# The Future of Artificial Intelligence

How robots, artificial intelligence and machine learning will affect our lives, and the innovators behind the algorithms

MaRS Discovery District

In partnership with

UNIVERSITY OF TORONTO

Karen Mazurkewich (above left) and Kara Collins at MaRS

# More artful than artificial: AI is already solving some of our most complex challenges

By Karen Mazurkewich, Editor, and Kara Collins, Managing Editor

Since artificial-intelligence software is encoded with the same basic 1's and 0's that have formed the building blocks of computer programs for decades, readers might ask, what makes AI "special"? What makes artificial intelligence more "intelligent" than, say, an old X-Box, or a gas-station microwave oven? And why all the hype?

AI is often confused with data analytics, but it's far more complex. It's a set of technologies that mirrors human intelligence, and augments it, by filtering vast inputs and options at speeds far beyond the computing ability of a human brain.

Geoffrey Hinton, a world-leading AI expert and chief scientific advisor at Toronto's Vector Institute, says programmers essentially teach computers how to solve a problem the way the human brain would if it could perform all the required computations: "You know the basic way to solve the problem. The computer is just doing the actual work."

Since machine learning involves much more than traditional computer algorithms, it promises to solve some of the world's most complex problems – from identifying traces of diabetes in ophthalmological scans to turning speech into text.

It is also a field that attracts a wide-ranging professional population. For example, developing machine-learning technology to identify early symptoms of cognitive impairment may require the combined expertise of computer programmers, psychiatrists, geriatricians and speech-language pathologists. And because the needed data sets may exist only in the hands of governments, universities or large corporations, these projects typically require collaboration among them all.

In 2000, long before the machine-learning technology revolution came of age, MaRS Discovery District was conceived as a hub where such forms of scientific and institutional cross-pollination would happen every day, across many different fields. With 1.5 million square feet of lab and office space, and a curated mix of diverse tenants, MaRS is the ideal platform to support multi-disciplinary discoveries. And its location in the heart of Canada's biggest city, next to where University of Toronto researchers have created many breakthroughs in AI, has already proven critical to forming one of the world's most productive and renowned centres for machine-learning incubation.

The pace of progress in the field of machine learning has been stunning. Many Canadians may be surprised to learn how many facets of their lives have already been affected by AI. In the articles that follow, MaRS and the University of Toronto have teamed up with Toronto's top AI institutes to give you an overview of these exciting developments, as well as an introduction to a few of the leading researchers and entrepreneurs whose work we support. ■

# Contents





**The Future of Artificial Intelligence**
**MaRS Discovery District**

This magazine was generously supported by members of the MaRS ecosystem that are making Toronto a global hub for artificial intelligence.

Our title sponsor, the University of Toronto, is a founding partner of MaRS and a global pioneer in machine learning.

We would also like to give special thanks to the Government of Ontario—a long-time supporter of MaRS—which also sponsored this magazine.

**EDITOR**
Karen Mazurkewich

**MANAGING EDITOR**
Kara Collins

**CONTRIBUTING EDITOR**
Jerry Johnson

**ART DIRECTOR**
Erik Mohr

**GRAPHIC DESIGNER**
Ian Sullivan Cant

**CONTRIBUTORS**
Gerry Flahive, Mary Gooderham, Tyler Hamilton, Jonathan Kay, John Lorinc, Glenn Lowson, David Paterson, Michael Rafelson, Nora Underwood

**GLOBE EDGE CONTENT STUDIO**

Jon Banack
Managing Director

Liz Massicotte
Program Manager

Isabelle Cabral
Production Co-ordinator, The Globe and Mail

Printing and pre-press by dm digital+1

**ALL CORRESPONDENCE:**
MaRS Discovery District
101 College St., Suite 401, Toronto, ON M5G 1L7
www.marsdd.com




# The intelligent choice

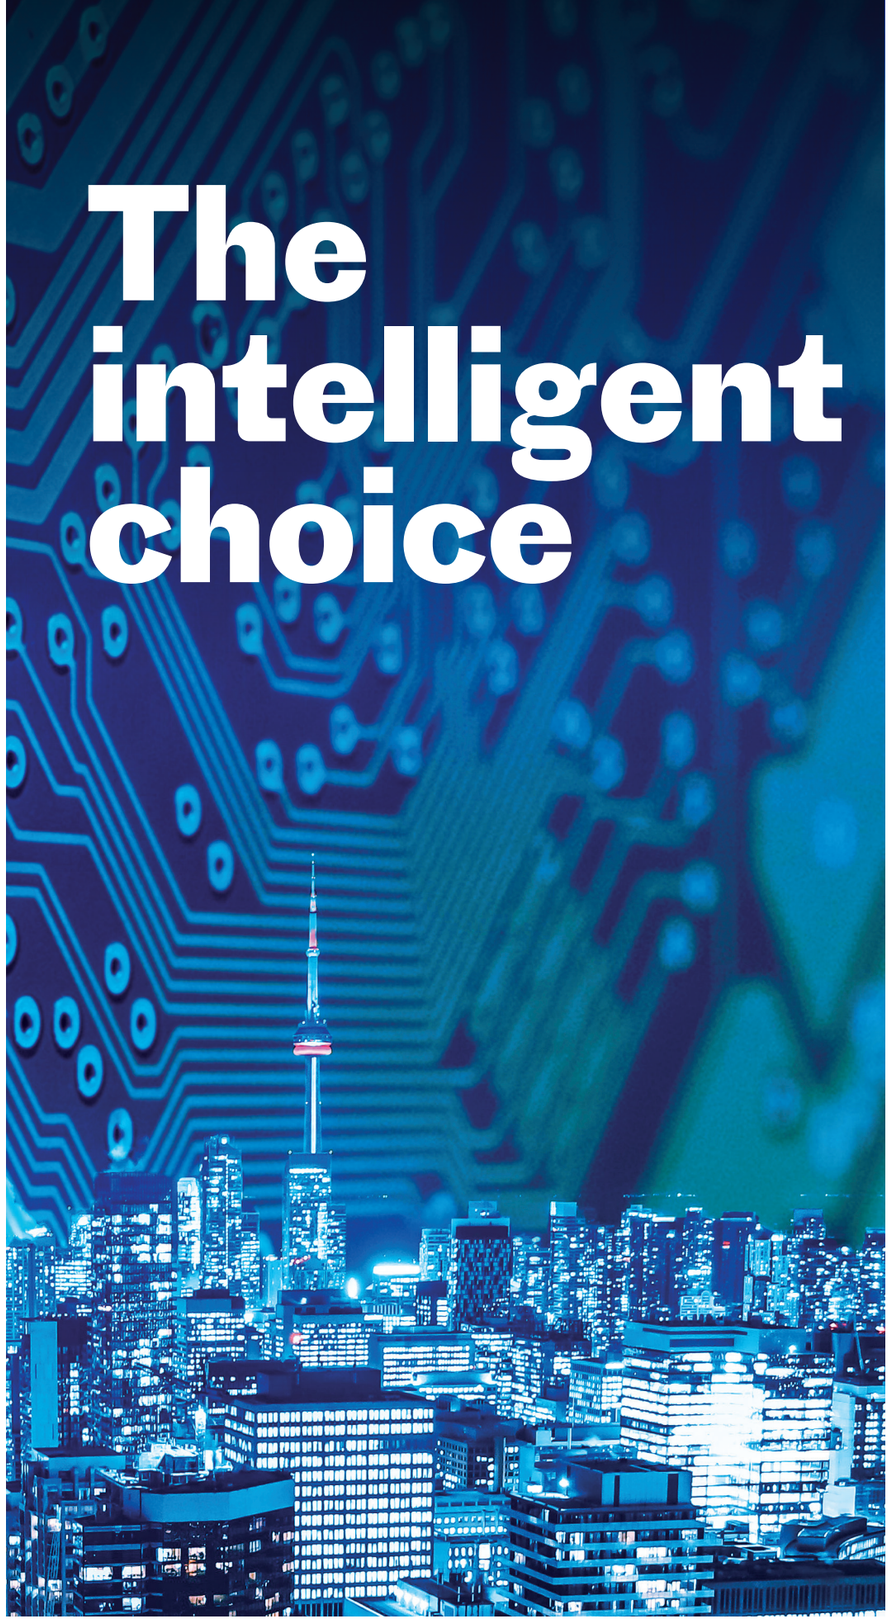

Whether it's labelled artificial intelligence, machine learning or big data, teaching computers to think now dominates the digital cutting edge, and Toronto is rapidly gaining a reputation as a major player. There are some good reasons for this, Jonathan Kay reports

When we talk about the benefit of diversity to our society, we often speak in general terms. But spend time with the entrepreneurs fueling the remarkable surge in artificial-intelligence (AI) research and development now taking place in Toronto, and the benefits of diversity become apparent in concrete form.

At a recent University of Toronto Rotman School of Management event showcasing business opportunities associated with the latest advances in AI, I saw venture demos from dozens of promising startups, all amid a crowd that was strikingly diverse even by Toronto standards. These things are awkward to quantify. But here's a random sequence of attendee surnames based on a scan of ID cards at the reception desk: Adejuwon, Ehrsam, Conde, Pal, Lepshokova, Dhamani, Kurian, Ing, MacGregor.

Of course, the Canadian technology sector has been benefiting from diversity for

generations. When I was doing postgraduate work at McGill's engineering department 25 years ago, my closest collaborators were recent arrivals from Iran, Turkey, India and China. In some cases, it was pointed out at grad seminars (in good humour) that I was the only Canadian-born engineer in the room.

But in too many cases, these colleagues would return to Asia or the Middle East to pursue careers as professors. Or, if they stayed in this country, they sometimes would have difficulty moving beyond narrow roles as research specialists. At the time, career development within science, engineering and information technology was still dominated by legacy tech companies struggling to get past old-fashioned attitudes.

A generation later, everything is different. Modern employers increasingly recognize the value of a diverse workforce. The sons and daughters of the immigrant scholars I collaborated with in the 1990s typically have broader, deeper networks (in part because of today's more collaborative approach to teaching science, technology, engineering and math), making it easier to bring their science into the market.

**$300K**

Salaries commanded by highly sought-after AI experts in today's competitive market

## Toronto's secret sauce

This pattern is playing out across Canada, with centres of AI excellence developing beyond the long shadow of Silicon Valley in Montreal, Ottawa and Edmonton, as well as in Toronto. However, Toronto's status as a larger and especially diverse city presents a special advantage that has led to a virtuous cycle: the city's brand helps to attract experts from other countries, who in turn help to lure in yet more talent.

"One of our advantages is that Toronto is simply a great place to live," says Richard Zemel, a computer science professor at the University of Toronto and the research director of the new Vector Institute for Artificial Intelligence. As *The New York Times* recently reported, experts in the AI field can command salaries upwards of $300,000 U.S. in a fierce global market. And Zemel often finds himself face to face with candidates who just as easily could decamp to Carnegie Mellon University in Pittsburgh, or the Berkeley Artificial Intelligence Research Lab in California. Sometimes, geography acts as a trump card: "When I'm recruiting new talent, one of my big selling points is that Toronto is simply a fun, lively place to work and live."

