

AI Judges and Juries

Artificial intelligence is changing the legal industry.

WHEN THE HEAD of the U.S. Supreme Court says artificial intelligence (AI) is having a significant impact on how the legal system in this country works, you pay attention. That's exactly what happened when Chief Justice John Roberts was asked the following question:

"Can you foresee a day when smart machines, driven with artificial intelligences, will assist with courtroom fact-finding or, more controversially even, judicial decision-making?"

His answer startled the audience.

"It's a day that's here and it's putting a significant strain on how the judiciary goes about doing things," he said, as reported by *The New York Times*.

In the last decade, the field of AI has experienced a renaissance. The field was long in the grip of an "AI winter," in which progress and funding dried up for decades, but technological breakthroughs in AI's power and accuracy changed all that. Today, giants like Google, Microsoft, and Amazon rely on AI to power their current and future profit centers.

Yet AI isn't just affecting tech giants and cutting-edge startups; it is transforming one of the oldest disciplines on the planet: the application of the law.

AI is already used to analyze documents and data during the legal discovery process, thanks to its ability to parse through millions of words faster (and more cheaply) than human beings. That alone could automate away or completely change the almost 300,000 paralegal and legal assistant jobs estimated to exist by the U.S. Bureau of Labor Statistics. However, that is just the beginning of AI's potential impact; it also is being used today to influence the outcomes of actual cases.

In one high-profile 2017 case, a man named Eric Loomis was sentenced to six years in prison thanks, in part, to recommendations from AI algorithms.



The system analyzed data about Loomis and made sentencing recommendations to a human judge on the suggested length of Loomis' sentence.

Make no mistake: AI-enhanced courtrooms may be more science fact than science fiction—for better or for worse.

The Predictable, Reliable Choice?

Artificial intelligence holds some promise for the world of legal decisions.

In Canada, Randy Goebel, a professor in the computer science department of the University of Alberta working in conjunction with Japanese researchers, developed an algorithm that can pass the Japanese bar exam. Now, the team is working to develop AI that can "weigh contradicting legal evidence, rule on cases, and predict the outcomes of future trials," according to Canadian broadcaster CBC. The goal is to use machines to help humans make better legal decisions.

This is already being attempted in U.S. courtrooms. In the Loomis case, AI was used to evaluate individual defendants. The algorithm used was created and built into software called Compas

by a company called Northpointe Inc. The algorithm indicated Loomis had "a high risk of violence, high risk of recidivism, [and] high pretrial risk." This influenced the six-year sentence he received, though the sentencing judges were advised to take note of the algorithm's limitations.

Criminal justice algorithms like the one in the Loomis case use personal data such as age, sex, and employment history to recommend sentencing, reports the Electronic Privacy Information Center (EPIC). The technology is relatively common in the U.S. legal system.

"Criminal justice algorithms are used across the country, but the specific tools differ by state or even county," says EPIC.

The case for using AI-based systems to assist in the legal process hinges on the perceived ability of machines to be more impartial than humans. "Humans can be swayed by emotion. Humans can be convinced. Humans get tired or have a bad day," says Tracy Greenwood, an expert in e-discovery, the process of using machines to per-

form legal discovery work faster and more accurately than humans.

“In a high crime city, a judge might start to hand out harsher sentences towards the upper end of the sentencing guidelines. In court, if a judge does not like one of the lawyers, that can affect the judge’s opinion,” says Greenwood.

The argument is that machines could potentially analyze facts and influence judgments dispassionately, without human bias, irrationality, or mistakes creeping into the process.

For instance, the Japanese bar exam AI developed by Goebel and his team is now considered “a world leader in the field,” according to CBC. It succeeded on the exam where at least one human failed: one of Goebel’s colleagues failed the Japanese bar exam.

Human fallibility is not an isolated problem in the legal field. According to an investigation by U.K.-based newspaper *The Guardian*, local, state, and federal courts in the U.S. are rife with judges who “routinely hide their connections to litigants and their lawyers.” The investigation learned that oversight bodies found wrongdoing and issued disciplinary action in nearly half (47%) of complaints about judge conflict of interest they investigated.

