: Final Report Abstracts*, U.S. FOOD & DRUG ADMIN., <http://www.fda.gov/AboutFDA/WorkingatFDA/FellowshipInternshipGraduateFacultyPrograms/CommissionersFellowshipProgram/ucm413253.htm>.

¹²⁸ Alex Measure, Machine Learning: How Bureau of Labor Statistics Did It (July 25, 2019), <https://digital.gov/event/2019/07/25/machine-learning-how-bureau-labor-statistics-did-it/>.

¹²⁹ Arti Kaur Rai, Machine Learning at the Patent Office: Lessons for Patents and Administrative Law (May 24, 2019), <http://dx.doi.org/10.2139/ssrn.3393942>.

¹³⁰ Karlan & Bankman, *supra* note __ (interview with David Engstrom).

¹³¹ Gerald Ray & Glenn Sklar, An Operational Approach to Eliminating Backlogs in the Social Security Disability Program 31–34 (June 2009), http://www.crfb.org/sites/default/files/An_Operational_Approach_to_Eliminating_Backlogs_in_the_Social_Security_Disability_Program.pdf.

improve increase the accuracy of finding violations of the Clean Water Act by 600 percent.¹³² A separate analysis of a machine-learning tool that could identify chemicals likely to be toxic showed that it could save the EPA nearly \$980,000 for every toxic chemical identified.¹³³

In addition to improving the allocation of scarce administrative resources, machine-learning systems may be able eventually help in reducing some of the inevitably biases that are found with unaided human judgment.¹³⁴ For example, in the context of the Social Security Administration's disability adjudications, some research suggests that human decisions reflect racial disparities that tend to disfavor claimants of color.¹³⁵ A study of just a single office within the Social Security Administration found vastly disparate rates of benefits awards: "judge grant rates in this single location ranged ... from less than 10 percent being granted to over 90 percent."¹³⁶ If machine-learning tools are used as either substitutes for or even just complements to human decision-making, they may be able to reduce inconsistencies and other foibles that permeate human judgment.

On the other hand, the use of machine learning in governmental settings has not escaped controversy. If the underlying data contain biases—particularly as administrative data can derive from human practices and systems that themselves reflect biases and prejudices—then machine learning might simply reify the inequities built into the data.¹³⁷ Moreover, if the underlying data

¹³² Miyuki Hino, Elinor Benami, & Nina Brooks, *Machine Learning for Environmental Monitoring*, 1 NATURE SUSTAINABILITY 583 (2018).

¹³³ Matthew Martin et al., *Economic Benefits of Using Adaptive Predictive Models of Reproductive Toxicity in the Context of a Tiered Testing Program*, 58 SYS. BIOLOGY REPROD. MED. 3, 4–6 (2012).

¹³⁴ See, e.g., DANIEL KAHNEMAN, THINKING, FAST AND SLOW (2011).

¹³⁵ U.S. GENERAL ACCOUNTING OFFICE, SOCIAL SECURITY: RACIAL DIFFERENCES IN DISABILITY DECISIONS WARRANTS FURTHER INVESTIGATION (1992), <https://www.gao.gov/assets/160/151781.pdf>; Erin Godtland et al., *Racial Disparities in Federal Disability Benefits*, 25 CONTEMP. ECON. POL. 27 (2007).

¹³⁶ TRAC, *Social Security Awards Depend More on Judge than Facts: Disparities Within SSA Disability Hearing Offices Grow* (2011), <https://trac.syr.edu/tracreports/ssa/254/>.

¹³⁷ For recent discussions, compare CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY (2016), with MICHAEL KEARNS & AARON ROTH, THE ETHICAL ALGORITHM: THE SCIENCE OF SOCIALLY AWARE ALGORITHM DESIGN (2019).

on which algorithms rely is limited, or they are not designed or tested well, they may lead to a false sense of accuracy—perhaps even making decision-making more error-prone. Furthermore, reliance on large amounts of data gives rise to privacy concerns and abuses of power by irresponsible or oppressive governmental actors.¹³⁸

In the governmental setting, machine-learning algorithms’ “black box” character seem to raise particular concerns about transparency and accountability. The ways that such algorithms optimize outcomes and the solutions they support may not be readily apparent to those who they affect, which has suggested to some observers that either they be avoided in the public sphere or that government officials take extra strides to explain what these algorithms do. For example, a school district in Boston worked with researchers at the Massachusetts Institute of Technology to use a machine-learning algorithm to help redesign student bus schedules in a way that would have saved the district up to \$15 million in annual expenses and produced schedules that were healthier for students, better for the environment, and more equitable for minority students.¹³⁹ But the bus schedule’s “overhaul was introduced with insufficient explanation or opportunity for citizen interaction with the model,” resulting in a “public pushback [that] was strong and swift.”¹⁴⁰ The school district eventually dropped the proposed scheduling changes.