Toronto is still dwarfed by Silicon Valley when it comes to leading the digital revolution more broadly — especially when it comes to enormous, well-established brands such as Apple and Facebook. But the nature of AI technology allows relatively small companies to create highly profitable, technologically ambitious projects. Many of those companies are putting down roots in Toronto.

## Diversity is key, both in coding and research

And while abstract issues such as diversity and inclusion may seem removed from the nuts and bolts of AI coding, Zemel's own research suggests otherwise. While he spends much of his time working on core mathematical issues associated with AI, he also now is focusing on ensuring that new technologies exhibit what he, in one research paper, calls "algorithmic fairness" — so that machine-learned algorithms do not unconsciously embed racist or sexist suppositions into "the setting of insurance rates, the allocation of police, the targeting of advertising, the issuing of bank loans, the provision of healthcare [or] the admission of students."

Many of the brightest stars in AI are young, with careers that began during the current AI boom. (At the Rotman event, I spoke with DarwinAI founder Alexander Wong. When I mentioned that he looked quite young, his colleague Mohammad Javad Shafiee politely responded that Wong already has published more than 400 refereed journals and conference papers in computational imaging, AI and related fields. Moments like this happen all the time in this field.) But some veterans I spoke to stressed that today's success began with far-sighted policies implemented decades ago.

**When I'm recruiting new talent, one of my big selling points is that Toronto is simply a fun, lively place to work and live.** — Richard Zemel

According to Philippe Beaudoin, co-founder of Montreal-based Element AI, the Canadian Institute for Advanced Research is "one of the big reasons Canadian researchers developed expertise." A small organization that many Canadians may have never heard of, "CIFAR encourages networks of researchers on high-risk, high-reward ideas, even if the ideas have no immediate payoff," Beaudoin says. "In the case of [AI], some of these projects were cutting against the grain of the academy, where there was a lot of skepticism of this technology."

"It's also a great example of why this kind of diversity — diversity in research — is also important," he adds. "Government shouldn't just fund one set of projects. You have to invest in a lot of different fields, because you never know. And if Canada is going to keep its AI advantage, we're going to have to keep investing in fundamental research, and in our students, and in making sure our companies are early adopters of the latest technology. That doesn't mean just developing tech, but applying it, too."

To understand why the last five years have been a particularly exciting time in AI, and why Toronto has found itself at the centre of it, there's no better source than University of Toronto professor emeritus of computer science Geoffrey Hinton. He is a legend in the field, thanks to his pioneering work in our understanding of neural nets, which comprise the basis for modern forms of what now is known as "machine learning." Hinton has been working in the field since the 1970s. But it wasn't until the mid-2000s that computer hardware was sufficiently advanced to perform the trillions upon trillions of iterative computations required to implement his models.

In 2006, he and Ruslan Salakhutdinov, a University of Toronto colleague at Carnegie Mellon since 2009, led a scientific breakthrough that opened the floodgates to many of the machine-learning applications we now take for granted, such as accurate speech and image recognition. For non-specialists, it's difficult to understand the significance of this moment. Suffice it to say that their ground-breaking work, summarized in a *Science* article titled Reducing the Dimensionality of Data with Neural Networks, helped to solve one of the foundational problems associated with many neural networks.

"People were tremendously excited," Hinton recalls. "But there were no big, major, practical breakthroughs until it became implemented in speech-recognition. You may remember in 2012, when suddenly Android software could understand you much better. A lot of that technology pretty much came right out of the work my students did right here at the University of Toronto. And it showed that neural nets were the way to go for speech recognition — and a lot of other things, too."

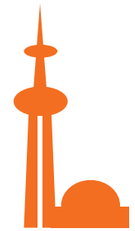

"The AI technologies that Hinton and his colleagues created and developed ... may do more for the city's growth, prosperity and reputation than Drake, Auston Matthews and Joe Carter put together."

It was an amazing moment in Canadian science, even if it was not widely lauded as such at the time. Hinton, who had been proselytizing neural nets for three lonely decades, all at once saw his brainchild being put to use in technology contained in gadgets that billions of people around the world carry in their pockets.

Now 69, Hinton is still at work, developing new neural-network technologies aimed at further streamlining and improving machine-learning processes. Among his colleagues at the University of Toronto and the offices of Google (he splits his time between the two), he is an intellectual celebrity. Yet, when he steps outside for lunch, few Torontonians recognize him.

Which is ironic, given Toronto's historic obsession with its place in the pantheon of "world-class" cities. For, in the long run, the AI technologies that Hinton and his colleagues created and developed at the University of Toronto may do more for the city's growth, prosperity and reputation than Drake, Auston Matthews and Joe Carter put together. ■

Naheed Kurji, president and CEO of Cyclica Inc., with his company's bioinformatics platform Ligand Express, which leverages AI to assess the safety and efficacy of drugs

# Faster. Cheaper. Better.

## When it comes to discovering life-saving drugs, AI kicks the "fast track" into overdrive

By Mary Gooderham

They say that speed kills, but in medical science it can also save lives.

For example, developing new drugs is lengthy, complicated and expensive. It requires that scientists identify proteins that cause disease, evaluate molecules that may target those proteins, and then go through rounds of clinical trials before receiving the regulatory approval needed to treat the sick and needy. The whole endeavour can take 15 years and cost billions.

Artificial intelligence, or machine learning, promises to accelerate this process—and enhance it. By finding the genetic mutations at the root cause of a disease, and rapidly predicting the effectiveness of treatments, AI can not only deliver drugs more quickly and more cheaply, it can also customize those drugs. Researchers and entrepreneurs in the field say that AI, particularly in combination with the latest research approaches, has the potential to guide the development of precision drugs and more effective therapies for a wide range of conditions.

"AI will be the enabling technology for medicine in the future," says genome biologist Brendan Frey, co-founder and chief executive officer of Deep Genomics, a Toronto startup that specializes in AI-driven genetic medicine. There is now what he calls a "huge wealth" of information on cell biology as well as on patients themselves – from fully sequenced genomes to the data gathered by fitness trackers and cell-phone apps – but it is inaccessible to family doctors.

"Currently, the system is broken," says Frey, who is also a co-founder of the Vector Institute, the new centre for AI research based at MaRS in Toronto, and a professor of engineering and medicine at the University of Toronto.

### AI will make drug discovery more intentional, less about luck

Human biology is "too huge and complex" for the mind to understand fully, he says. In the past, hypothesis-driven pharmaceutical development, using a "look-and-see approach," led to major breakthroughs, such as the discovery of insulin by Canadians Frederick Banting and Charles Best. "But what's crucial there is a huge amount of luck," Frey notes. "How many diseases and effective therapies are out there that we haven't come across by accident?"

The key, he adds, is to understand the biology of a disease and then engineer a therapy for it. This changes drug discovery from a "surprising type of endeavour" to something "very much data-driven and intentional."

Deep Genomics is applying this "intentional" approach to rare Mendelian disorders, a class of inherited diseases that result from a single genetic mutation, so they are considered relatively simple for scientists to target. Although rare, there are many such disorders; indeed, they are estimated to affect 350 million people around the world. One example, spinal muscular atrophy, is a leading cause of infant mortality, for which the U.S. Food and Drug Administration approved a new drug called SPINRAZA in December 2016, despite the fact it was still in the final phase of clinical trials.

Even with the FDA fast-track, the drug took 13 years to develop, and it costs patients $750,000 U.S. in the first year, then $375,000 annually for life. Frey suggests the use of modern technology, including AI, can reduce both the time frame and cost. "At Deep Genomics, the hope is to make better medicines available to more people, more quickly."

As well as rare diseases, AI is on the trail of mass murderers: the Ontario Institute for Cancer Research is using machine learning both to develop genomic tools for early cancer detection and to categorize tumours, so clinicians can devise more effective treatments, says Philip Awadalla, director of computational biology at OICR.

"There's a lot of data to be mined that have been under-exploited or under-utilized," he explains, and OICR has access to mountains of it collected both by the Ontario Health Study, its landmark attempt to help researchers better understand the environmental, lifestyle and genetic factors that cause cancer and other chronic diseases, and the International Cancer Genome Consortium, a global agency based at the institute. Analyzing all this information could bring to light, for instance, compounds that are used for one condition that could be applied to another.

As for drug discovery, says Awadalla, who also teaches population and medical genomics at the University of Toronto, AI can speed up the assessment of new compounds for their effectiveness, toxicity and ability to be metabolized. Researchers use "pre-existing data to predict potentially new interactions," he says, which screens out molecules with undesired effects and prioritizes more effective ones. This process takes seconds or minutes using high-speed computing, compared with months and years of experimentation.

But the process doesn't stop there, he says. It is still necessary "to take something into the lab, and you still need to go through the various steps of clinical trials. Sometimes there is something to be said for approaches that may be more traditional."

### AI is only one part of the drug-discovery model

What researchers need is a hybrid, according to Naheed Kurji, president and CEO of Cyclica Inc., a Toronto life-sciences technology startup.

"AI will play an important role," but is "only one piece of the puzzle," he says. "The nexus between knowledge-based approaches, like AI, and structure-based approaches ... is critical in ensuring that the best medicines are brought to the market faster and at a lower cost."

Brendan Frey, founder and CEO of Deep Genomics, which uses AI to build life-saving genetic medicines

AI is augmenting traditional methods of drug discovery, Kurji says, noting that previous approaches could not always take into account all possible side-effects. However, "the limitation of AI is that its predictive strength is based on the availability of high-quality data," which can be difficult to come by. Without enough basic information, "the predictive power of AI falls off a cliff" and, in the rush to harness AI, it is easy to forget that, while "blockbuster" diseases such as cancer, Alzheimer's and HIV have much available data, many that are highly complex or less prevalent do not.

Kurji says that coupling knowledge-based and structure-based research makes a "demonstrable difference." His company focuses on "off-target" interactions of new medicines, which can lead to side-effects. He's had some "powerful results" while using his platform to help researchers better understand the impact of whatever they're working on—be it drugs and nutritional supplements or even cosmetics and personal-care products—and try to avoid possible side-effects.

AI must become "a way of thinking" in drug discovery, says Frey, at Deep Genomics. He calls it "integral to the culture" at his company, whose AI platform "supports all the different stages of drug development and can lead to the ability to unlock new classes of medicines."

Beyond drug development, Frey adds, AI has the potential to reduce barriers all over the medical system. What's ultimately needed is a "direct pipeline" between scientists and patients, allowing them to share data more effectively and rapidly, in collaboration with medical practitioners and with proper safeguards in place.

