

Should an Algorithm Make Our Decisions?

CASE STUDY

Darnell Gates of Philadelphia had been jailed for running a car into a house in 2013 and later for violently threatening his former domestic partner. When he was released in 2018, he was initially required to visit a probation officer once a week because he had been identified as “high risk” by a computer algorithm. Gates’s probation office visits were eventually stretched to every two weeks and then once a month, but conversations with probation officers remained impersonal and perfunctory, with the officers rarely taking the time to understand Gates’s rehabilitation. What Gates didn’t know was that a computer algorithm developed by a University of Pennsylvania professor had made the “high risk” determination that governed his treatment.

This algorithm is one of many now being used to make decisions about people’s lives in the United States and Europe. Predictive algorithms are being used to determine prison sentences, probation rules, and police patrols. It is often not clear how these automated systems are making their decisions. Many countries and states have few rules requiring disclosure of algorithms’ formulae. And even if governments provide explanation of how the systems arrive at their decisions, the algorithms are often too difficult for a layperson to understand.

The Rotterdam city government has used such algorithms to identify risks for welfare and tax fraud. A program called System Risk Indication scans data from various government authorities to flag people who may be claiming unemployment benefits when they are working, or receiving a housing subsidy for living alone when they are actually living with several others. According to the Ministry of Social Affairs and Employment, which runs this program, data examined include income, education, property, rent, debt, car ownership, home address, and the welfare benefits received for housing, health care, and children. The algorithm produces “risk reports” on individuals who should be investigated. When the system was used most recently in Rotterdam, it produced 1,263 risk reports in two neighborhoods. People who are on the list are not told how the algorithm is making its decisions and whether they are on the register, nor are they shown ways to appeal. An October 2019 report by the United Nations

special rapporteur on extreme poverty criticized the Dutch system for creating a “digital welfare state” that turns crucial decisions about people’s lives over to machines.

Rotterdam citizens, along with privacy rights groups, civil rights lawyers, and the largest national labor union rallied against System Risk Indication. A district court ordered an immediate halt to the risk algorithm’s use. The court stated that the System Risk Indication program lacked privacy safeguards and that the government was insufficiently transparent about how it worked. The decision can be appealed.

Bristol, England is using an algorithm to compensate for tight budgets that reduced social services. A team that includes representatives from the police and children’s services meets weekly to review results from an algorithm that tries to identify youths in the city who are most at risk for crime and children who are most in need. In 2019, Bristol introduced a software program that creates a risk score based on data extracted from police reports, social benefits, and other government records. The risk score takes into account crime, school attendance, and housing data, known ties to others with high risk scores, and whether a youth’s parents were involved in a domestic incident. The scores fluctuate, depending on how recently a youth had an incident such as a school suspension.

There is evidence that Bristol’s algorithm is identifying the right people, and there are human decision makers governing its use. Charlene Richardson and Desmond Brown, two city workers who are responsible for aiding young people flagged by the software, acknowledge that the computer doesn’t always get it right, so they do not rely on it entirely. The city government has been open with the public about the program. The government has posted some details online and staged community events. However, opponents believe the program isn’t fully transparent. Young people and their parents do not know if they are on the list, and have no way to contest their inclusion.

Studies of algorithms in credit scoring, hiring, policing, and health care have found that poorly designed algorithms can reinforce racial and gender

biases. According to Solon Barocas, an assistant professor at Cornell University and principal researcher at Microsoft Research, algorithms can appear to be very data-driven, but there are subjective decisions that go into setting up the problem in the first place.

For example, a study by Ziad Obermeyer, a health policy researcher at the University of California, Berkeley, and colleagues, published in October 2019 in the journal *Science*, examined an algorithm widely used in hospitals to establish priorities for patient care. The study found that black patients were less likely than white patients to receive extra medical help, despite being sicker. The algorithm has been used in the Impact Pro program from Optum, United Health Group's health services arm, to identify patients with heart disease, diabetes, and other chronic illnesses who could benefit from more attention from nurses, pharmacists, and case workers in managing their prescriptions, scheduling doctor visits, and monitoring their overall health.

