Chapter 20

Getting Smart on the Government and AI

By Robin Nunn

Emerging technologies such as artificial intelligence (AI) have captured the attention of the private sector. From start-ups to large companies, many businesses are making use of or actively investigating such technologies to make their businesses more productive, increase efficiencies, and improve their bottom line.

Governments around the world are also actively assessing how AI can be deployed. AI has beneficial applications in many areas of the public sector. For example, with data collected in real time from traffic lights, CCTV cameras, and other sources assist officials with traffic management. Customer service centers manned by robots that use AI help answer questions. Algorithms and machine-learning techniques, in which computers analyze large amounts of data to detect statistical patterns and develop models that can be used to make accurate predictions, are rapidly becoming key tools for governments.

However, despite the obvious opportunities for efficiency and effectiveness, the role of AI, automation, and robotics in government policy and service delivery remains challenging. There are more questions than answers at this point. For example can regulators prevent algorithms based on historical data from perpetuating or reinforcing decades of bias? When is it acceptable to use black box deep-learning models, where the logic used for decisions cannot possibly be explained or understood even by the data scientists designing the underlying algorithms? The United States is investing in enhancing its AI capabilities, but there may be associated risks for individuals and companies subject to government decisions made with AI, especially because many government processes are premised on principles of non-discrimination that may not always be fully reflected when AI is used to enable government action. This chapter examines the relevant underlying technologies, the trends in government adoption, and the corresponding legal touchpoints for challenging government use of AI on the basis of algorithmic bias.

I. DEVELOPMENT OF AI TECHNOLOGY

Leading up to recent incidences of government adoption of AI, the private sector has experienced exponential growth in the investment, development, and implementation of AI technologies in the last two decades.[1] AI algorithms are often not blazingly new, many are decades old. But we now have comparatively huge volumes of data that can be stored and processed cheaply, such that the performance of and ability to further research AI systems has greatly increased. There are several different AI applications that feature prominently as the public sector begins to adopt AI:

- *Machine Learning:* a subset of AI that often uses statistical techniques to give machines the ability to "learn" from data without being explicitly given the instructions for how to do so. This process is known as "training" a "model" using a learning "algorithm" that progressively improves model performance on a specific task;
- *Neural Networks:* multiple layers of weighted nodes, including at least one "hidden layer;" which can be "trained" to perform certain tasks (e.g., facial recognition, detecting fraud, predicting stock performance) via large data sets and means of rewarding and penalizing desired and undesirable outcomes; and
- *Expert Systems:* computer systems, typically rule-based, emulating human experts' decision-making ability (e.g., for medical diagnoses).

Especially in the context of machine learning, AI systems operate by finding relationships between input features and known outcomes in a set of training data.[2] What generally makes these systems different from human decision-making models is that the machines themselves, without direct human intervention, develop the rules that best predict the known outcomes based on the input data—this set of rules is known as the "model."[3] This model is then applied to future, unobservable cases of interest and predicts the results.[4]

AI and neural networks are designed to learn from data fed to them. This is how many agencies can accurately target and predict future outcomes. The downside to this is that AI does not necessarily understand the data or the nuances of social structures. This can mean that AI can acquire operator and societal biases, resulting in discrimination in administration.

---

1. Louis Columbus, "10 Charts That Will Change Your Perspective On Artificial Intelligence's Growth," Forbes, (January 12, 2018), available at https://www.forbes.com/sites/louiscolumbus/2018/01/12/10-charts-that-will-change-your-perspective-on-artificial-intelligences-growth/#15ae625d4758.

