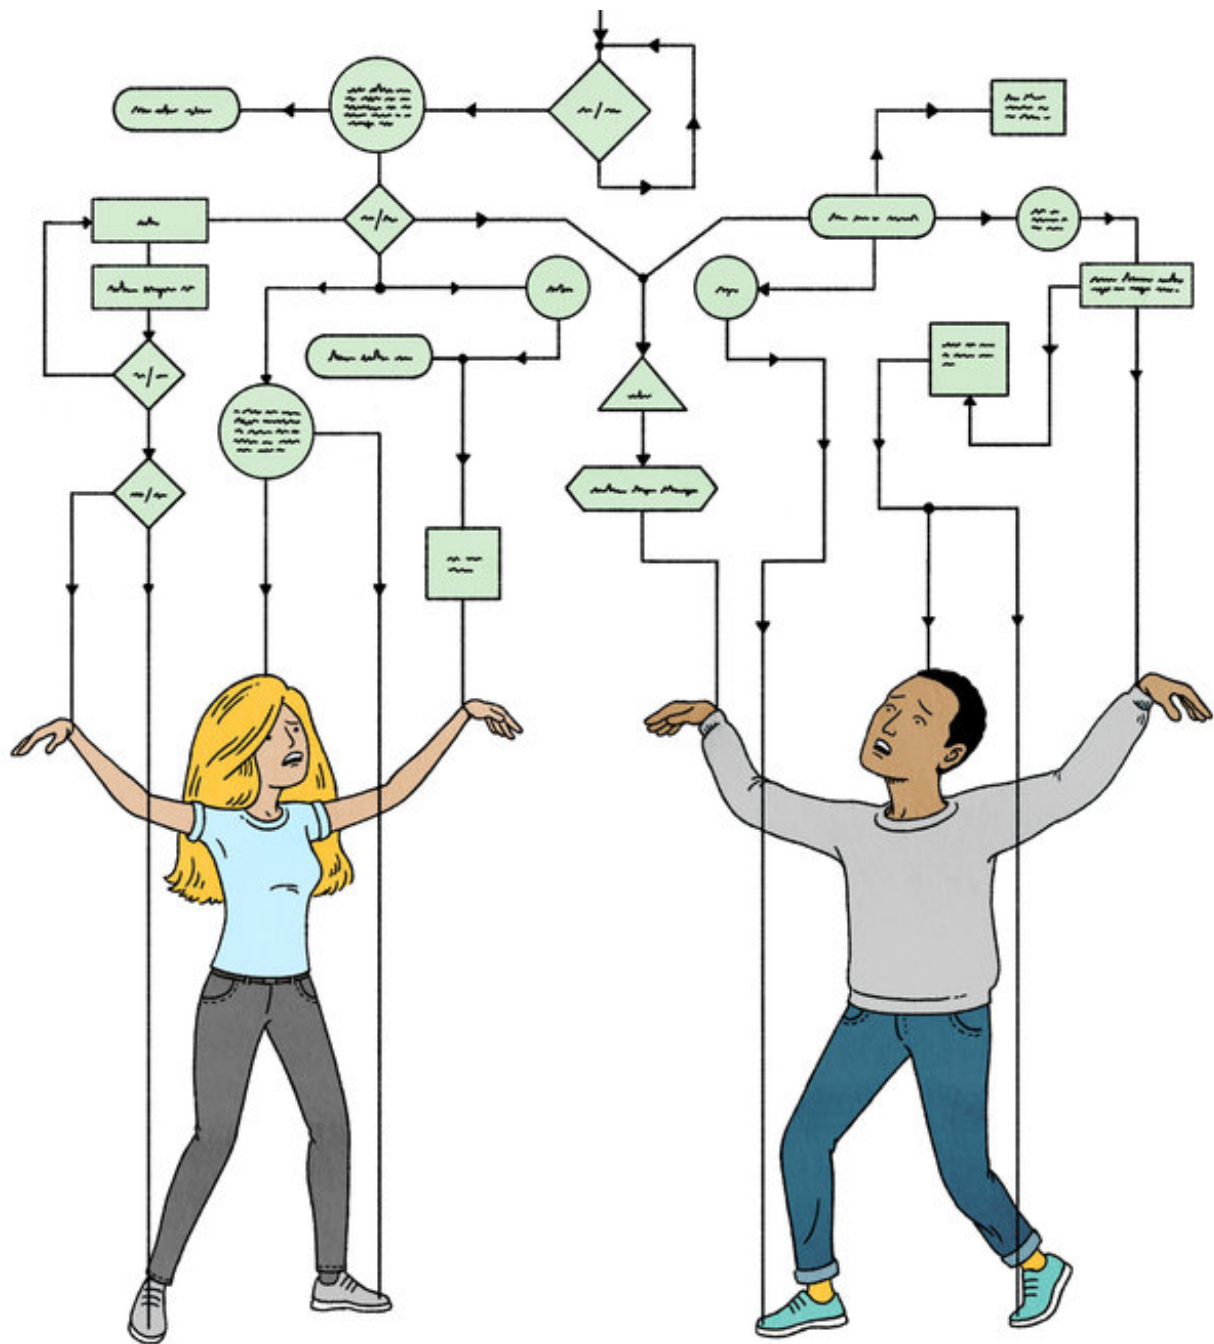


Don't Fear the Robots

nytimes.com/2015/10/25/sunday-review/dont-fear-the-robots.html

October 24, 2015



MENTION artificial intelligence and the image that most quickly springs to mind is an anthropomorphic automaton, a robot. It's a recurring metaphor that peaks at times of torrid technological change and angst about where technology is taking us.

