As governments adopt artificial intelligence, there's little oversight and lots of danger

April 18 2019, by James Hendler

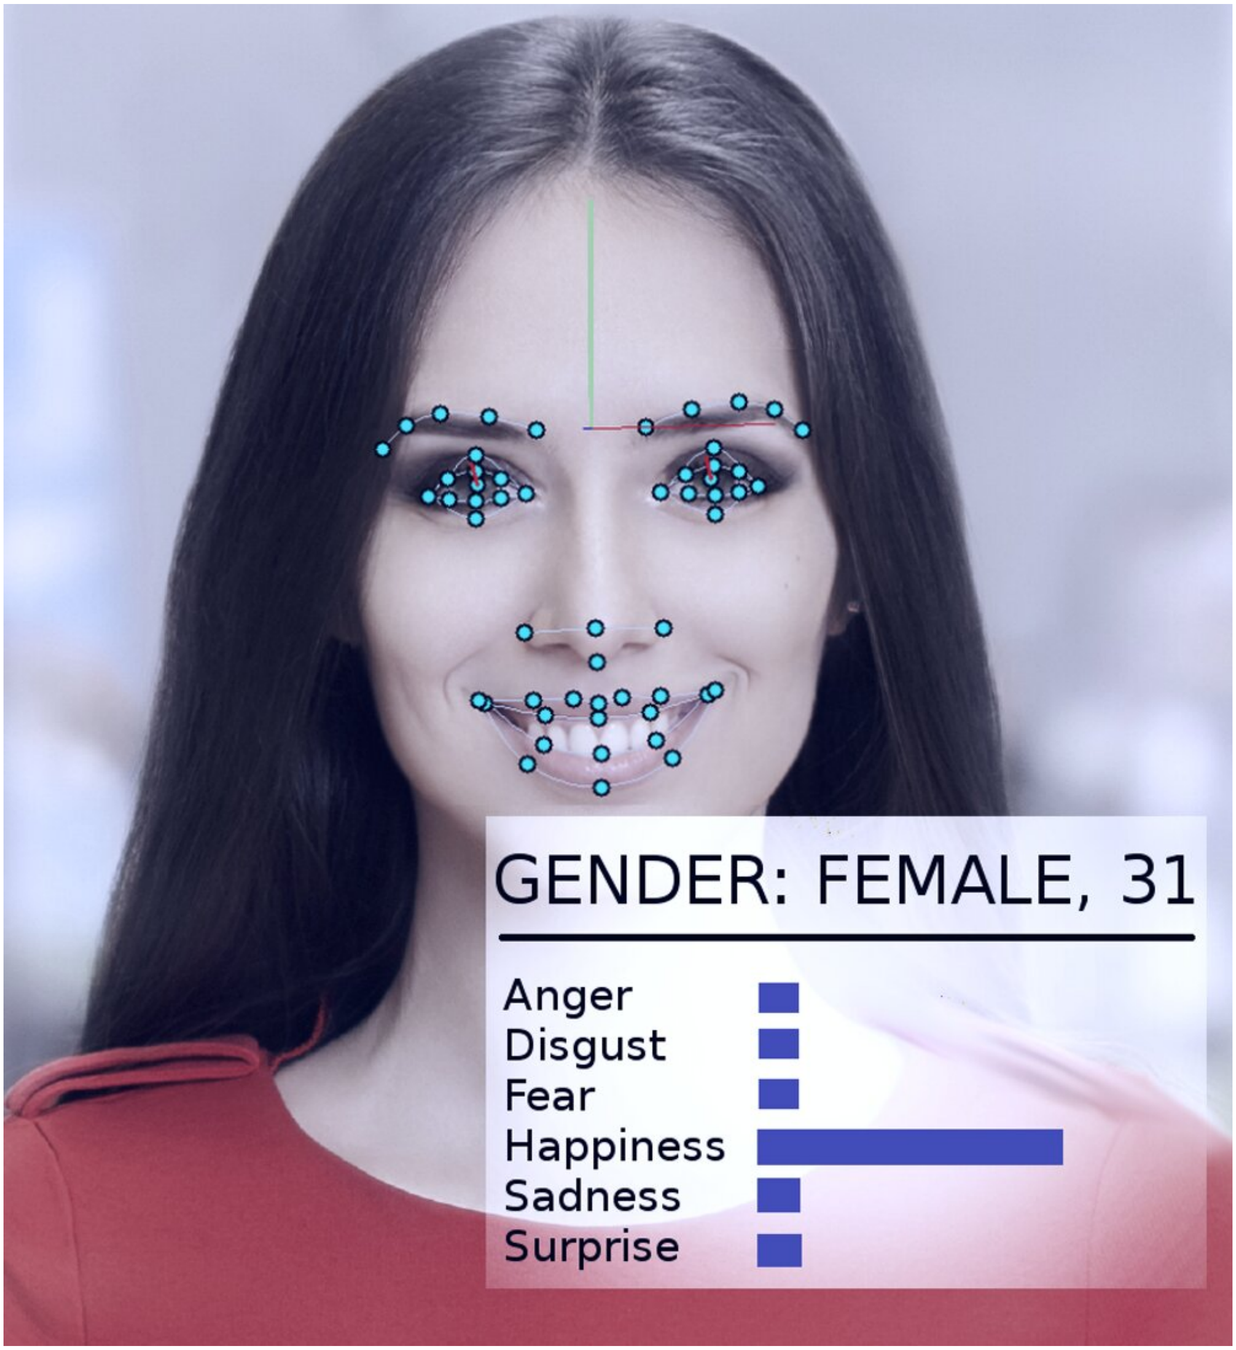

A face tracking and analysis system takes a look at a woman's face. Credit: Abyssus/Wikimedia Commons, CC BY-SA

Artificial intelligence systems can – if properly used – help make government more effective and responsive, improving the lives of citizens. Improperly used, however, the dystopian visions of George Orwell's "1984" become more realistic.