The algorithm assigned healthier white patients the same ranking as black patients who had one more chronic illness as well as poorer laboratory results and vital signs. To identify people with the greatest medical needs, the algorithm looked at patients' medical histories and how much was spent treating them, and then predicted who was likely to have the highest costs in the future. The algorithm used costs to rank patients.

The algorithm wasn't intentionally racist—in fact, it specifically excluded race. Instead, to identify patients who would benefit from more medical support, the algorithm used a seemingly race-blind metric: how much patients would cost the health-care system in the future. But cost isn't a race-neutral measure of health-care need. Black patients incurred about \$1,800 less in medical costs per year than white patients with the same number of chronic conditions; thus the algorithm scored white patients as equally at risk of future health problems as black patients who had many more diseases.

When the researchers searched for the source of the scores, they discovered that Impact Pro was using bills and insurance payouts as indicators for a person's overall health. However, health care costs tend to be lower for black patients, regardless of their actual well-being. Compared with white patients, many black patients live farther from their hospitals, making it harder to visit them regularly. They also tend to have less flexible job schedules and more child care responsibilities.

As a result, black patients with the highest risk scores had higher numbers of serious chronic conditions than white patients with the same scores, including cancer and diabetes, the researchers reported. Compared with white patients with the same risk scores, black patients also had higher blood pressure and cholesterol levels, more severe diabetes, and worse kidney function.

By using medical records, laboratory results, and vital signs of the same set of patients, the researchers found that black patients were sicker than white patients who had a similar predicted cost. By revising the algorithm to predict the number of chronic illnesses that a patient will likely experience in a given year—rather than the cost of treating those illnesses—the researchers were able to reduce the racial disparity by 84 percent. When the researchers sent their results to Optum, it replicated their findings and committed to correcting its model.

Civil rights lawyers, labor unions, community organizers, and some government agencies are trying to find ways to push back against the growing dependence on automated systems that remove humans from the decision making process. Media Mobilizing Project in Philadelphia and MediaJustice in Oakland, California have compiled a nationwide database of prediction algorithms. Community Justice Exchange, a national organization supporting community organizers, has issued a guide advising organizers on how to confront the use of algorithms. Idaho passed a law in 2019 specifying that the methods and data used in bail algorithms should be publicly available so that the general public can understand how they work.

Concerns about biases have always been present whenever people make important decisions. What's new is the much larger scale at which we rely on algorithms in automated systems to help us decide, and even to take over the decision making for us. Algorithms are helpful in making predictions that will help guide decision makers, but decision making requires much more. Good decision making requires bringing together and reconciling multiple points of view and the ability to explain why a particular path was chosen.

As automated systems increasingly shift from predictions to decisions, focusing on the fairness of algorithms is not enough because their output is just one of the inputs for a human decision maker. One must also look at how human decision makers interpret and integrate the output from algorithms and under what conditions they would deviate from

an algorithmic recommendation. Which aspects of a decision process should be handled by an algorithm and which by a human decision maker to obtain fair and reliable outcomes?

Sources: Cade Metz and Adam Satariano, "An Algorithm That Grants Freedom, or Takes It Away," *New York Times*, February 6, 2020; Irving Wladawsky-Berger, "The Coming Era of Decision Machines," *Wall Street Journal*, March 27, 2020; Michael Price, "Hospital 'Risk Scores' Prioritize White Patients," *Science*, October 24, 2019; Melanie Evans and Anna Wilde Matthews, "Researchers Find Racial Bias in Hospital Algorithm," *Wall Street Journal*, October 25, 2019.

CASE STUDY QUESTIONS

- 12-13** What are the problems in using algorithms and automated systems for decision making?
- 12-14** What management, organizational, and technology factors have contributed to the problem?
- 12-15** Should automated systems be used to make decisions? Explain your answer.