2. Andrew Selbst, *A Mild Defense of Our New Machine Overlords*, 70 Vand. L. Rev. En Banc 87, 90 (2017).

3. *Id.*

4. *Id.*

II. GOVERNMENT ADOPTION OF AI

In recent years, computing has begun to shift from merely relieving governments of routine work such as data entry, to a new era involving the automation of tasks previously thought to require human judgment. As part of this new era, governments have increasingly begun to incorporate AI technology into the regulatory process. New York City, for example, has established a Mayor's Office of Data Analytics, which, among other things, is working with the city's fire department to use machine learning to decide where to send building inspectors.[5] The Internal Revenue Service has launched an Information Reporting and Document Matching program, which applies algorithms to credit card and other third-party data to predict tax underreporting and non-filing by businesses.[6]

Outside the United States, the government of the People's Republic of China (China) has been active in the AI space. For instance, China has been in the process of building facial recognition technology that relies heavily on artificial intelligence to monitor its citizens. It has also set up large surveillance efforts to track members of the Uighur Muslim minority and map their relationships with family and friends.[7] The Chinese government will often point to instances of bike theft or jaywalking prevention as the rationale for giving the Chinese government more control over its citizens through these efforts.[8]

In Europe, the European Commission adopted a plan on December 7, 2018, to foster the development and use of AI in Europe, with a focus on becoming the world-leading region on cutting-edge, ethical and secure AI.[9]

---

5. Mayor Bloomberg and Fire Commissioner Cassano Announce New Risk-Based Fire Inspections Citywide Base on Data Mined from City Records, City of New York (May 16, 2018), available at https://www1.nyc.gov/office-of-the-mayor/news/163-13/mayor-bloomberg-fire-commissioner-cassano-new-risk-based-fire-inspections-citywide#/10. As discussed further below, the City of New York has also passed a local law that establishes an "Automated Decision Systems Task Force" which will explore how the City uses algorithms. The task force, the first of its kind in the United States, will work to develop a process for reviewing "automated decision systems:" commonly known as algorithms, through the lens of equity, fairness and accountability." 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.

6. David DeBarr & Maury Harwood, *Relational Mining for Compliance Risk*, Presented at the Internal Revenue Service Research Conference (2004), available at http://www.irs.gov/pub/irs-soi/04debarr.pdf.

7. Paul Mozar, *One Month, 500,000 Face Scans: How China Is Using A.I. to Profile a Minority*, New York Times (April 14, 2019), available at https://www.nytimes.com/2019/04/14/technology/china-surveillance-artificial-intelligence-racial-profiling.html.

8. Paul Mozar, *Inside China's Dystopian Dreams: A.I., Shame and Lots of Cameras*, New York Times (July 8, 2018), available at https://www.nytimes.com/2018/07/08/business/china-surveillance-technology.html.

9. Press Release, European Commission, Member States and Commission to work together to boost artificial intelligence "made in Europe" (December 7, 2018), available at http://europa.eu/rapid/press-release_IP-18-6689_en.htm.

These are just a few examples of governments adopting AI to supplement traditional government activities. And while in the U.S. context none of these examples amount to facially illegal government action, small changes in facts or government process could easily yield a transgression on an individual right, as explained below.

III. RISKS OF GOVERNMENT AI ADOPTION

With this rise in adoption of AI in the governmental process, government agencies, companies, and their counsel will have to increasingly confront how decades-old processes and procedures apply in the new era. For example, certain properties of AI, and especially machine learning, combine to distinguish it from other analytical techniques and give rise to potential concerns about the greater reliance on machine learning by regulatory agencies.

The first is machine learning's self-studying property. The results of algorithms do not depend on humans specifying in advance how each variable is to be factored into the predictions; indeed, as long as learning algorithms are running, humans are not really controlling how they are combining and comparing data. Machine learning systems "learn" from the data, meaning that these algorithms find patterns or correlations between variables in a set of data, which can then be used to make predictions.[10]

The second key property is machine learning's black box nature. Unless specifically designed to ensure transparency, the results of many machine learning systems are not intuitively explainable and cannot support causal explanations of the kind that underlie the reasons traditionally offered to justify governmental action. As a result, it can be difficult to explain exactly how or why a machine-learning algorithm keys in on certain correlations or makes the predictions that it does. As such, legal commentators have lamented the black box nature of machine learning-based algorithms, arguing that if we cannot see the code or interact with it, we cannot appropriately, or legally, make use of it.[11] And in many cases, due to trade secrecy or other reasons for lack of access, such access might prove impossible.[12]

Finally, machine learning, as with other computational strategies in today's digital era, can be fast and automatic, supporting uses in which the algorithm produces results that can shorten or potentially bypass human deliberation and decision making. All

10. Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a "Right to Explanation" is Probably Not the Remedy you are Looking For?*, 16 Duke L. & Tech. Rev. 18, 25 (2017).