However, oversight bodies rarely look into complaints at all—90% of over 37,000 complaints investigated were dismissed by state court authorities “without conducting any substantive inquiry,” according to the investigation.

Conflict of interest is not the only human bias that plagues the U.S. legal system; racial bias, explicit or implicit, also is common.

“Minorities have less access to the courts to begin with, and tend to have worse outcomes due to systemic factors limiting their quality of representation, and subconscious or conscious bias,” says Oliver Pulleyblank, founder of Vancouver, British Columbia-based legal firm Pulleyblank Law.

Intelligent machines, however, do not carry the same baggage. Acting as dispassionate arbiters looking at “just the facts,” machines hold the potential to influence the legal decision-making process in a more consistent, standardized way than humans do.

The benefits would be significant.

“To introduce a system with much greater certainty and predictability

would open up the law to many more people,” says Pulleyblank. The high cost and uncertain outcomes of cases discourage many from pursuing valid legal action.

“Very few people can afford to litigate matters,” says Pulleyblank, “even those who can generally shouldn’t, because legal victories are so often hollow after all the expenses have been paid.”

However, when you look more deeply at machine-assisted legal decisions, you find they may not be as impartial or consistent as they seem.

“Unbiased” Machines Created by Biased Humans

In the Loomis algorithm-assisted case, the defendant claimed the algorithm’s report violated his right to due process, but there was no way to examine how the report was generated; the company that produces the Compas software containing the algorithm, Northpointe, keeps its workings under wraps.

“The key to our product is the algorithms, and they’re proprietary. We’ve created them, and we don’t release them because it’s certainly a core piece of our business,” Northpointe executives were reported as saying by *The New York Times*.

This is the so-called “black box” problem that haunts the field of artificial intelligence.

Algorithms are applied to massive datasets. The algorithms produce results based upon their “secret sauce”—how they use the data. Giving up the secret sauce of an algorithm is akin to giving up your entire competitive advantage.

The result? Most systems that use AI are completely opaque to anyone

“To introduce a system with much greater certainty and predictability would open up the law to many more people.”

except their creators. We are unable to determine why an algorithm produced a specific output, recommendation, or assessment.

This is a major problem when it comes to using machines as judge and jury: because we lack even the most basic understanding of how the algorithms work, we cannot know if they are producing poor results until *after* the damage is done.

ProPublica, an “independent, non-profit newsroom that produces investigative journalism with moral force,” according to its website, studied the “risk scores,” assessments created by Northpointe’s algorithm, of 7,000 people who were arrested in Broward County, FL. These scores are used to determine release dates and bails in courtrooms, as they purportedly predict the defendant’s likelihood to commit crime again.

As it turns out, these algorithms may be biased.

In the cases investigated, ProPublica says the algorithms wrongly labeled black defendants as future criminals at a rate nearly twice that of white defendants (who were mislabeled as “low risk” more often than black defendants).

Because the algorithms do not operate transparently, it is difficult to tell if this was an assessment error, or if the algorithms were coded with rules that reflect the biases of the people who created them.

In addition to bias, the algorithms’ predictions just are not that accurate.

“Only 20% of the people predicted to commit violent crimes actually went on to do so,” says ProPublica. Fewer violent crimes committed is a good thing, but based on this assessment, decisions were made that treated 80% of defendants as likely violent criminals when they were not.

Critics claim that algorithms need to be far more transparent before they can be relied on to influence legal decisions.

Even then, another huge problem with having AI take on a larger role in the legal system is that there is no guarantee machines can handle the nuances of the law effectively, says Pulleyblank.

“Many legal problems require judges to balance distinct interests against

each other,” he says. He cites the example of a sexual assault victim bringing a case against their attacker. The judge is required to balance the victim’s need for privacy with the principle that justice should take place in the open for all to see. There’s no easy answer, but the decision to publish the victim’s name or keep the proceedings behind closed doors is one a judge has to make—and one that has major effects on the case.

“What it depends on is not ‘the law,’” says Pulleyblank. “There is no clear legal answer to how those values will be balanced in any given case. Rather, it depends on the judge.”