"If you talk to experts," he adds, "you find that it's not possible to imagine a future of medicine without AI." ■

PHOTO: JEFFREY BEARDALL

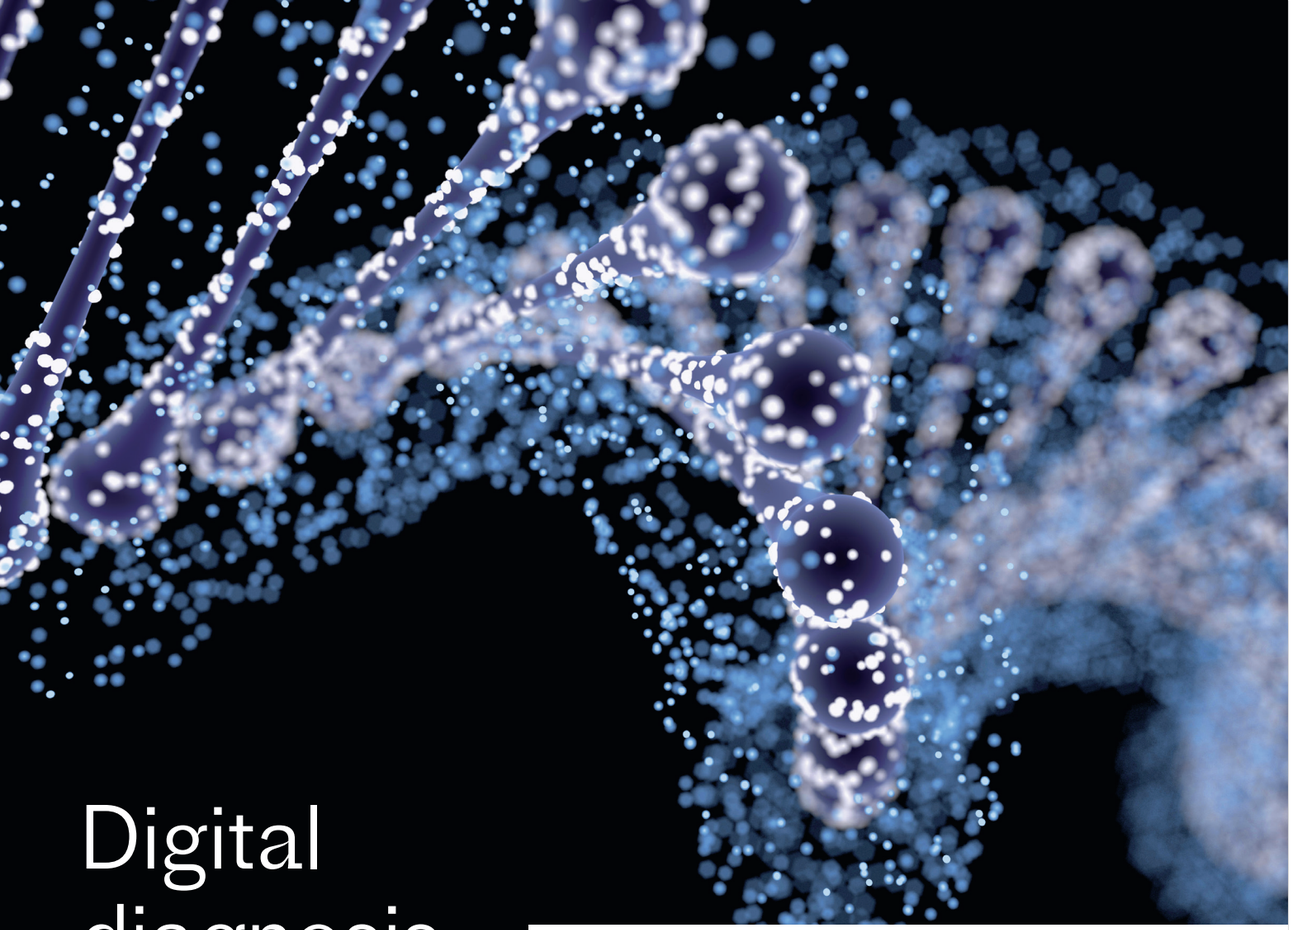

# Digital diagnosis

## Machine-assisted medicine is so good at figuring out what ails you that frail patients need not undergo arduous tests and Alzheimer's disease no longer comes as a surprise

By Jonathan Kay

To understand the revolutionary innovations in medical technology being unleashed by artificial intelligence, it's useful to start with the lives — and deaths — of two legendary actors: Paul Newman and Gene Wilder.

They both died, eight years apart, at the age of 83. But while Newman kept up an active public life until shortly before his death, Wilder spent his final years in seclusion. Only later did his family reveal that he had suffered from Alzheimer's disease.

Wilder was able to keep that fact private, but AI-powered technology developed in Canada now makes it possible to analyze interviews he recorded, from the early 1970s onward, and chart linguistic symptoms associated with cognitive impairment. For instance, as the years went by, he tended to use shorter noun phrases and fewer clauses per sentence. He also swapped out nouns for pronouns with greater frequency.

By contrast, analysis of Newman's interviews over a similar period does not reveal such a pattern. He died in 2008, with no signs of cognitive impairment.

The comparison of the two screen stars was conducted using software created by WinterLight Labs Inc., a startup based out of Johnson & Johnson's JLABS incubator in Toronto that brings together experts in speech, dementia, neurology and computer science. Using machine learning — a subset of AI that allows computers to self-construct decision-making algorithms through the recursive analysis of banked data — WinterLight's software parses recordings for hundreds of characteristics, including the length of pauses, the types of verbs that are used, irregular frequency and loudness, changes in vowel acoustics, reduced syntactic complexity and instances of repetition.

### AI-powered systems could give advance warning of illnesses

Each variable, taken in isolation, may have little or nothing to say about a patient's cognitive state. But when all are processed through a matrix that incorporates a mathematical model of their inter-relationship, the system can alert caregivers to markers for Alzheimer's or Parkinson's disease, as well as depression, multiple sclerosis and schizophrenia — sometimes years before a patient exhibits overt symptoms.

In one 2016 study, co-authored by WinterLight co-founder Frank Rudzicz, the approach achieved more than 80-per-cent accuracy in distinguishing individuals with Alzheimer's. Computer scientist Kathleen Fraser, also a WinterLight founder, has applied similar tools to achieve almost perfect accuracy in detecting patients with primary progressive aphasia, another degenerative neurological disease.

WinterLight has been enlisted in pilot projects at senior-care homes operated by Revera, VHA Home HealthCare and Shannex, as well as in clinical trials conducted by two large pharmaceutical companies. In so doing, its scientists have come to understand that accurate results aren't enough: The science has to be presented in a way that patients and caregivers find useful.

"People want to know *why* our software produced a certain result," says chief executive officer Liam Kaufman. "You have to explain what behaviours were measured. It isn't enough just to provide the numerical output."

WinterLight has yet to receive approval from regulators to market its software as a diagnostic medical device — a process that can move much more slowly than the creation of new software technology. But Kaufman and Rudzicz are optimistic.

**Like all modern workers, health providers and researchers are grappling with how much of their professional role in our society will migrate from human agency to computer algorithm.**

### Computers already outperform medical professionals in some kinds of diagnostic tests

The use of machine-learning technology to assist in medical investigations isn't new. Computers now outperform dermatologists in scanning lesions for skin cancer. And a Stanford-led group has created an algorithm that trumps cardiologists in detecting heart arrhythmias on the basis of electrocardiograms. But these applications represent an extrapolation of existing consumer-oriented technologies — such as facial recognition and photo classification — that focus on a single diagnostic artifact. WinterLight exemplifies the ongoing expansion of this machine-learning approach to broader and more complex types of inputs.

**14 Trillion**

Number of data points used (in big-data analytics) to assess a single tissue sample.

MarketsandMarkets, 2017

A product developed by Analytics 4 Life Inc., another startup based out of the JLABS incubator in Toronto, illustrates the same ambitious approach. The company's CorVista device and software package applies machine-learning techniques and three-dimensional imaging to detect coronary artery disease (CAD) on the basis of skin-surface electrode measurements and other physiological data. The technology is still being tested in clinical studies. But the ultimate goal, says Shyam Ramchandani, vice-president of clinical affairs, is to allow doctors to investigate the presence of CAD without subjecting patients to radiation, injections or exhausting clinical therapies.

"Existing diagnostic techniques in this area typically require the

injection of radioactive dyes and other contrast elements into a patient's bloodstream," Ramchandani explains. The whole process is uncomfortable, and can take many hours. And only a small percentage of these people even need treatment. We're creating a better alternative."

### Machine-learning diagnostics could sometimes be the kinder, gentler option

In pediatric medicine, in particular, the less invasive and less arduous diagnostic strategies facilitated by machine learning will offer significant improvements in the way patients receive care. During a presentation this past summer, Anna Goldenberg, an assistant professor of computer science in the University of Toronto's computational biology group, described the rigorous oncological monitoring regime required by sufferers of Li-Fraumeni syndrome, a hereditary condition that compromises the body's ability to suppress tumours. This regime includes regular full-body MRI scans, which are stressful for adults, and sometimes almost impossible for young children, who cannot lie still for the duration of the test.

Responding to this challenge, Goldenberg and her colleagues used machine-learning software to identify those Li-Fraumeni patients most likely to be diagnosed with cancer before age 6. Already, her non-invasive AI-powered surveillance model is approaching the accuracy level of traditional diagnostic methods, all without the associated cost and trauma.

Shyam Ramchandani, vice-president of clinical affairs for Analytics 4 Life, holds the company's CorVista diagnostic device. CorVista is limited by federal law to investigational use; it is not available for commercial distribution.

Goldenberg's presentation was titled *Will Dr. Robot Ever See You?* It's an apt question: Like all modern workers, health providers and researchers are grappling with how much of their professional role in our society will migrate from human agency to computer algorithm. While doctors once were assumed to be largely protected from the trend toward automation, that is changing: As machine-learning technology is used to automate the search for symptoms of diabetic retinopathy in eye scans, for instance, we may one day need fewer ophthalmologists.

Some of the new technologies also may change the relationship between patients and caregivers — sometimes in unsettling ways. For example, WinterLight's software may lead some patients to fear that every word they utter will be scrutinized for evidence of mental deterioration.

"I wonder about the dynamic between the nurse and the patient," says Samir Sinha, director of geriatrics at Mount Sinai Hospital in Toronto. "When people are losing their memory, they get paranoid, they get anxious. They get upset and depressed — because they know that this can be used as evidence to take away their freedoms. And if you start co-opting human communication, a fundamental way that people get pleasure and companionship, they might just keep their mouths shut."

Liam Kaufman, WinterLight's CEO, has thought about such issues. "White coat syndrome is a real problem, and not just with cognitive assessment," he acknowledges. "Just going to the hospital, it turns out, can make your blood rate go up."

But in his view, much of the stress arises from the irregular, high-stakes nature of medical visits. "Right now, someone may go to the hospital once per year — and they stress out and clam up. Our theory is that, by doing your assessment in the comfort of your own home, and by doing it frequently — every few months — it becomes a habit. And the stress actually goes down."

Because this technology is new, both patient-response theories are untested. But these issues will need to be explored, as AI-enabled tools take a more prominent place in our health systems. In geriatrics, as in most areas of medicine, the future likely will take the form of a partnership between the old and the new, humans and machines — with doctors informing their judgment on the basis of improved diagnostic analysis, without alienating patients in the process. ■

# Smart money

## Challenged by nimble startups, Canada's big banks are embracing AI, both to cut costs and compete. From chatbot service reps to predicting what will go viral online, the changes will be staggering.