The robot trope is riding high these days, in such books as "The Rise of the Robots," and in movies like "Ex Machina" and "Terminator Genisys." A recent cover of Foreign Affairs carried the headline, "Hi, Robot."

A vision of robots threatening jobs and perhaps humanity itself is fueling dire warnings of future trouble — “summoning the demon,” in the evocative phrase of the technologist-entrepreneur Elon Musk.

Yet the current obsession puts the mechanized cart before the algorithmic horse, steering attention away from the here-and-now promise and peril of A.I.

Artificial intelligence is already all around us, but it’s mostly software. Google search and ad targeting, movie and product recommendations on Netflix and Amazon, Apple’s Siri digital assistant and IBM’s Watson question-answering system are all animated by artificial intelligence. Fed by vast amounts of digital data from sources like the web, sensors, smartphones and genomics, the software actually learns, in its way. The more raw data that is ingested, the smarter the artificial intelligence becomes.

The physical world where robots exist, meanwhile, is multidimensional and messy. So it is in the digital realm where the machine-learning algorithms of artificial intelligence have made their greatest strides — in tasks like facial recognition, language translation, prediction and decision making.

“A lot of the problems are easier in software, so a lot of the action is around intelligent software — softbots rather than robots,” said Oren Etzioni, a computer scientist and executive director of the Allen Institute for Artificial Intelligence in Seattle.

An interesting reality check took place in June, when two dozen teams of leading robotics engineers gathered in Pomona, Calif., for a competition sponsored by the Pentagon’s research agency. Their robots had to navigate mocked-up hazardous environments, like a contaminated nuclear plant, and do simple tasks — walk up steps, turn a valve, operate a power drill. The chores would take a human five minutes, or 10 at most. The winning robot took 45 minutes.

Most struggled badly, falling down steps and taking long pauses to figure things out, even with remote control assistance. Turning a knob to open a door proved daunting for many. One young man in the audience observed, “If you’re worried about the Terminator, just keep your door closed.”

Robots will surely get better. Google’s self-driving cars, for example, are impressive, but in heavy rain or snow, a human had better take the wheel. And robots with more limited ambitions are already taking over backbreaking work on factory floors and assisting surgeons for greater precision and control in operating rooms.

But the greatest progress has been in software, which is rapidly moving into the mainstream of the economy. So far, the largest commercial use of learning software has been in marketing, where it improves the odds of making a sale — tailored marketing, targeted advertising and personalized product recommendations.

Big companies and start-ups are beginning to use learning software in higher-stakes decisions like medical diagnosis, crime prevention, hiring selections and loan approvals.

The idea is that an A.I. turbocharger can be applied to all kinds of decisions, making them smarter, fairer and less prone to human whim and bias. The goal could be saving money or

saving lives.

Still, even enthusiasts have qualms.

Take consumer lending, a market where several start-ups are using big data and algorithms to assess the credit risk of borrowers. It's a digital-age twist on the most basic tenet of banking: Know your customer. By harvesting data from many sources, including social network connections, even observing how an applicant fills out online forms, lenders say algorithms can more accurately predict whether a candidate will repay than by simply looking at a person's credit history.

The promise is more efficient loan underwriting and pricing, saving consumers billions of dollars. But the new A.I. lending essentially amounts to a digital black box that pores over mountains of data. "A decision is made about you, and you have no idea why it was done," said Rajeev Date, a former deputy director of the Consumer Financial Protection Bureau. "That is disquieting."

Dr. Herbert Chase, a professor at Columbia's College of Physicians and Surgeons, was asked as an unpaid researcher to try out IBM's Watson software when the company's scientists were adapting the technology for medicine. To try to stump Watson, Dr. Chase recalled a case decades earlier, when he made a correct diagnosis of adult rickets for a young woman, but only after extensive tests and months of being baffled. He fed Watson a few symptoms and the program, which ranks diagnoses by probability, swiftly replied and ranked adult rickets second.

Not perfect, but Dr. Chase was impressed that the technology was so close and so fast. He thinks Watson-like software, which can scan and mine many thousands of medical articles in a few seconds, will be part of the future of medicine, assisting doctors inundated with information and short of time. But his nagging concern is that over time "we come to trust the technology too much, that it becomes the medical equivalent of blindly following a GPS system on your car down a dead end," Dr. Chase said.

This, of course, is a central issue as these intelligent systems — whether software programs or robots — evolve. Will they be servants or masters?

"Even if these systems are consistently right, slavishly following their instructions robs us of our ability to decide on our own," said Kristian Hammond, an artificial intelligence expert at Northwestern University. "Then you actually get the world no one wants. The machine algorithm tells us what to do."

The antidote is what Mr. Hammond calls "transparency" and others call "storytelling" — an explanation of the data ingredients that go into an automated decision and how it is made. University scientists and corporate researchers are working on A.I.-monitoring technology that ranges from data audit trails to English-language accounts of an algorithm's chain of reasoning.

Meanwhile, new research initiatives have sprung up that seem to have in mind that distant date when robots might achieve independence from their human masters. At Stanford, a group of leading scientists has begun a 100-year study on artificial intelligence. Mr. Musk

has put millions of dollars into research grants sponsored by the [Future of Life Institute](#), which is focused on guiding A.I. down beneficial paths.

But for the near term, it's those unseen algorithms — far more than robots — that deserve a watchful human eye. The stakes, some experts say, extend far beyond mere technology. “We need to make sure that the data and algorithms are continuously reviewed and vetted by a broad class of people,” said Alex Pentland, a computational social scientist at the M.I.T. Media Lab. “Think of representative democracy, forging algorithms rather than laws.”