On their own and urged by a new presidential executive order, governments across the U.S., including state and federal agencies, are exploring ways to use AI technologies.

As an AI researcher for more than 40 years, who has been a consultant or participant in many government projects, I believe it's worth noting that sometimes they've done it well – and other times not quite so well. The potential harms and benefits are significant.

An early success

In 2015, the U.S. Department of Homeland Security developed an AI system called "Emma," a chatbot that can answer questions posed to it in regular English, without needing to know what "her" introductory website calls "government speak" – all the official terms and acronyms used in agency documents.

By late 2016, DHS reported that Emma was already helping to answer nearly a half-million questions per month, allowing DHS to handle many more inquiries than it had previously, and letting human employees spend more time helping people with more complicated queries that are beyond Emma's abilities. This sort of conversation-automating artificial intelligence has now been used by other government agencies, in cities and countries around the world.

Flint's water

A more complicated example of how governments could aptly apply AI can be seen in Flint, Michigan. As the local and state governments struggled to combat lead contamination in the city's drinking water, it became clear that they would need to replace the city's remaining lead water pipes. However, the city's records were incomplete, and it was going to be extremely expensive to dig up all the city's pipes to see if they were lead or copper.

Instead, computer scientists and government employees collaborated to analyze a wide range of data about each of 55,000 properties in the city, including how old the home was, to calculate the likelihood it was served by lead pipes. Before the system was used, 80% of the pipes dug up needed to be replaced, which meant 20% of the time, money and effort was being wasted on pipes that didn't need replacing.

The AI system helped engineers focus on high-risk properties, identifying a set of properties most likely to need pipe replacements. When city inspectors visited to verify the situation, the algorithm was right 70% of the time. That promised to save enormous amounts of money and speed up the pipe replacement process.

However, local politics got in the way. Many members of the public didn't understand why the system was identifying the homes it did, and objected, saying the AI method was unfairly ignoring their homes. After city officials stopped using the algorithm, only 15% of the pipes dug up were lead. That made the replacement project slower and more costly.

Distressing examples

The problem in Flint was that people didn't understand that AI technology was being used well, and that people were verifying its findings with independent inspections. In part, this was because they didn't trust AI – and in some cases there is good reason for that.

In 2017, I was among a group of more than four dozen AI researchers who sent a letter to the acting secretary of the U.S. Department of Homeland Security. We expressed concerns about a proposal to use automated systems to determine whether a person seeking asylum in the U.S. would become a "positively contributing member of society" or was more likely to be a terrorist threat.

"Simply put," our letter said, "no computational methods can provide reliable or objective assessments of the traits that [DHS] seeks to measure." We explained that machine learning is susceptible to a problem called "data skew," in which the system's ability to predict a characteristic depends in part on how common that characteristic is in the data used to train the system.

So in a database of 300 million Americans, if one in 100 people are, say, of Indian descent, the system will be fairly accurate at identifying them. But if looking at a characteristic shared by just one in a million Americans, there really isn't enough data for the algorithm to make a good analysis.

As the letter explained, "on the scale of the American population and immigration rates, criminal acts are relatively rare, and terrorist acts are extremely rare." Algorithmic analysis is extremely unlikely to identify potential terrorists. Fortunately, our arguments proved convincing. In May 2018, DHS announced it would not use a machine learning algorithm in this way.

Other worrying efforts

Other government uses of AI are being questioned, too – such as attempts at "predictive policing," setting bail amounts and criminal sentences and hiring government workers. All of these have been shown to be susceptible to technical issues and data limitations that can bias their decisions based on race, gender or cultural background.

Other AI technologies such as facial recognition, automated surveillance and mass data collection are raising real concerns about security, privacy, fairness and accuracy in a democratic society.

As Trump's executive order demonstrates, there is significant interest in harnessing AI for its fullest positive potential. But the significant dangers of abuse, misuse and bias – whether intentional or not – have the potential to work against the very principles international democracies have been built upon.

As the use of AI technologies grows, whether originally well-meant or deliberately authoritarian, the potential for abuse increases as well. With no currently existing government-wide oversight in place in the U.S., the best way to avoid these abuses is teaching the public about the appropriate uses of AI by way of conversation between scientists, concerned citizens and public administrators to help determine when and where it is inappropriate to deploy these powerful new tools.

Provided by The Conversation

Citation: As governments adopt artificial intelligence, there's little oversight and lots of danger (2019, April 18) retrieved 3 May 2024 from https://phys.org/news/2019-04-artificial-intelligence-oversight-lots-danger.html