11. *See, e.g.*, Frank Pasquale, The Black Box Society 3-4 (2015); Brenda Reddix-Smalls, *Credit Scoring and Trade Secrecy: An Algorithmic Quagmire or How the Lack of Transparency in Complex Financial Models Scuttled the Finance Market*, 12 U.C. Davis Bus. L. J. 87 (2011); Frank Pasquale, *Restoring Transparency to Automated Authority*, 9 J. on Telecomm. & High Tech. L. 235, 237 (2011).

12. Mikella Hurley & Julius Adebayo, *Credit Scoring in the Age of Big Data*, 18 Yale J.L. & Tech. 148, 196-98 (2016).

three of these factors combine to make machine learning techniques appear qualitatively more independent from humans when compared to other statistical techniques.[13]

One problem with the government adopting AI involves data management and proper oversight. Letting computers make decisions could cause serious problems if there is no additional level of oversight. Algorithms learn by being fed certain data, often chosen by engineers, and the system builds a model of the world based on that data. So, for instance, if a system is trained on photos of people who are overwhelmingly white, it will have a harder time recognizing nonwhite faces, leading to the emergence of problematic biases baked into predictions. Indeed, this is precisely why IBM recently announced plans to release a database of more than 1 million facial images to academics, public interest groups and competitors.[14] Release of this information is intended to improve training of machine learning applications used in facial recognition systems.

One example that particularly exemplifies the dangers associated with government use of AI is "predictive policing," or the use of predictive and analytical techniques in law enforcement to identify potential criminal activity. ProPublica published an investigation that found that widely used software that assessed the risk of recidivism in criminals was twice as likely to mistakenly identify blacks as being at a higher risk of committing future crimes. It was also twice as likely to incorrectly flag whites as low risk.[15] In a recent test of Amazon's facial recognition system "Rekognition", the ACLU tested the software on members of Congress and found that the software incorrectly matched 28 members of Congress as persons who have been arrested for a crime. It also found that the false matches were disproportionately of people of color.[16]

Although efforts have been recently undertaken to remove the bias from datasets, many of the initial datasets used to train these algorithms were based on disproportionately higher numbers of images of white men and therefore had greater accuracy for white men than for women or minorities.[17]

Similarly, judges are increasingly turning to artificial intelligence in making bail decisions. While there are serious concerns regarding judges relying on intuition and personal preference in setting bail, there are also grave concerns about the use of

13. Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 Geo. L. J. 1147, 1167 (2017).

14. Aleksandra Mojsilovic & John Smith, *IBM to Release World's Largest Annotation Dataset for Studying Bias in Facial Analysis*, IBM (June 27, 2018), available at https://www.ibm.com/blogs/research/2018/06/ai-facial-analytics/.

15. Julia Angwin, et. al., *Machine Bias*, ProPublica (May 23, 2016), available at https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

16. Jacob Snow, *Amazon's Face Recognition Falsely Matched 28 Members of Congress With Mugshots*, aclu.org (July 26, 2018), available at https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28.

17. Steve Lohr, *Facial Recognition Is Accurate, if You're a White Guy*, New York Times (February 9, 2018), available at https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html.

algorithms for bail decisions without appropriate oversight, safeguards, and transparency regarding their inner workings.[18]

Certain studies show that improperly trained machines will mistake correlation for causation, which may be the reason for mistakes in prediction.[19] Unfortunately, because these algorithms are secret, or are based on proprietary information, we do not know why these predictions ended up being so skewed and wrong.

Today, judges and even police departments across the United States are relying on machine-driven risk assessments in different ways—some may use them regularly while others discount them entirely—but there is little these government officials can do to understand the logic behind them. As "predictive policing" crime prevention efforts gain traction, cities' methods of policing are adapting. More and more cities could start to rely on software analyses of large sets of historical crime data to forecast where future crime hot spots are most likely to emerge; the police are then directed to those areas.