By John Lorinc

The feeling of dread is nearly universal in our credit-addled society: Even as bills pile up and that retirement fund demands care and feeding, temptation rears its head — a pricey outfit, an impulse weekend in New York, that new iPhone. Some of us are highly disciplined when it comes to managing our finances, but others avoid the math and end up overdrawn, if not deeply in the red.

ATB Financial, an Alberta credit union, recently launched ATB Trackit Mobile, a smart-phone-based personal banking assistant whose artificial intelligence (AI) technology effectively allows customers to outsource the management of their cash flow to an app. Designed by Vancouver-based startup Finn.ai, the assistant automatically performs such tasks as creating budgets, tracking debt and sequencing payments by making predictions based on the customer's spending and saving habits so they avoid bounced cheques or overdrafts.

According to Jake Tyler, chief executive officer at Finn.ai, banks may benefit because their sustainably solvent customers will have more savings since they can better gauge how much consumer debt they can handle. Tyler believes they'll also gain a greater sense of brand loyalty, pay less in interest and avoid over-draft fees. "Every bank in Canada is looking at how to deploy part of what we're talking about," he says. "We're going to see the market move quite quickly."

Such partnerships illustrate how some fintech firms are increasingly inserting themselves into the space between financial institutions and their customers, many of whom now do much of their day-to-day banking on mobile phone apps with relatively limited functionality. Others are going a step further, grabbing ever-larger chunks of the consumer-lending business.

Foteini Agrafioti, (facing page) head of Borealis AI and chief science officer of RBC, in the Borealis AI office space

### Without AI, banks could soon see their own "Kodak" moment

Eyeing these rapid changes, the giants of the banking industry are looking to AI for help. They hope these intelligent computing systems will allow them to automate certain customer-service and wealth-management functions, reduce their geographical footprint and develop more precise techniques for detecting fraud and money-laundering.

AI systems will transform the way banks interact with their customers both "in big and small ways," says Ruby Walia, Head of Mobile and Online Banking for TD Bank.

Some observers predict that, for routine queries, online customer-service bots will replace personnel at branches and call centres within a decade. Others feel the shedding of jobs will happen even sooner. A survey of 600 international bankers by the consulting firm Accenture found that three-quarters believe bots will be ubiquitous within three years.

There are even those who argue that, if banking doesn't find religion with AI, it may face what former Barclay's chief executive Antony Jenkins recently described as its "Kodak moment" – turning the fallen photo giant's famous slogan into a synonym for death by disruption.

### Improving the customer experience

For example, Jordan Jacobs, co-founder and CEO of Toronto AI firm, Layer 6, says the lumbering sector has done much in recent years to meld pools of customer data gathered from savings accounts, mortgage and lending operations, and portfolio management. "Banks have spent a fortune building these data lakes," he says. "But they're incapable of doing anything predictive with it."

Layer 6 is marketing a set of algorithms that continuously scan customers' interactions with banks for patterns to help predict how they will react. As Jacobs explains, if a customer has just had a negative experience – such as a conflict with a call-centre service rep over a credit-card transaction – the system can serve up measures meant to improve that customer's next encounter, and thus seek to ensure the person doesn't switch banks.

Yet, despite Jacobs' skepticism, most Canadian banks are investing in AI capabilities, either directly or indirectly. This year, former TD Bank CEO Ed Clark was instrumental in setting up the Vector Institute, a partnership between government, institutions like the University of Toronto, and the private sector to advance AI research, and drive its application, adoption and commercialization. Located in the MaRS building, Vector's public support comes from the Ontario government

*(cont'd on page 15)*

# Banking on AI

**Why does Canada, a leader in AI research, fall behind when it comes to creating marketable products? Often, it's because our best and brightest head south. Banking giant RBC aims to change that, luring talent back with the promise of big datasets and greater research freedom.**

*BY DAVID PATERSON*

Borealis AI is billed as half theoretical academic research institute, half applied machine learning space. But it looks all startup. There are basketball hoops, exposed brick walls and a meeting room where the benches and table slide on tracks. Sparkling Italian sodas are the drink of choice among the staff, and they keep a broad selection on hand.

Aside from a familiar lion and globe logo on the phone screens, there's nothing to indicate that Borealis AI is actually an offshoot of one of Canada's oldest companies, banking giant RBC.

From their space on the third floor of MaRS, an innovation hub in downtown Toronto, Borealis AI's 35 staff are trying to apply the vast analytical abilities of machine learning to a host of problems. They've already taken AI developed to process machine language and converted it to recognize malicious sequences in streams of computer code, potentially offering a way to strengthen cybersecurity. And they are also working on reinforcement learning, a technique for teaching AI in the same way you might train a pet.

Among their moonshot ideas are using AI to better predict climate change and creating a system that could monitor global events and determine how they might affect stock markets or economies.

Founded a year ago, Borealis AI is RBC's response to a persistent question in Canada's artificial intelligence industry: Why does Canada, which is the birthplace of the latest advancements in machine learning, lead the world in academic research but lag countries like the U.S. and U.K. in turning those discoveries into marketable products?

Graham Taylor, an associate professor in machine learning at the University of Guelph, says that American investors were quicker to spot the potential of AI and began pouring money into the sector much earlier than in Canada.

Until recently, a brain drain of Canadian graduates headed south of the border sapped this country's ability to commercialize breakthroughs. According to Foteini Agrafioti, a brilliant researcher and former startup founder tapped by RBC to head Borealis AI, Canada was in danger of losing so much expertise and intellectual property, it would have been unable to catch up. Foreign companies would profit from commercializing research paid for by Canadian taxpayers – an unsustainable proposition.

"Research cannot be the destination," says Agrafioti. "Academic research is funded by the resources of this country, so it has to give back and generate wealth for the Canadian population."

In the past year, major investments from the public and private sectors, such as the Vector Institute and a potential AI innovation "supercluster" in Montreal, have helped reverse the flow of talent.

Borealis AI's lure for in-demand talent is the offer of a much broader degree of research freedom than they would get working for a product-focused tech firm, along with access to the computing resources and enormous data sets of Canada's biggest company. Staff pursue a mix of curiosity-driven fundamental research and applied machine learning projects. Engineers are encouraged to collaborate with universities and startups, and share their results with the scientific community.

As competition for talent intensifies, organizations will face growing pressure to allow their research-minded staff to rove between the theoretical and practical in their work. Taylor says that AI graduates want to work in places where they won't be chained to constant product development.

"One of the ways to keep people is [to create] these open source and open publication models, because that's the kind of environment they look for and want to work in," he says.

If these models succeed, they should smooth the path from discovery to product, creating and retaining intellectual property in Canadian hands, instead of letting it slip through our fingers.

That's something to raise a glass of Italian soda to. ■

*(cont'd from page 13)*

and the federal government's Pan-Canadian AI Strategy. It involves five of Canada's six large banks, plus several national insurers.

### Fraud detection, cybersecurity, the stock market — all impacted by AI

As well as backing Vector, the Royal Bank of Canada (RBC) has, like Google, set up its own university-style research lab with a mandate to publish papers and push boundaries, called Borealis AI. TD, says Walia, has formed a licensing partnership with Kasisto, a machine-learning spinoff of the famed Stanford Research Institute (SRI) in Silicon Valley. Based in New York, Kasisto specializes in customer-service messaging platforms, and will help TD create a mobile chat application.

Scotiabank, meanwhile, has paired with another Toronto AI startup, DeepLearni.ng, on a machine learning application designed to improve debt-collection systems. The algorithm, says Michael Zerbs, Scotiabank's Chief Technology Officer, searches for patterns that can help the bank contact account holders who have fallen behind in a way that prompts them to pay without getting angry.

As well, RBC has allied itself with high-profile AI researcher Matt Taylor from Washington State University, hiring him to lead their Edmonton lab, where he'll conduct research and development in a range of fields, from reinforcement learning to autonomous agents.

The bank is also eyeing the stock market. Foteini Agrafioti, head of Borealis AI and chief science officer of RBC, says her group is analyzing historical news articles to predict whether tweets or other digitally transmitted messages will go viral on social-media networks. The purpose is to provide real-time analysis of incidents that may affect share prices of companies in the RBC investment banking stable. "We're in a news-driven business," she explains.

Agrafioti's team of computer science and engineering PhDs is also conducting fundamental and applied academic research on various aspects of machine learning science, with an eye to developing applications in fraud detection and cybersecurity.

### Robo-advisors will automatically adjust your portfolio

Banks are also edging cautiously into using AI for wealth management through low-fee robo-advisors. So far, among bank-owned investment dealers, only BMO Nesbitt Burns has begun to offer such a service, known as SmartFolio. The field is dominated by independents such as Wealthsimple, a three-year-old Toronto "intelligent investment" startup backed by $100 million from Silicon Valley and Power Financial, a subsidiary of Quebec's Power Corp.

These services, geared at younger investors not wealthy enough to qualify for premium investment advice, use machine-learning algorithms to shape and adjust portfolios according to changing market conditions.

Jacobs of Layer 6 notes that, eventually, machine learning will affect everything from generic internal applications – such as hiring – to the huge volumes of work involved in investment banking. One example: algorithms capable of evaluating large numbers of commercial contracts as part of the due diligence process. Traditionally, Jacobs observes, such work is done by teams of lawyers and auditors, and can be extremely time-consuming, as well as costly.

Implicit in all this, of course, is the prospect of huge job cuts. Antony Jenkins has jumped to fintech since leaving Barclays, and says he anticipates a 50-per-cent drop in banking employment in coming years.

### Robots taking over? Not yet

Such predictions have drawn the attention of policy-makers at the highest levels, including Carolyn Wilkins, Senior Deputy Governor of the Bank of Canada. In a speech last April, she warned that AI "has raised the spectre of technological unemployment — the dystopian vision of an economy in which machines make many workers obsolete." She also cited studies suggesting job cuts nearly as high as those forecast by Jenkins. When accountants and investment bankers are included, Wilkins said, technology could eliminate up to 40 per cent of "tasks performed by humans."

**Some observers predict that, for routine queries, online customer-service bots will replace personnel at branches and call centres within a decade.**

**30%**

Estimated reduction in banking jobs due to automation over the next decade.

Frost & Sullivan, 2017

But from his perch, Scotiabank's Michael Zerbs says there is a substantial implementation gap between the entrepreneurial visions of startups and the institutional realities of large and complex organizations, where full deployment of advanced tech is "still many years away."

Nothing can happen, he adds, until banks figure out how to deliver reams of machine-readable data suitable for AI algorithms. And still more turns on the ability of these organizations to put the much-hyped technologies into operation.

And, as the ocean liners of banking slowly change course, consumers may find themselves increasingly reaching for AI-based fintech services as they seek to manage their finances on their own. ■

# Cleaner, greener... and smarter

## From heating homes and trading electricity to boosting wind power production and keeping pipelines safe, machine-learning algorithms and big data are giving the energy sector a serious boost

By Tyler Hamilton

When Ontario's energy ministry launched its "green bank" this past summer to help homeowners and small businesses become more energy efficient, one of its first actions was to give away and install 100,000 smart thermostats.