These types of contextual considerations crop up constantly in all manner of cases. “Machines are good at identifying what has been tried and what has not been tried, but they lack judgment,” says Greenwood. He says machines may produce consistent results, but lack other critical skills to ensure justice is served. “A machine will not lecture a defendant in a criminal case and tell him to get his life together.”

Pulleyblank agrees that making the law more “predictable” using machines may cause more problems than it solves. “Whenever you seek to make the law more predictable, you risk sacrificing fairness,” he says.

In ProPublica’s investigation, the algorithm assessed two defendants. One was a seasoned criminal; the other a young girl with a prior misdemeanor. Both had stolen items of the same value, but the machine failed to contextualize the fact that the young girl had stolen a bicycle and had no serious criminal record. She was deemed a likely repeat offender, just like the career criminal. To the machine, these two people had both committed crimes and had past charges. It failed to contextualize; as a result, the algorithm used in this case got the situation very wrong.

Yet introducing context and circumstance inherently reduces the predictability and consistency of the law’s application, so the balance between machine predictability and human judgment is a tenuous one.

“This is the order versus fairness dichotomy that has long been the subject of legal thought,” says Pulleyblank.

This leads both Pulleyblank and Greenwood to the same conclusion: machines probably will come to heavily

“In order to allow predictable non-human judicial decisions, the law would have to change in a fairly fundamental way.”

assist humans in the legal profession. The industry will transform as a result, but to completely replace humans in the legal process would likely require changing the law itself.

“In order to allow predictable non-human judicial decisions, the law would have to change in a fairly fundamental way,” says Pulleyblank, “and if the law does not change, there is simply too much discretion inherent in the law as it exists for the public to accept that discretion being exercised by machines.”

While machines might have superior predictive power, humans will issue the final verdict on their use. **■**

Further Reading

Liptak, A.

Sent to Prison by a Software Program’s Secret Algorithms, *The New York Times*, May 1, 2017, <https://www.nytimes.com/2017/05/01/us/politics/sent-to-prison-by-a-software-programs-secret-algorithms.html>

Angwin, J.

Machine Bias, *ProPublica*, May 26, 2016 <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

Algorithms in the Criminal Justice System, *Electronic Privacy Information Center*, <https://epic.org/algorithmic-transparency/crim-justice/>

Snowden, W.

Robot judges? Edmonton research crafting artificial intelligence for courts, *CBC*, Sept. 19, 2017, <http://www.cbc.ca/news/canada/edmonton/legal-artificial-intelligence-alberta-japan-1.4296763>

Logan Kugler is a freelance technology writer based in Tampa, FL, USA. He has written for over 60 major publications.

© 2018 ACM 0001-0782/18/12 \$15.00

Milestones

Estrin Awarded MacArthur ‘Genius Grant’

Cornell Tech computer science professor Deborah Estrin was among the 25 people named 2018 MacArthur Fellows. The MacArthur Foundation said Estrin was chosen for “designing open source platforms that leverage mobile devices and data to address socio-technological challenges such as personal health management.”

Estrin, 58, earned her bachelor of science degree from the University of California at Berkeley, and master of science and doctorate degrees from the Massachusetts Institute of Technology. She taught at the University of Southern California and the University of California at Los Angeles prior to joining the faculty of Cornell Tech in 2012, where she is a professor in the Department of Computer Science and associate dean.

Estrin was named an ACM fellow in 2000 “for significant contributions to the design of scalable Internet protocols, and for service to the networking community.” In 2006, she received the ACM Athena Lecturer Award, bestowed annually in celebration of women researchers who have made fundamental contributions to computer science.

The MacArthur Foundation said Estrin “has demonstrated a remarkable ability to anticipate the applicability of technological advances to a variety of fields. She made fundamental contributions to improving the scalability and broader utility of the emerging Internet through her work on network routing (the process that determines how data is forwarded from source to destination). She then went on to build the foundational protocols for wireless sensor networks—that is, connectivity among distributed autonomous sensors that record conditions in a specific environment. “

“I was and remain very humbled and grateful,” said Estrin. “I feel a sense of commitment to do good by it, and to live up to it.”