This scenario clearly sets up the distinct possibility of perpetuating an already vicious cycle, in which the police increase monitoring in certain locations they are already monitoring, thus increasing the probability of arrests from those same areas. While humans may be able to identify and prevent this type of biased outcome, smart algorithms, unless they are prompted to account for the unique characteristic of data inputs, may not. And, in the United States, this could result in more surveillance in poor, non-white neighborhoods, while rich white neighborhoods are left alone. Humans are inherently biased by virtue of their limited ability to absorb and reason, and predictive programs are only as good as the data they are trained on. When algorithms are trained by humans on this data that has a complex history, they learn these inherent biases and bring them forward into their own beliefs.

IV. EFFORTS TO ADDRESS RISKS OF AI

As AI has become more ubiquitous, policymakers are looking more closely at developing policies that foster the great potential for AI applications and systems, and leveraging these tools to improve government efficiency and operations. In 2016, the Obama Administration sought to control risks of unconstitutional government action involving AI. To that end, the administration convened a series of workshops and published two separate reports outlining strategies for supporting the long-

18. Sam Corbett-Davis, et. al., *Even Imperfect Algorithms Can Improve the Criminal Justice System*, New York Times (December 20, 2017), available at https://www.nytimes.com/2017/12/20/upshot/algorithms-bail-criminal-justice-system.html; Jason Tashea, *Courts are Using AI to Sentence Criminals. That Must Stop Now*, Wired (April 17, 2017), available at https://www.wired.com/2017/04/courts-using-ai-sentence-criminals-must-stop-now/.

19. *See, e.g.,* Robin Wigglesworth, *Spurious Correlations are Krytonite of Wall St's AI Rush*, Financial Times (March 14, 2018), available at https://www.ft.com/content/f14db820-26cd-11e8-b27e-cc62a39d57a0.

term development of AI through increased R&D, access to public datasets, greater collaboration between industry and government, and other strategies.[20] More recently, in May of 2018, the Trump Administration convened an Artificial Intelligence Summit and announced several initiatives that the administration has underway, including increasing funding for AI, removing barriers to development, and using AI "to improve the efficiency of government services."[21] In addition, on February 11, 2019, President Trump signed an executive order to encourage the investment in research and development of AI. The executive order intends to educate workers, improve access to services and data needed to build AI systems. Although many of the details of this American AI Initiative remain yet to be determined, it directs federal agencies to prioritize investments in research and the development of AI.[22]

Congress has also introduced several pieces of legislation at the federal level, including the National Security Commission on Artificial Intelligence Act, which would create an independent National Security Commission on Artificial Intelligence; the House-passed the Self Drive Act, which addresses the safety of automated vehicles; the AV Start Act, a bipartisan Senate companion that similarly tackles driverless cars; the Future of AI Act, a bipartisan Senate bill that would create an advisory committee on AI issues; and the AI in Government Act of 2018, which directs certain agencies to specifically research and consider AI, and to create an advisory board to address AI related issues.[23]

Notably, there is some early movement by the government to address concerns regarding the use of these new tools for government work. The National Defense Reauthorization Act of 2019 establishes a National Security Commission on AI and directs the Department of Defense to conduct an in-depth review of how AI may be used in defense systems, including ethical considerations.[24] Additionally, the City of New York has passed a local law that establishes an Automated Decision Systems Task Force which will explore how New York City uses algorithms. The task force, the first of its kind in the United States, will work to develop a process for reviewing "automated decision systems," commonly known as algorithms, through the lens of

20. Ed Felten, *Preparing for the Future of Artificial Intelligence*, The White House Blog (May 3, 2016), available at https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.

21. *Artificial Intelligence for the American People*, The White House (May 10, 2018), available at https://www.whitehouse.gov/briefings-statements/artificial-intelligence-american-people/.

22. *Accelerating America's Leadership in Artificial Intelligence*, The White House, Office of Science and Technology Policy (February 11, 2019), available at https://www.whitehouse.gov/articles/accelerating-americas-leadership-in-artificial-intelligence/.