On the surface, it seemed like an expensive publicity stunt. But dig a little deeper and the campaign, with its $40-million price tag, is arguably one of the biggest efforts in Canada to give artificial intelligence (AI) a foothold in the home.

Chances are, the thermostat hanging on your wall right now isn't much of a thinker. It turns on and off when you tell it to, but doesn't really "know" anything about you. It doesn't sense and react to its surrounding environment.

Smart thermostats, on the other hand, are always learning. These clever devices, like models sold by Toronto-based startup ecobee or Google-owned Nest, never stop collecting data on local weather, and they're constantly taking note of when we come and go from our homes. They track the times we go to bed and wake up and, with the help of occupancy sensors, know which rooms we tend to use most and when we typically use them.

The more data that flows into these wall-mounted gadgets, the more accurate they get. Using machine-learning algorithms to tune into our changing routines and behaviours, they use the least amount of energy possible to keep us comfortable in our homes.

### Big changes coming to the electricity sector

It's just a small taste of how AI is starting to transform our relationship with energy as we make the transition to a low-carbon economy – from how and when we consume electricity and fuels in buildings, vehicles and industry; to how we produce, deliver, store and even trade it. The impact will be felt most in the electricity sector, which is expected to become more efficient, reliable, secure and safe as AI algorithms play more critical roles in an increasingly complex show.

"It may lead to a world where power generation, distribution and transmission operations are automatically optimized, where the grid is balanced independently of any human interventions, where trading and arbitrage decisions are made in nanoseconds at a scale that only machines could tackle, and where [customers] never have to worry about searching for a better supplier or changing the temperature manually," McKinsey, the international consulting firm, envisioned in a discussion paper last June.

It's unclear when the power grid will no longer need humans, but AI is already having a measurable impact. General Electric says it can boost energy production from wind farms by as much as 20 per cent using machine learning to anticipate and better respond to wind direction and speed, and to monitor wear and tear on parts, allowing for proactive maintenance. Google has slashed 15 per cent from its power bill by employing AI alongside a network of data-collecting sensors to make better use of fans, cooling systems and other equipment in its data centres.

### AI will turn your electric car into a power broker

Own an electric car? If companies like Microsoft and GE have their way, AI could one day be your own personal power broker, watching the electricity market and charging your vehicle only when the price is low or the wind is blowing; then selling the electricity in your car battery back to the grid when it can fetch a higher price (and you don't need it).

For car owners, it saves money – and sometimes makes it. For utilities trying to reduce fossil-fuel use, particularly during periods when electricity is in high demand, it's a way of harnessing an expanding network of vehicle batteries to create virtual power plants that automatically spring into action when needed.

Toronto-based Kelvin Thermal Energy is pursuing similar AI capabilities for big industrial customers that want to store cheap electricity as heat inside graphite blocks, and then extract that energy when it's most needed. "It's important that we have the ability to take, and stop taking, electricity when it's most advantageous for the customer, so we're leveraging all the great work that's going on in AI and predictive analytics to achieve that," says Stephen

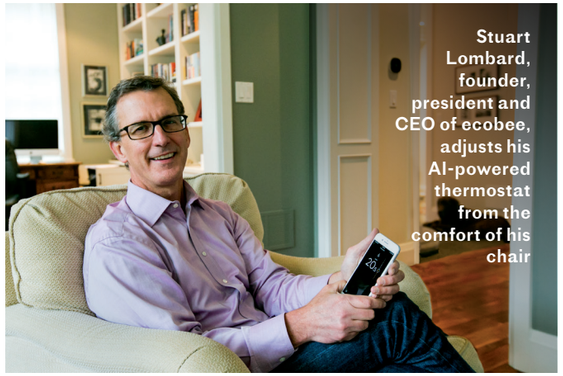

Stuart Lombard, founder, president and CEO of ecobee, adjusts his AI-powered thermostat from the comfort of his chair

Roloff, the company's Chief Strategy Officer.

Ecobee, for its part, is just getting started. Founder and CEO Stuart Lombard has big plans for the rich pool of data his company's smart thermostats collect. Within the last year, he says, ecobee has hired nearly a dozen employees with PhDs in machine learning to explore new applications: "The fight for talent is big, but we'll continue to grow that team over time."

One service ecobee could soon offer, likely in partnership with utilities, is remote home-energy audits. Using AI to analyze thermostat data, the company can look at patterns and deduce what a homeowner can do to lower energy use. "Our goal would be to get within 10 per cent of the accuracy of in-home audits," says Lombard.

Likewise, ecobee can help utilities tailor government incentives such as appliance rebates to the needs of specific customers, resulting in a more targeted and ultimately cost-effective use of scarce public dollars.

## Good AI requires good data

The risk, of course, is that we put too much trust in the decision-making capability of this increasingly intelligent gadgetry. An underlying machine-learning algorithm might be functionally sound, but its effectiveness depends on the quality of data being used to train it. That may not be a big problem for a smart thermostat that can be easily replaced if it doesn't strike the right comfort-savings balance. But if AI is relied on to juggle supply and demand on critical infrastructure like the power grid, bad data will lead to dropped balls, whether that means damage to expensive equipment or community-crippling blackouts.

"Artificial Intelligence techniques draw conclusions from large masses of data, which may or may not include garbage data," explains analyst Reinoud Kaasschieter of Capgemini, the global technology consultancy. "At a certain moment in time, it becomes impossible to determine on which data elements these predictions are based. In this way, artificial intelligence becomes black-box technology."

The more we come to depend on AI, the more important it will be to test the veracity of the data being used.

## Keeping pipelines flowing

Moving from electrons to molecules, this is something the petroleum industry will need to consider as it turns to AI for solving big problems. If you're the CEO of an oil company, for example, pipeline breaks are bad – even worse when they happen in the middle of populated or environmentally sensitive areas. It's why more providers of pipeline monitoring services are beginning to tap the power of machine learning to improve accuracy and reliability.

Calgary-based Ingu Solutions Inc. has developed a golf-ball-sized sensor that it inserts into an operating pipeline. As it's carried through the pipe, the sensor gathers a tremendous amount of data and can detect potential problems such as corrosion, cracks and the buildup of sediment.

"All of these anomalies or problems have unique data signatures. Every leak is different. Every case of sediment buildup is different. So what we need to do is sift through terabytes of data and scan for a signature that might represent a risk," explains Ingu co-founder and CEO John van Pol.

No human, he adds, can do the task efficiently or effectively, which is why the company is experimenting with AI algorithms – using platforms such as IBM Watson and Amazon Web Services – to identify and flag signatures in the noise of all that data.

However, like the average homeowner or power-grid operator, Ingu and its customers crave dependability. "What's extremely important is that we avoid false-positives because if we tell a customer there's a leak, they will dig up the pipe to do a manual inspection or repair. If they dig it up and nothing is wrong, they'll never use our technology again," says van Pol.

Thinking much like the Ontario government with its smart-thermostat giveaway, he adds: "AI is definitely something we need." ■

# Energetic upstarts

## Three cleantech startups that are cleaning up thanks to AI

By Tyler Hamilton

### General Fusion

Based in Burnaby, B.C., the company aims to build the world's first commercial nuclear fusion reactor, and is using AI to accelerate development plans.

A major challenge is understanding how to create a controlled fusion reaction with plasma fuel. That means running simulations of how plasma behaves under a seemingly infinite combination of operating conditions and variables. Here, AI algorithms are being used to speed up analysis of simulations and even reduce the number of simulations required.

Chief Technology Officer Michael Delage says it allows General Fusion to optimize its processes and reactor settings faster than ever. "These tools are just critical if we're going to take advantage of all the data we're collecting."

### energyX Solutions

With headquarters in Toronto, energyX Solutions has developed an online, fully automated service called MyEnergyXpert, which lets home and business owners assess how they use energy and learn about the actions they can take to lower energy bills.

The company says its digital scan, which takes only 15 minutes, is just as accurate as those done by engineering firms but at one-tenth the cost. On average, it identifies improvements, such as installing new energy-efficient windows or replacing old appliances, capable of saving 25 per cent of current energy consumption for customers.

Co-founder and Chief Technology Officer Alex Corneglio says the next step is to enhance the accuracy of the service. Using machine learning, he says, the digital tool will do more than recommend actions for customers – it will reliably make decisions for them.

### Thermo.AI

Machine-learning algorithms help Thermo.AI create the conditions in a power plant required for complete combustion of such fuels as coal or natural gas. Based in New York, the company employs special sensors to analyze moisture content in fuel, atmospheric pressure, interaction with air and other factors, making adjustments as necessary to assure the perfect burn.

AI helps customers get more energy out of their fuel, which not only reduces carbon emissions but extends the life and lowers maintenance costs of equipment by reducing wear and tear. Company co-founder Carolina Chaves Gonalez says the best way to make a huge impact is to make our current energy infrastructure more efficient.

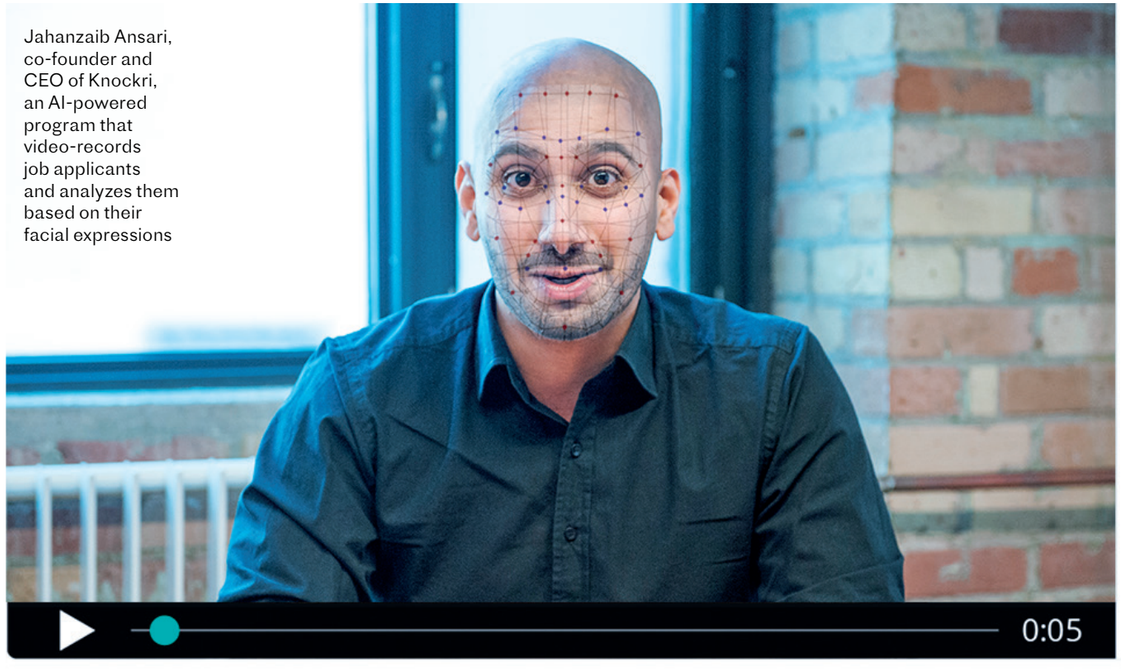


Jahanzaib Ansari, co-founder and CEO of Knockri, an AI-powered program that video-records job applicants and analyzes them based on their facial expressions

# R.I.P. résumés

## AI-assisted screening of job applications reduces sexism, racism and ageism in the workplace

By Nora Underwood

Four years ago, much to his mother's chagrin, Jahanzaib Ansari dropped out of university to pursue his entrepreneurial dreams.