Chapter 12 References

- Breuker, Dominic, Martin Matzner, Patrick Delfmann, and Jörg Becker. "Comprehensible Predictive Models for Business Processes." *MIS Quarterly* 40, No. 4 (September 2016).
- Brynjolfsson, Erik, Tomer Geva, and Shachar Reichman. "Crowd-Squared: Amplifying the Predictive Power of Search Trend Data." *MIS Quarterly* 40, No. 4 (December 2016).
- Davenport, Thomas H. "Analytics 3.0." *Harvard Business Review* (December 2013).
- Davenport, Thomas H., and Jeanne G. Harris. *Competing on Analytics: The New Science of Winning: Updated, with a New Introduction*. Boston: Harvard Business Review Press (2017).
- Day, George S., and Paul J. H. Schoemaker. "How Vigilant Companies Gain an Edge in Turbulent Times." *MIT Sloan Management Review* 61 No. 2 (Winter 2020).
- Dietvorst, Berkeley J. "When People Don't Trust Algorithms." *MIT Sloan Management Review* (July 5, 2017).
- Grau, Jeffrey. "How Retailers Are Leveraging 'Big Data' to Personalize E-commerce." *eMarketer* (2012).
- Haglu, Andrei and Julian Wright. "When Data Creates Competitive Advantage." *Harvard Business Review* (January–February 2020).
- Hardin, Andrew, Clayton A. Looney, and Gregory D. Moody. "Assessing the Credibility of Decisional Guidance Delivered by Information Systems." *Journal of Management Information Systems* 34, No. 4 (2017).
- Harris, Michael, and Bill Tayler. "Don't Let Metrics Undermine Your Business." *Harvard Business Review* (September–October 2019).
- Hernandez, Morela, Roshni Raveendhran, Elizabeth Weingarten, and Michaela Barnardt. "How Algorithms Can Diversify the Startup Pool." *MIT Sloan Management Review* 61 No. 1 (Fall 2019).
- Hu, Yuheng, Anbang Xu, Yili Hong, David Gal, Vibha Sinha, and Rama Akkiraju. "Generating Business Intelligence Through Social Media Analytics: Measuring Brand Personality with Consumer-, Employee-, and Firm-Generated Content." *Journal of Management Information Systems* 36 No. 3 (2019).
- Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux (2011).
- Kaplan, Robert S., and David P. Norton. "The Balanced Scorecard: Measures That Drive Performance." *Harvard Business Review* (January–February 1992).
- Luca, Michael, Jon Kleinberg, and Sendhil Mullainathan. "Algorithms Need Managers, Too." *Harvard Business Review* (January–February 2016).
- Marchand, Donald A., and Joe Peppard. "Why IT Fumbles Analytics." *Harvard Business Review* (January–February 2013).
- Martens, David, Foster Provost, Jessica Clark, and Enric Janqué de Fortuny. "Mining Massive Fine-Grained Behavior Data to Improve Predictive Analytics." *MIS Quarterly* 40, No. 4 (December 2016).
- Mintzberg, Henry. "Managerial Work: Analysis from Observation." *Management Science* 18 (October 1971).
- Porter, Michael E., and Nitin Nohria. "How CEOs Manage Time." *Harvard Business Review* (July–August 2018).
- Sharda, Ramesh, Dursan Delen, and Efraim Turban. *Analytics, Data Science, & Artificial Intelligence: Systems for Decision Support*, 11e. New York: Pearson (2020).
- Shi, Donghui, Jian Guan, Jozef Zurada, and Andrew Manikas. "A Data-Mining Approach to Identification of Risk Factors in Safety Management Systems." *Journal of Management Information Systems* 34, No. 4 (2017).
- Simchi-Levi, David. "The New Frontier of Price Optimization." *MIT Sloan Management Review* (Fall 2017).
- Simon, H. A. *The New Science of Management Decision*. New York: Harper & Row (1960).
- Tan, Aaron. "Singapore's JTC taps GIS Platform in Punggol Digital District." *Computer Weekly* (March 23, 2020).
- Tversky, Amos, and Daniel Kahneman. "Rational Choice and the Framing of Decisions." *Journal of Business* (1986).
- Venna, Srivani. "SNCF Uses IBM IoT For Improved Rail Safety." *Railway Technology*, February 19, 2017.
- Zarembski, Allan. "Better Railroading through Big Data." *Railway Age*, April 24, 2018.