23. *AI in Government Act of 2018*, S. 3502, 115th Cong. (2017-2018), available at https://www.congress.gov/115/bills/s3502/BILLS-115s3502is.pdf.

24. The National Defense Reauthorization Act of 2019, H.R. 5515, 115th Cong. (2017-2018), available at https://www.congress.gov/bill/115th-congress/house-bill/5515.

equity, fairness and accountability.[25] Other states such as Washington, Massachusetts, and Illinois have considered bills that are intended to ensure accountable and fair use of automated decision systems in state government.[26]

In addition, there are many other ethics-setting organizations, including AI Now, an institute comprising leading researchers and developers, who are beginning to weigh in on government use of AI.[27]

Outside the United States, efforts to protect data privacy are promising. The EU General Data Protection Regulation (GDPR) may have an impact on government use of AI. As of May 25, 2018, all organizations doing business in the European Union must comply with the new European privacy legislation GDPR—arguably this includes both European and even non-European governments.[28] Article 22 of the GDPR says, that the person "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." In practice, this means that, in the case of important decisions like mortgages, loans, job applications, school admissions, judicial decisions, etc., one must always offer the person the choice to have the decision made by a human being or with significant human involvement. For example, this passage has been interpreted to require an appeal for decisions that are made with no human involvement, such as when a machine learning algorithm decides if you are eligible for a loan, for example, in order to prevent discrimination. Such expectations may apply to government agencies when they target citizens in Europe.

V. CHALLENGING GOVERNMENT ACTION

AI technology remains controversial in the context of government action, in part, because algorithms are not always clear on their decision-making logic. It was troubling enough when Flickr, which applies automatic labels to pictures in digital photo albums, was labeling images of black people as gorillas. Or when Google search results for black-sounding names are more likely to be accompanied by ads about criminal activity than search results for white-sounding names. Or when Microsoft released a

25. *Mayor De Blasio Announces First-In-Nation Task Force to Examine Automated Decision Systems Used by the City*, City of New York (May 16, 2018), available at 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.

26. Christopher Mims, *Our Software Is Biased Like We Are. Can New Laws Change That?*, March 23, 2019, available at https://www.wsj.com/articles/our-software-is-biased-like-we-are-can-new-laws-change-that-11553313609.

27. *See, e.g.,* Dillon Reisman, et. al., *Algorithmic Impact Assessments*, AI Now Institute (April 2018), available at https://ainowinstitute.org/aiareport2018.pdf.

28. Dawn Kawamoto, *Will GDPR Rules Impact States and Localities?,* Government Technology (May 3, 2018), available at http://www.govtech.com/data/Will-GDPR-Rules-Impact-States-and-Localities.html.

chatbot as part of an experiment in conversational understanding and relied on data from users interacting with the chatbot, but then noticed that the chatbot started tweeting about racism, anti-Semitism, dictators and other profane content within 24 hours of being released.[29]

But what about when AI is used by the government to determine which restaurants should be inspected; make judgments on where the next crime might happen; or even decide the length of a prison sentence? AI Now has called for the government to stop the use of certain types of AI until the technology is better understood and made "available for public auditing, testing, and review, and subject to accountability standards."[30] Microsoft recently declined a government request for facial recognition software due to human rights concerns.[31] Further analysis is presented below on how to address government use of AI.

A. Lack of Relevant Case Law

Thus far, most of the rules implicated by AI are not particular to AI at all. Rather, they are existing and sometimes longstanding privacy, cybersecurity, unfair and deceptive acts and practices, due process, and health and safety rules that cover technologies that now happen to concentrate on AI. These include rules about holding, using and protecting personal data, guidance on how to manage the risks caused by financial algorithms, and protections against discrimination.