His first effort, a bespoke tailoring business founded with two friends, was a success, but Ansari left. He had personal reasons but also wanted to try his hand at selling technology — in this case, a device that could create drinking water from moisture in the air. When that didn't pan out, he went to Amsterdam seeking to work with another tech startup.

Nothing materialized and, by November 2015, Ansari was back in Canada, broke and looking for a job. To his surprise, he did not get a single response to the many résumés he sent to prospective employers. "I've got a good soft-skills set, I've gone to school in Canada — I felt like I should have been getting a lot more job interviews," he recalls.

Then he came across studies showing that members of minorities who anglicized their names were more likely to get a second look. Jahanzaib became Jay, and "I got so many job interviews, I was astounded."

One of them paid off, but not for long. Dismayed at what he'd had to do just to get the job, Ansari decided to leave it and try to change the system. He and two partners founded Knockri, and joined the ranks of those now using cutting-edge technology to help employers find the best person for a job — without bias getting in the way.

One reason that racism, sexism and all the other "isms" are so pervasive in society is the fact that they aren't necessarily intentional: nearly everyone is guilty of unconscious bias. Even the most determinedly open-minded can't fight brain science, according to psychologist Timothy Wilson.

Wilson, who teaches at the University of Virginia, says that we are constantly exposed to millions of bits of information, but the brain can process only 40. To cope, "it creates shortcuts and uses past knowledge to make assumptions," he told *Fast Company* magazine in 2014.

In other words, our experiences and our culture can influence how we assess others, and we don't even realize it.

### Why your résumé isn't giving you a foot in the door

Artificial intelligence can help solve the problem by providing a bias-free screening tool. This saves employers valuable time and the need to sift through résumés, which are not always a reliable reflection of an applicant.

"All the data on that résumé doesn't actually predict success," explains Caitlin

**One reason that racism, sexism and all the other "isms" are so pervasive in society is the fact that they aren't necessarily intentional: nearly everyone is guilty of unconscious bias.**

MacGregor, chief executive officer of Plum.io, a Kitchener-based company with a different approach to applicant screening.

In fact, she says, résumés introduce a whole slew of biases — where the applicant went to school, past experience, name, gender. Only when "you don't use that as your short-listing factor" do you have a set of criteria "that matters."

Plum, like Knockri, works closely with clients to identify what they are really searching for in a job candidate. These qualities are programmed into the AI, and job seekers are filmed as they respond to questions that have been formulated with the help of psychologists to tease out whether they possess the desired attributes.

For each question, Knockri allows 10 seconds of prep time and 45 seconds to provide an answer that is analyzed with IBM AI, as well as the company's own proprietary technology, which examines muscle movements. Videotaping may reveal a person's gender and race, but Ansari explains that they, like sexuality, are not considered "desirable measures."

The motion technology, he adds, "just sees facial muscular contractions and analyzes tonality and style of response." It will register whether an applicant grimaces when talking about a previous boss, for example, and can detect undertones of aggression or enthusiasm, providing information that can be more meaningful than entries on a résumé.

### AI as job matchmaker

Obviously, all of these features benefit whoever is doing the hiring, but it could also benefit applicants, helping them avoid altogether a job for which they're simply a bad fit. There is evidence that as many as 70 per cent of workers may be in positions for which they are not well suited.

To Jamie Schneiderman that is a "huge problem" and it's one of the reasons why he launched Toronto-based Clearfit Analytics. "We set out with a mission to get people into the right jobs," he says, adding that the end result is "happier people and a more productive company."

Schneiderman knows from personal experience how important personal satisfaction can be. After earning a master's degree in business administration from Harvard, he worked with some large, well-known companies but "I found my career unbelievably frustrating."

Neither employers nor employees, he adds, want to make poor decisions, but avoiding them requires that they do things differently. Clearfit uses artificial intelligence to play matchmaker. Every applicant (the company says it has worked with 5,000 companies in the past 11 years and collected data on 1.5 million people) answers a series of questions. "Then our system can combine people together to build 'success profiles' on a role-company-industry basis," Schneiderman explains. "I can take a person who's applying to a company and automatically match them with the job for them."

### AI helps employers truly see their candidates

Also, like Knocki and Plum, Clearfit says its process increases diversity because applicants are shortlisted strictly on the basis of how they perform during the screening.

"Our system doesn't know how old you are, what sex you are, your ethnicity, your religion," Schneiderman says. "It can't see you. It can't judge you. It's solely looking at every single person in an equal, unbiased way."

Plum's MacGregor describes AI-assisted screening as "the top of the funnel" — at the other end, employers see only a list of applicants who meet, or are close to meeting, the criteria they have set out. In some cases, qualified candidates are right under their noses.

The psychometric assessments that are used can bring to the fore people already on staff who would otherwise escape notice. In such cases, companies avoid the time, energy and financial burden of employee turnover. "We have candidates who were previously overlooked surfacing," MacGregor says. "The short list of candidates ends up being far more diverse."

Of course, when humans take over, bias can enter the process, but at least the pool of available talent is better.

AI alone has not made these advancements in hiring possible, as MacGregor notes. Instead, to work properly, the technology has to be married to the lessons of industrial-organizational psychology — the study of how we behave in the workplace — or the machine will just replicate the errors of the past.

Rather than simply automating something that is faulty and outdated, she says, "our belief is you have to redesign the system to get a better result." Just as training can combat unconscious bias in the workplace, AI "provides us the opportunity" to overhaul the hiring process."

Schneiderman agrees and says that, because it can learn, AI can adapt by adjusting on the basis of the information it is given. Like a child, "it doesn't have inherent biases."

The potential catch, then, is that AI still depends on the humans who input the data. But the companies using it say they strive to ensure that their criteria remain as objective as possible. There is no point in promising to find the best person for the job if someone like Jahanzaib Ansari is disqualified simply because of his name. ■

# The very real ways that artificial intelligence is reshaping medicine

**A Q&A with Dr. Guna Rajagopal, Ph.D., Global Head of Computational Sciences, Discovery Sciences, at Janssen Research & Development**

Artificial intelligence presents boundless opportunities to transform human health. Already, it's aiding in the discovery of new drugs and helping to create better, faster medical diagnostics. It's not difficult to imagine a time in the near future when A.I. is able to reliably predict and intercept disease before symptoms ever arise.

However, with so much hype surrounding A.I., it's easy to forget that we've only scratched the surface in understanding what's possible with machine learning, and how it may be applied to healthcare, says Guna Rajagopal, Ph.D., Global Head of Computational Sciences, Discovery Sciences, at Janssen Research & Development.

Dr. Rajagopal leads a team of data scientists who are using high-performance computing to assist in the creation of new drugs and healthcare products. He's also currently working with Johnson & Johnson Innovation, JLABS to evaluate nominees to the Artificial Intelligence for Drug Discovery QuickFire Challenge, which will award up to $100,000 in grants and one year of JLABS residency to individuals or teams with the best ideas for using artificial intelligence to advance healthcare.

With the excitement of the QuickFire Challenge selections looming (expect an announcement in mid-December), we spoke with Dr. Rajagopal to learn more about how A.I. is changing the future of medicine.

**Q.** Artificial intelligence means different things to different people. How do you define A.I., especially as it relates to healthcare?

**A.** A.I. is the science of building and programming a machine that's able to imitate human cognition. The machine can learn from experience and generalize, which is where the intelligence part comes in. Most of us already interact with A.I. in our daily lives, whether it's Amazon giving us personalized suggestions of products we might like, to real-time alerts of potential credit card fraud. Regardless of the industry, the purpose of A.I. is to guide humans to make informed decisions based on enormous amounts of data. As it pertains to healthcare, this data may include an individual's genetic and environmental factors, activity trackers and bio-sensors, blood samples and electronic health records. The amount of health and genomic data that we're generating is growing exponentially every year. While our A.I. tools are improving by the day, we also have to recognize that big-data research is still very much a nascent field. We have a lot to learn.

**Q.** What are some of the biggest transformations that A.I. will bring to healthcare?

**A.** It's impossible to predict all the ways that A.I. will change health and medicine over time, but I believe it will be truly fundamental and sweeping. I'll defer to the famous quote from American futurist Roy Amara: "We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run." This is true of A.I., just as it has been for other transformational technologies. If you look at the impact of A.I. for healthcare, you have to consider not only how it relates to the development of new medicines and faster clinical decision-making, but also elements such as pharmacy and supply chain – how we optimize our resources to get the right products to patients most efficiently.

**Q.** 'Precision medicine' just may be the healthcare buzzword of the decade. How might A.I. impact the development of drugs that treat patients in a more precise way?

**A.** We've only begun to realize the benefits of precision medicine to treat disease, with notable early successes in cancer. The underlying idea is that we can use patients' personalized information, such as genetic or molecular profiles, to determine what treatment approaches will work best for them as individuals. As we move forward, our success in advancing precision medicine will depend not just on collecting and storing vast datasets on diverse patient populations, but also on our ability to develop sophisticated machine-learning algorithms that can mine this data to answer specific healthcare questions. We see A.I. as a tool to help us analyze these factors and bring clarity to patients earlier in the process.

**Q.** In the latest QuickFire Challenge for Johnson & Johnson Innovation, JLABS, you've put a call out for entrepreneurs and innovators that are using artificial intelligence to advance drug discovery and development in some way. What do you hope comes as a result of this type of competition?

**A.** Even though Johnson & Johnson is among the world's largest healthcare companies, we realize that we can't do it all on our own. Our goal is to create an ecosystem of innovation around A.I. in healthcare, partnering with the best and thinking holistically about how solutions can be applied to various aspects of drug development and patient care. If you think end-to-end about our healthcare system – from the lab where drugs are created, to the patient's bedside – the one thing that ties it all together is data. I'm optimistic about what A.I. can do to bring meaning to our growing pools of information, aiding in our ability to intercept, diagnose and treat disease.



# What is possible vs. what is right

## How far should we let AI go? There's a growing campaign to make sure society comes out a winner

By John Lorinc

The transformative power of artificial intelligence has come to preoccupy big business and government as well as academics. But as AI's potential sinks in, a growing number of policy experts — along with some leading figures in technology — are asking tough questions: Should these cutting-edge algorithms be regulated, taxed or even, in certain cases, blocked?

Consider what AI can do in the workplace. For example, managers realize that office politics, stress and other pressures take a toll on employees. They also know that standard-issue job-satisfaction surveys "don't provide a true gauge of what's going on" around the water cooler or in the staff lunchroom, says Jonathan Kreindler, Chief Executive Officer of Receptiviti.ai.