In Houston, a school district ended up in court after relying on a complex algorithm—albeit not a machine-learning one—to rate teachers’ performance and justify the dismissal of teachers who rated poorly. The district relied on a private consulting firm to develop and run the

¹³⁸ See, e.g., Shibani Mahtani, *Chicago Police Take a Page From “Minority Report,”* WALL ST. J. (May 12, 2017), <https://www.wsj.com/articles/chicago-police-take-a-page-from-minority-report-1494581400>; Ali Winston & Ingrid Burrington, *A Pioneer in Predictive Policing Is Starting a Troubling New Project,* THE VERGE (Apr. 27, 2018), <https://www.theverge.com/2018/4/26/17285058/predictive-policing-predpol-pentagon-ai-racial-bias>.

¹³⁹ Ellen Goodman, *The Challenge of Equitable Algorithmic Change,* REG. REV. (2019), <https://www.theregreview.org/wp-content/uploads/2019/02/Goodman-The-Challenge-of-Equitable-Algorithmic-Change.pdf>

¹⁴⁰ *Id.* at 3, 7.

algorithm, but the firm considered its “algorithms and software as trade secrets, refusing to divulge them to either [the district] or the teachers themselves.” The teachers’ union and several teachers took the district to court, arguing that the algorithm deprived them of procedural due process. They argued that, without “access to the computer algorithms and data necessary to verify the accuracy of their scores,” the district deprived them of their constitutional rights. The trial court issued only an interim decision, ruling that the procedural due process claim could possibly have merit and the teachers were entitled to take their case to a jury. The court held that, “without access to . . . proprietary information—the value-added equations, computer source codes, decision rules, and assumptions—[the teachers’] scores will remain a mysterious ‘black box,’ impervious to challenge.” Although the court recognized that the consulting firm may well have been in its rights to keep its algorithms secret, a jury could consider whether “a policy of making high stakes employment decisions based on secret algorithms [is] incompatible with minimum due process.” Of course, the preliminary nature of the trial court’s decision cannot rule out the possibility that, had the matter gone to a jury, the school officials might have been able to put forth additional evidence that, while still protecting the firm’s trade secrets, could still have satisfied the teachers’ due process rights.

The Houston case is the closest of any court decision of which we are aware to raise legal issues of the kind that might be raised by the administrative use of machine-learning algorithms.¹⁴¹ Concerns about algorithmic bias could manifest in legal claims raising claims under federal statutes, such as Title VI of the Civil Rights Act of 1964, which prevents state and local governments that receive federal financial assistance from engaging in practices that have disparate impacts on protected classes. The due process and equal protection clauses of the

¹⁴¹ The cases discussed in Part I of this report addressing judicial use of algorithms are also obviously relevant to administrative use of algorithms. However, those cases also did not address any truly *machine-learning* algorithm.

Constitution's Fourteenth and Fifth Amendments would also prevent state and federal governments, respectively, from engaging in intentionally discriminatory practices. If they are not careful, governmental entities could certainly use machine-learning tools in ways that offend principles of constitutional or administrative law. However, as one of us has written elsewhere, the responsible use of machine learning can probably be readily accommodated under existing principles of U.S. law.¹⁴²

CONCLUSION

Although the day when a judge's role is fully supplanted by an algorithm is surely still far into the future, if it should ever completely come,¹⁴³ the building blocks that could eventually give rise to a world of increased use of artificial intelligence by governmental entities have started to manifest themselves in state and federal legal systems across the United States. The widespread adoption of risk assessment tools in criminal cases in courts at every level of government appears to reflect an increasing comfort in allowing algorithms to inform decisions. Increasing digitization of court records could potentially provide judicial managers with troves of data for artificial intelligence programs that could analyze and possibly even facilitate future automated adjudication. The growing adoption of online dispute resolution by both the private organizations and the courts could also eventually make the public more comfortable with fully computerized and automated adjudication. The opportunities for successful application of artificial intelligence seem even greater in administrative agencies, and they are already starting

¹⁴² Coglianesse & Lehr, *supra* note 1; Coglianesse & Lehr, *supra* note 12.

¹⁴³ Dave Orr and Colin Rule, Artificial Intelligence and the Future of Online Dispute Resolution, https://static1.squarespace.com/static/59bb4f5dc027d80ba3e5be79/t/5ac1625e70a6ad85f4d37fbf/1522623070923/OrrRule_AI.pdf ("We are still a long way away from giving an AI Lexis-Nexis access and then asking it to serve on the Supreme Court.").

to rely on machine learning tools to inform enforcement decisions, allocate social services, and manage programs.

Overall, these tools have great promise. As with any tool, of course, if they are not used with care, they may give rise to further problems, including conflict and litigation. Public concerns have already arisen over the use of algorithms in predictive policing and in the criminal justice system more generally. The few cases decided to date do not suggest that the courts will ultimately disapprove of responsibly designed and implemented machine-learning tools—and it is certainly beyond the limits of any kind of intelligence, human or artificial, to forecast with precision what the future will hold for governmental use of machine learning in the United States. Yet with the continued reliance on machine learning in other spheres of life, the public acceptability of, if not demand for, its use in the governmental sector may only increase.