There are cases that reference the risk of government adoption of AI. For example, in the unpublished California state appellate case *County of Riverside v. Perone,* the court referenced the challenges in a government agency adopting AI to generate recruitment lists by matching job qualifications supplied by the requesting department with skills and education found on resumes that had been scanned into the system.[32] Or on the federal side, amicus briefs were filed in the U.S. Supreme Court case *Gill v. Whitford,* arguing that the increasing adoption of machine learning for analyzing voter data and behavior poses the threat of increasingly precise and discriminatory gerrymandering.[33]

29. Ingrid Angulo, *Facebook and YouTube Should Have Learned from Microsoft's Racist Chatbot*, CNBC (March 17, 2018), available at https://www.cnbc.com/2018/03/17/facebook-and-youtube-should-learn-from-microsoft-tay-racist-chatbot.html.

30. *See, e.g.,* Alex Campolo, et. al., *AI Now 2017 Report* (2017), available at https://ainowinstitute.org/AI_Now_2017_Report.pdf.

31. Joseph Menn, *Microsoft Turned Down Facial-Recognition Sales on Human Rights Concerns*, (April 19, 2019), available at https://www.reuters.com/article/us-microsoft-ai-idUSKCN1RS2FV .

32. *Cty. of Riverside v. Perone,* 2006 WL 245319 (Cal. Ct. App. February 2, 2006).

33. *Gill v. Whitford,* Brief of Amici Curiae Political Science Professors in Support of Appellees and Affirmance, 2017 WL 4311101 (S.Ct. September 5, 2017).

However, there are limited cases in which a citizen or company has directly challenged the government's use of AI.[34] The dearth of litigation may be that the use of AI is still preliminary and minimal (though increasing) in the regulatory context. It may be that lawyers are still grappling with how to apply old laws and principles to new technology. Whatever the reason, it is likely that such challenges will increase in the future. As governments adopt these technologies, several risks to the regulatory process emerge and will be essential for private parties to consider when facing potential unlawful and unconstitutional administrative actions. Namely, private parties may face situations where they challenge government use of AI in the regulatory context on the basis of principles of transparency, due process, nondelegation, and nondiscrimination. This article focuses primarily on challenging government use of AI on the basis of algorithmic bias.

B. Potential Legal Framework to Address Discriminatory Government Action

There are laws that could provide some protection against government action that may be based on algorithmic bias. Current anti-discrimination laws in sectors like education, housing, and employment prohibit both intentional discrimination—called "disparate treatment"—as well as unintentional "disparate impact," which happens when neutral-sounding rules disproportionately affect a legally protected group (e.g., on the basis of sex, age, disability, race, etc.).

Since the civil rights movement, a body of law has emerged around claims that an institution intentionally treated a protected class of individuals less favorably than other individuals. In 1971, the term "disparate impact" was first used in the Supreme Court case *Griggs v. Duke Power Company.*[35] The Court ruled that, under Title VII of the Civil Rights Act, it was illegal for the company to use intelligence test scores and high school diplomas—factors which were shown to disproportionately favor white applicants and substantially disqualify people of color—to make hiring or promotion decisions, whether or not the company intended the tests to discriminate. A key aspect of the *Griggs* decision was that the power company could not prove their intelligence tests or diploma requirements were actually relevant to the jobs they were hiring for.

34. As one of the few examples of challenge of government use of AI, defendant Eric Loomis was found guilty for his role in a drive-by shooting. During intake, Loomis answered a series of questions that were then entered into Compas, a risk-assessment tool developed by a privately held company and used by the Wisconsin Department of Corrections. The trial judge gave Loomis a long sentence partially because of the "high risk" score the defendant received from this risk-assessment tool. Loomis challenged his sentence, because he was not allowed to assess the algorithm. The state supreme court ruled against Loomis, reasoning that knowledge of the algorithm's output was a sufficient level of transparency. *Wisconsin v. Loomis*, 371 Wis.2d 235 (July 13, 2016).

35. *Griggs v. Duke Power Co.*, 401 U.S. 424 (1971).

More recently, the government and other plaintiffs have advanced disparate impact claims that focus more on the effect instead of the intent of lending policies. Recently, the Supreme Court's decision in *Texas Department of Housing and Community Affairs v. Inclusive Communities Project* affirmed the use of the disparate impact theory.[36] The Inclusive Communities Project had used a statistical analysis of housing patterns to show that a tax credit program effectively segregated Texans by race.