To tap into more candid expressions of employee sentiment, his company, a three-year-old Toronto-based startup, has created an AI algorithm grounded in the research of James Pennebaker, a University of Texas social psychologist who has found that the way employees communicate with each other can provide insight into their behaviour and state of mind. So, the Receptiviti algorithm scans internal messages for particular words and expressions that Pennebaker says indicate dissatisfaction.

Anyone who works for a large organization will know their emails

are not their own; nonetheless, Receptiviti's service raises questions about privacy, even though Kreindler is quick to point out its goal is simply to find out how employees really feel: "We're not reading emails to understand what people are talking about."

He also stresses that "there are no regulations preventing us from doing what we're doing." Still, employees may wonder what's really going on, and firms such as Wall Street's Goldman Sachs, he adds, have already been called out for scanning emails for certain keywords.

## Does computing need human oversight?

Even with the best intentions, a growing number of firms like Receptiviti find themselves thrust into the increasingly high-profile debate about the social implications of AI. Previous waves of digital development have triggered dialogues about everything from Internet regulation to online bullying and censorship, but AI seems to have raised the stakes. After all, it points to a future in which computing doesn't necessarily require human oversight.

Some critics, as well as AI experts themselves, have focused on the mid-term impact on employment.

**75%**

The number of workers globally who suspect their employers are capturing data about them without their knowledge.

Deloitte

An alarming loss of jobs has been predicted based on the implementation of AI-driven systems, including the replacement of everyone from call-centre staff to lawyers.

Personal privacy isn't the only source of friction, says Joe Greenwood, a specialist in data analysis at MaRS Discovery District in Toronto. Other sensitive areas range from the quality of the data that is used by AI to what happens as robotic devices replace human beings. In the latter case, even Microsoft founder Bill Gates has suggested special taxes and a slower pace of adoption while authorities ask themselves, as he said in one interview, "Okay, what about the communities where this has a particularly big impact?"

Consent is another issue, Greenwood says. AI systems rely on a variety of sources of consumer data, but have the consumers given their permission to having personal information combined with less sensitive statistics?

Another major concern is reliability. For example, algorithms are being developed to read and interpret legal contracts and make certain types of medical diagnoses (such as skin cancers). "Are we confident enough that [the algorithm] has done that well?" Greenwood asks. "Or are we taking it on blind faith?" As for human oversight, he points out that medical staff no longer double-check the readings of heart monitors; the same could well happen with diagnostic AI.

And what if the AI has bad data? What if a person unjustly runs afoul of a facial-recognition device or "predictive policing," which is designed to anticipate whether people will be victims of crime, or commit it?

Writing in TechCrunch last year, Kristian Hammond, a Northwestern University computer scientist, identified several types of algorithmic bias, some of it absorbed from the human actors who interact with AI. (One oft-cited example: Tay, an AI-driven Twitter account created by Microsoft, became so aggressively racist and misogynist when let loose online that it was shut down within 24 hours.)

Calls for regulatory oversight have begun. The high-profile tech figures expressing concerns about machine learning are led by Tesla's Elon Musk. Early this year, he spearheaded a high-level effort to develop guiding principles for intelligent automation through the Future of Life Institute, a Boston think-tank that sponsors research on promoting a safe future for humanity. Dozens of leading researchers and entrepreneurs — such as physicist Stephen Hawking, Skype co-founder Jaan Tallinn, and Montreal AI researcher Yoshua Bengio — put their names to the Asilomar Principles, a list of 23 tenets they consider a necessary ethical foundation for this technology (Asilomar is the California conference centre where the summit took place).

## How can we hold AI accountable?

As Greenwood puts it, the principles pose the question: Do AI-based systems require accountability mechanisms if something goes wrong with the way the technology functions? As well, AI Now, a New York University-based think-tank founded by data scientists Kate Crawford and Meredith Whittaker, holds annual conferences on a range of issues related to machine learning. This summer, the group teamed up with the American Civil Liberties Association to identify algorithmic bias in AI-affected fields as disparate as money lending and parole rulings.

North of the border, the Canadian Institute for Advanced Research (CIFAR), which funds the newly established Vector Institute through the federal $125-million Pan-Canadian AI Strategy, promotes AI and has the additional mandate to investigate its social, legal and ethical implications, says Brent Barron, CIFAR's director of public policy.

Barron says CIFAR is tracking the debates taking place within AI as it develops a policy framework for Canadian research and development, and will convene advisory panels to scope out what should be addressed. But he also says "it's early days" and, for now, many AI startups will have to play cat-and-mouse with the evolving regulatory environment in jurisdictions across the country.

Kathryn Hume, vice-president of products and strategy for Toronto-based startup Integrate.ai, says her company has done this already. Founded less than a year ago, Integrate is testing a machine-learning system that draws on disparate data pools maintained by large companies with extensive consumer bases. The algorithm is meant to come up with solutions that aim to optimize customer engagement.

**100 million**

Estimated number of new jobs needed in the US to replace those lost to AI in the next two decades.

Future of Life Institute, 2017

The underlying idea – how can we keep our customers happy by analyzing data patterns – seems straightforward enough. But Hume points out that, when the algorithm also draws on third-party sources of data, such as credit scores, to develop conclusions, it immediately confronts questions about privacy and consent. Technically, she says, the company will erect a "mathematical wall" between the analysis and the underlying data – an emerging concept known as differential privacy – because it is seeking any patterns that data contains, not anything associated with any particular individual.

Despite those assurances, Hume points out that, in such regions as the European Union, regulatory authorities have already begun to promulgate AI-related policies, such as right-to-know laws and policies mandating AI firms to be prepared to explain just how their algorithms generated particular outcomes.

At the moment, she says, the difficult questions about the implications of AI adoption outnumber the answers: "The legal community and the regulators have a lot of work to do to figure out what all this means." ■

# Guiding principles for AI

The Future of Life Institute says the Asilomar Principles, the following 23 measures to guide the advent of the intelligent machine, "offer amazing opportunities to help and empower people in the decades and centuries ahead." Adopted by delegates to the institute's annual gathering this year at the Asilomar conference centre in Monterey, Calif., they fall into three main categories:

### RESEARCH

1. The intelligence created should be beneficial.
2. Investments should be accompanied by funding to ensure AI is used well, including such thorny questions as how we can prevent systems from malfunctioning and increase prosperity while maintaining people's resources and purpose.
3. There should be a constructive exchange between AI researchers and policy-makers.
4. A culture of co-operation, trust and transparency should be fostered among AI developers.
5. Teams developing AI should strive to work together and not cut corners on safety.

### ETHICS AND VALUES

6. AI should be verifiably safe and secure, where feasible.
7. If a system causes harm, it should be possible to ascertain why.
8. Any involvement by an autonomous system in judicial decision-making should provide a satisfactory explanation auditable by a competent human authority.
9. Creators of advanced AI have a stake in its use and misuse, and are responsible for its outcomes.
10. AI's goals and behaviour should align with human values.
11. It should be compatible with ideals of human dignity, rights, freedoms and cultural diversity.
12. People should have the right to access, manage and control the data they generate, given AI's ability to analyze and utilize that data.
13. The application of AI to personal data must not unreasonably curtail personal liberty.
14. AI should benefit and empower as many people as possible.
15. Prosperity created by AI should benefit all of humanity.
16. Humans should choose how and whether to delegate decisions to AI systems, to accomplish human-chosen objectives.
17. Rather than subvert, AI should respect and assist the processes on which a healthy society depends.
18. An arms race in lethal autonomous weapons should be avoided.

### LONG TERM

19. We should avoid strong assumptions regarding upper limits on AI capabilities.
20. AI could represent a profound change in the history of life on Earth, and should be planned for and managed with commensurate care and resources.
21. Risks posed by AI, especially those that are catastrophic or existential, must be subject to planning and mitigation efforts commensurate with their expected impact.
22. AI designed to improve or reproduce itself rapidly must be subject to strict controls.
23. Superintelligence should be developed only in the service of widely shared ethical ideals, and for the benefit of all humanity.

Image created by Google's DeepDream, a program that feeds pictures through a neural network, and then uses AI to enhance and build on some of their features

# Wanted: Creative spark

Can a computer write a great novel, paint a masterpiece? Not yet, but algorithms are beginning to prove very helpful in humanity's quest for creativity. And it won't be long, some predict, before they're making that quest on their own.

By Gerry Flahive

You're designing an office environment for 250 software workers. Parameters? Many. Time and resources? Limited. Number of employees who would like to have their desk near a window? All of them.

And then your creative collaborator, having considered all the options, comes back to you with 10,000 possible floor plans.

That's a lot, unless the collaborator is a program, not a person.

In fact, this surfeit of choices was just what Autodesk, a global creator of 3D design, engineering and entertainment software, wanted from Project Discover, the artificial-intelligence program it developed to help design its new two-storey office space at MaRS, an innovation hub located in downtown Toronto.

Organizations have long relied on a boardroom filled with people powered by coffee and surrounded by flip charts to generate ideas. But old-fashioned brainstorming has fallen into some disrepute. The fit and feasibility of the ideas it produces often can't be assessed quickly or easily, given the ethos of "every idea is a good idea."

What's more, the brainwaves are limited to the ingenuity of who happens to be in the room. Azam Khan, director of complex systems research at Autodesk, initiated the project, and says "the team's intention in designing the system was for them to discover something that they *couldn't* do on their own."

As well as that, he adds, the system assesses the ideas it produces, narrows them down "to a manageable set of outcomes," and then "presents those to the human designers from The Living, an Autodesk studio that wrote the original code to help them understand the tradeoffs that you wouldn't normally experience unless you did 10,000 designs."



**Although AI can already produce poetry, pop music or movie scripts, "there's the sense that something is fundamentally missing. In a way, the algorithm doesn't really get it."** — Inmar Givoni

## AI can see things that humans can't

In the process, Project Discover, well, discovered, that a "generative" design system can see things we can't. And because the algorithm "uses concepts found in natural evolution," according to its creators, it also can ensure survival of the fittest, by "gradually promoting the best options" for serious consideration.

What Autodesk has done reveals the immense potential of using AI to turbo-charge the creative process: It can come up with a huge number of possible solutions, but then winnow them down to a practical number. By rapidly running through virtual prototypes of solutions, a learning machine becomes a kind of serendipity engine that powers creativity.

But does this serendipity engine still need a driver?

The creative approach adopted by Project Discover, says Khan, has sparked "ravenous interest" in all sorts of applications of the software the project is to produce. He sees it being used "to optimize, for example, factory layouts, electronics designs and even entire neighbourhoods."

But AI-driven creation is far from a turnkey operation. There are still some big philosophical hurdles ahead.

For one thing, what exactly is creativity?

## AI might be productive, but does it *get* creativity?

Machine-learning specialist Inmar Givoni calls it "an abstract concept that is hard to nail down and properly define. It's one of those things where we know it when we see it. Or more accurately, we *think* we know it when we see it."

Givoni, who works for Kindred, a Toronto company trying to "enable robots to understand and participate in our world," says that, although AI can already produce poetry, pop music or movie scripts, "there's the sense that something is fundamentally missing. In a way, the algorithm doesn't really get it.