The fundamental validation of disparate impact theory by the Court in the *Inclusive Communities* case remains a wake-up call for technology and government agencies. An algorithm that inadvertently disadvantages a protected class continues to have the potential to create due process concerns.

C. Assertion of Disparate Impact Claims

In asserting a disparate impact theory, plaintiffs must prove that they were disproportionately and negatively affected by a government policy or practice.[37] Where a disparate impact is shown, the government may defend itself either by challenging the plaintiff's evidence (usually by attacking the statistics used to demonstrate the disparate impact) or by proving that the AI policy is necessary to achieve a valid interest. If the government can't prove that, then a plaintiff's claim of disparate impact must prevail.

If the government could demonstrate that the AI in question has a demonstrable relationship to the requirements, or a "business necessity," the plaintiff can still win by providing that the government refuses to adopt an alternative practice with a less discriminatory effect. One route would be to argue that alternative methods were equally effective without being discriminatory. Since algorithms are proprietary and frequently protected under non-disclosure agreements, organizations that use them, including government agencies, may not have the legal right to conduct independent testing. This would force the government to either admit it considered no alternatives or force an examination of the algorithm, which is currently quite challenging. The ability to audit an algorithm would answer some questions about bias, but there is a group of algorithms that move beyond our current abilities to analyze them. Artificial neural networks is one example. They are fed huge amounts of data and, through a process of breaking it down into much smaller components and searching for patterns, essentially come up with their own algorithms which generally are incomprehensible to humans.

36. *Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Communities Project*, 135 S.Ct. 2507 (2015).

37. Some civil rights laws, such as Title VI of the Civil Rights Act of 1964, do not contain disparate impact provisions creating a private right of action, although the federal government may still pursue disparate impact claims under these laws. *See Alexander v. Sandoval*, 532 U.S. 275 (2001).

D. Additional Points Regarding Discrimination and Valuing Diversity

In putting forth a disparate impact case, or similar legal theory, there are many additional points to be made to reinforce the theory of discriminatory treatment—two key points are included below. First, a plaintiff may draw light to the failure of government to hire and involve diverse stakeholders in developing the algorithm and data inputs. Diversity is crucial to unbiased outcomes, to preventing neutral data points to the ability to compete in the technologically-advanced global marketplace.[38] Government agencies that succeed in diversity and inclusion are those that have a formal diversity hiring strategy, including formal hiring and training programs and a commitment by agency executives to accomplish diversity hiring objectives. Showing a failure to recruit diverse employees, beyond the basic demonstration of a failure to create a high-functioning organization, would highlight a fertile environment for impartial algorithms.

Further, a plaintiff may consider showing the government failure to engage human oversight to test and check algorithmic decisions. These individuals should document and validate data inputs as a part of their process. What this means is that employees should be assigned to review data sets to ensure the data is fair and accurate, establish best practices for auditing algorithmic decision-making, and include specific guidance on addressing questions of disparate impact.

VI. CONCLUSION

AI has a tremendous potential to positively impact all manner of life, from credit to employment to health care and safety. But addressing the risks associated with the technology needs to be a priority, especially for government agencies. To the extent risks are not addressed proactively and/or comprehensively, individuals and companies may have opportunities to assert novel legal theories to preserve their rights in the face of discriminatory treatment, and other illicit treatment, through government use of AI.

38. Karen Hao, *AI's White Guy Problem Isn't Going Away*, (April 17, 2019) available at https://www.technologyreview.com/s/613320/ais-white-guy-problem-isnt-going-away/?utm_campaign=the_download.unpaid.engagement&utm_source=hs_email&utm_medium=email&utm_content=71836962&_hsenc=p2ANqtz--hGziUG3-VQAJT0mVc8Fbk2VKmBmPnbc2I_tD8Fmk2EcjJqUMauo2s6hyjunl_XbgkANIv4JUYd2Yt6JTsPW0xinh9qbY0WDJbB0hhkgVHrbldpWc&_hsmi=71836962.