"It doesn't understand what makes sense and what doesn't, and when is it interesting to *not* make sense."

Sanja Fidler at the University of Toronto says the advent of AI that "gets it" may just be a matter of time. The assistant professor of computer science and her colleagues have a program that generates pop music — they call it a "neural karaoke" — from photos, as well as another program focused on neuro-esthetics (it can calculate what's in fashion on the basis of thousands of clothing images from social media).

## The future of storytelling

Now she is trying to determine if such neural networks can go "beyond the data you give them." For example, "if you are asking AI to generate, say, 50 to 100 new stories based on data from thousands of books, what it creates is going to be somewhat biased because it's taking stuff from that existing pool of information." Because there is no AI that is "embodied" — able to "just go


View of Autodesk's AI-designed office space in Toronto

around and build its own stories" — all that technology can do "is read about other people's experiences, and kind of blend them together."

But that won't last forever, Fidler says. Today, the story — perhaps the most fundamental "unit" of human creativity – is too complex to be reduced to narrative patterns or data sets that a computer can use to generate satisfying results. But in the future, "there are going to be embodied AI agents that build their own world of experiences," she predicts. And when that happens — when machines can explore the world and gather the stuff of stories — we may see work that isn't generated by the human imagination.

Meanwhile, as scientists struggle to endow a computer with that ethereal quality, the imagination, AI continues to expand its assistive reach into practical realms, from refining Netflix recommendations to detecting email spam and transcribing interviews. Highly useful, but not creative.

But back at Autodesk, moviemaking veteran Hilmar Koch looks ahead and envisions a creative middle ground: AI that can do more than assist human storytellers; it can empower them.

Koch recently joined the company, after a career as a visual-effects pioneer on such films as *Avatar, Star Wars: The Force Awakens* and *Jurassic World*, in a new position: director of research and development for the future of storytelling.

He, like Fidler, sees technology's limitations. "Are computers becoming creative? That is questionable," he says. "They don't know about the human condition. They don't know what it's like to be jobless, or a parent, or to swim in the ocean." As a result, "I do not look at [AI] at all as here to replace people's jobs."

## AI stretches the realm of possible for human storytellers

What it can do, however, "is open up possibilities, as the canvas gets stretched out beyond the edges of what we know right now."

So, the team Koch leads at Autodesk, called Project Narrativa, is exploring what lies ahead for the story in any art form. And with the emergence of AI that can generate endless choices, "the role of the creative, of the artist, has to change," he says. "You are going to have to guide the process. You need to prune this decision tree at a very rapid pace, or it could very much overwhelm you. How are we going to behave in the presence of too many possibilities?"

In other words, 10,000 machine-generated screenplays aren't going to advance the art of cinema, but Koch can foresee using AI to develop what he calls "the story-information model – everything that you need to know about your story."

An example: Television and cinema are now home to hugely complex story worlds – think of *The Walking Dead* and its multi-platform spinoffs, or the many layers of *Star Wars* productions, now so dense with plots, characters and "rules" that Lucasfilm has to employ archivists just to keep track. One of them is called the Keeper of the Holocron, a database that now houses more than 80,000 entries.

Such ambitious properties, with so many humans involved in their creation, could benefit from the all-seeing technology that is to come, Koch says. AI could be a valued partner "as we search along a guided path through otherwise overwhelming data sets, and carve the story out of that big mass of marble, the potential story that's already in there." Before computers can do anything creative with all this information, he adds, "we need to shape the data so that it is better for humans to deal with."

He says the magic-button allure of AI doesn't cloud his faith in, and passion for, the human creative exchange. "What makes me tick is that ... creative spark, the handing over of an idea, one person talking to another." From that spark, "our ideas multiply and build on one another, and they become that much stronger."

As long as the creative force is liberated, Koch says, "I do not care what technology is running behind the curtain." ■

PHOTO: BEN RAHN/A-FRAME INC.

UNIVERSITY OF TORONTO

# WHERE IDEAS COME TO LIFE.

As Canada's most innovative university, the University of Toronto is the best place for entrepreneurially minded students and researchers. The university hosts nine incubators across its three campuses, and recently celebrated the opening of ONRamp, a new community space where entrepreneurs can collaborate.

In the past five years, U of T has fostered the development of more than 150 companies, making it Canada's #1 engine for research-based startups. With the country's deepest pool of research expertise, U of T offers multiple paths to personal and entrepreneurial success – whatever your idea might be.

**To learn more, visit UofT.me/entrepreneurs**

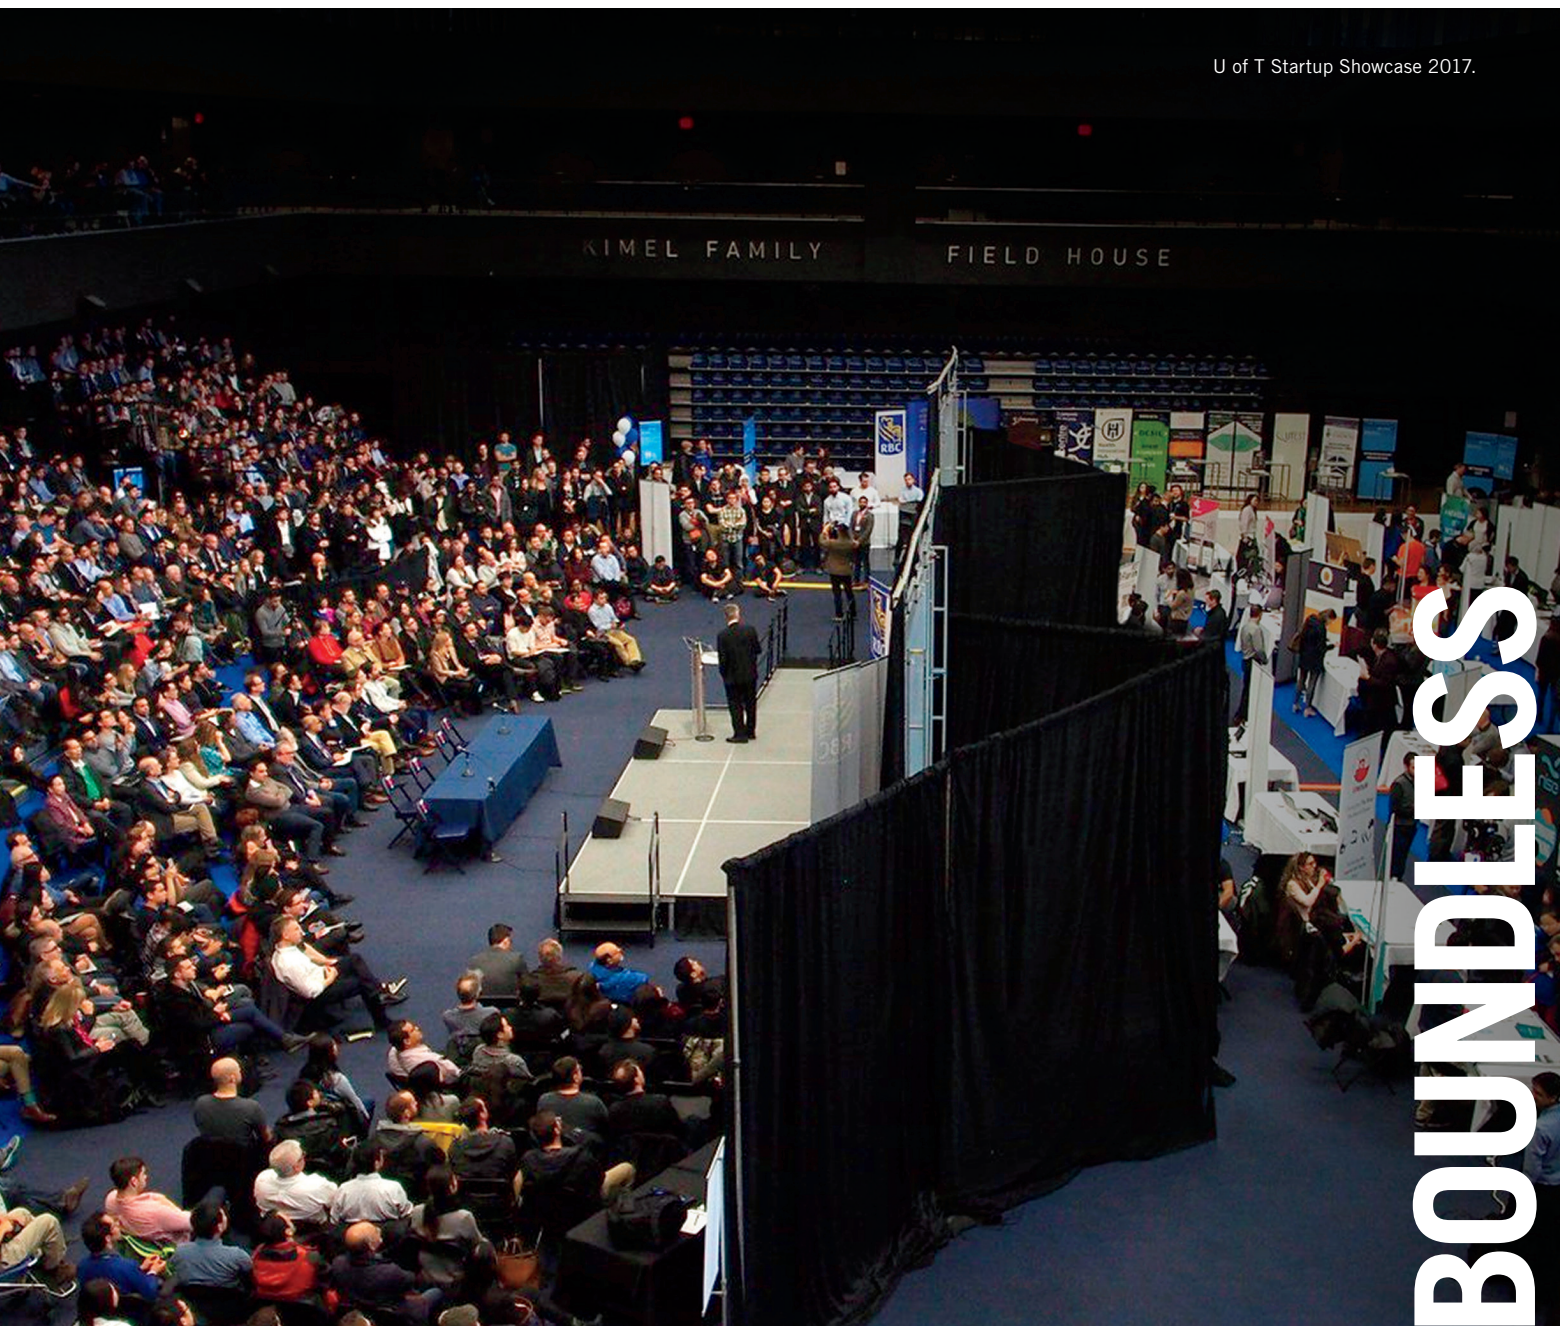


U of T Startup Showcase 2017.