#1 TOP AI BLOG
TOP 50 BLOGS

aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:

Dr. Ajay Agrawal

Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:

Dr. Russell Greiner

Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:

Dr. Foteini Agrafioti

Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:

Yoshua Bengio

MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

aitrends

The Business and Technology of Enterprise AI

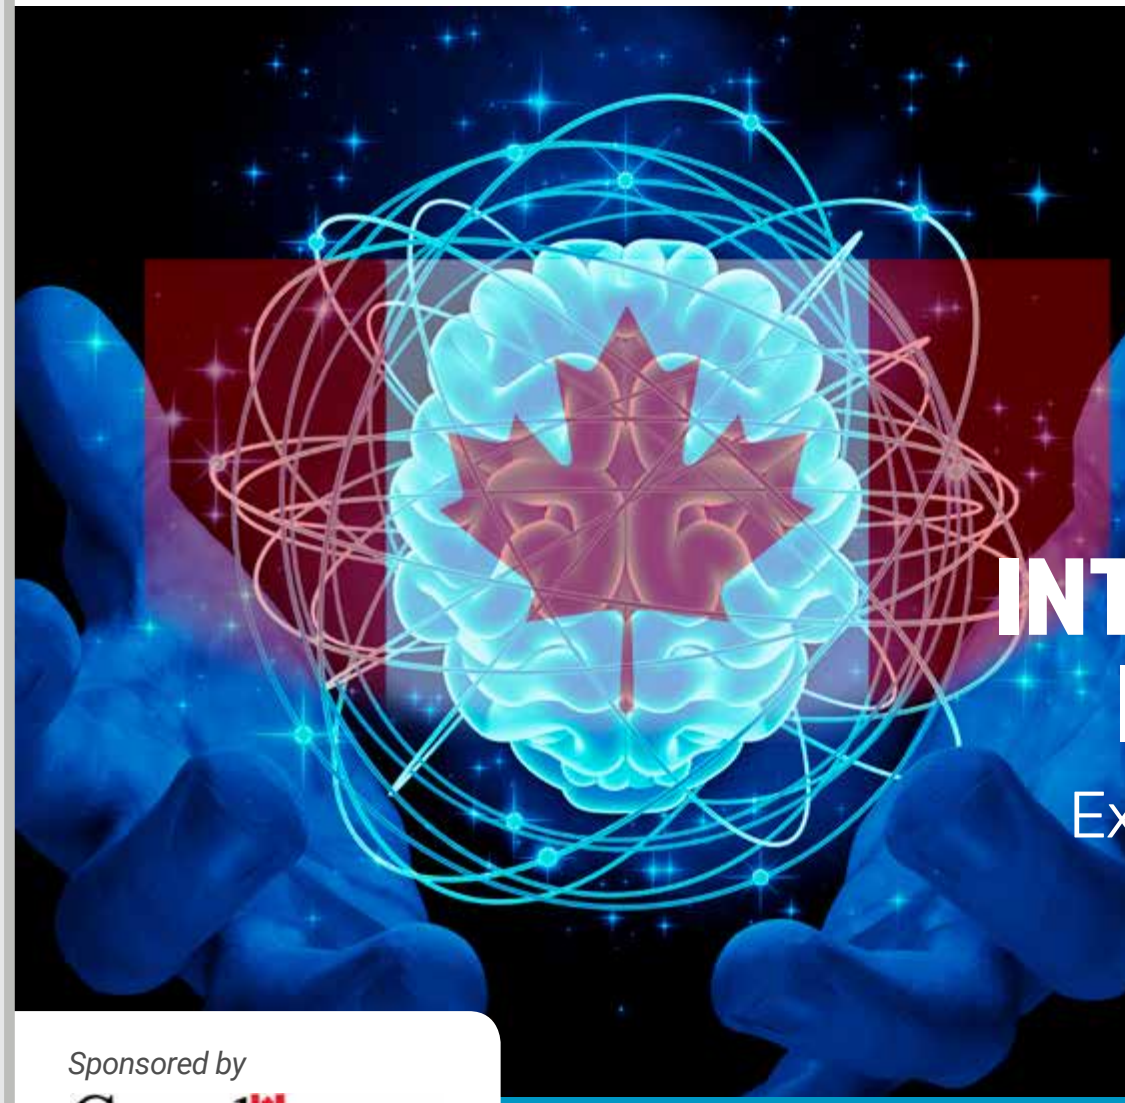

ARTIFICIAL INTELLIGENCE IN CANADA:

Executive Interviews on the State of AI

Sponsored by

Canada
Consulate General of Canada in Boston
Consulat général du Canada à Boston

Produced by
Cambridge Innovation Institute

#1 TOP AI BLOG
TOP 50 BLOGS

aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

ELIOT WEINMAN
CHAIR/FOUNDER,
AI WORLD, EXECUTIVE EDITOR, AI TRENDS

FROM THE EDITOR

Artificial Intelligence in Canada

The race for nations to become a global leader in artificial intelligence (AI) has begun. Canada's federal government was the first country to release an AI Strategy, setting aside $125 million in funding with the aim of promoting collaboration between three major academic centers of AI expertise–Montreal Institute for Learning Algorithms (MILA), Alberta Machine Intelligence Institute (AMII) and Vector Institute. A majority of this funding, along with $50 million from the Ontario government and $80 million from the private sector, helped establish the new Vector Institute, an independent research facility for AI located in Toronto. The Quebec government has also invested $100 million in Montreal's AI community, while the Canada First Research Excellence Fund donated $93.6 million to Université de Montréal, Polytechnique Montréal and HEC Montréal for cutting-edge research in deep learning.

As a result, Canada has become a hotbed of AI activity where the public sector, academic community and private sector are working together to advance the possibilities of AI. This a the first of a special three-part report on AI in Canada.

The Canadian government is an International Sponsor of AI World, the industry's largest independent conference and expo covering the state of the practice in enterprise AI. Be sure to attend AI World, Dec 3-5, 2018 in Boston – aiworld.com – and hear from some of these leaders and startups, and learn first-hand how your enterprise can take advantage of all that Canada's AI initiatives have to offer.

ABOUT AI TRENDS | AITrends.com

AI Trends is the leading industry media channel focused on the business and technology of enterprise AI. It is designed for business executives wishing to keep track of the major industry business trends, technologies and solutions that can help them keep in front of the fast-moving world of AI and gain competitive advantage.

This index is provided as an additional service. The publisher does not assume any liability for errors or omissions.

Subscriptions: Address inquires to Cambridge Innovation Institute, 250 First Avenue, Suite 300, Needham, MA 02494 888-999-6288 or e-mail kfinnell@cambridgeinnovationinstitute.com

Reprints: Copyright © 2018 by AI Trends, All rights reserved. Reproduction of material printed in AI Trends is forbidden without written permission. For reprints and/or copyright permission, please contact Lisa Scimemi, 781-972-5446, lscimemi@cambridgeinnovationinstitute.com.

CONTACTS:

PRESIDENT
Phillips Kuhl

PUBLISHER
Lisa Scimemi
lscimemi@cambridgeinnovationinstitute.com

EXECUTIVE EDITOR
Eliot Weinman
ew@aiworld.com

EDITOR
John Desmond
jd@aiworld.com

VICE PRESIDENT, BUSINESS DEVELOPMENT
Angela Parsons
aparsons@cambridgeinnovationinstitute.com

MARKETING ASSOCIATE
Tom Conley
tconley@cambridgeinnovationinstitute.com

MARKETING ASSISTANT
Brooke Trefry
btrefry@cambridgeinnovationinstitute.com

MARKETING ASSISTANT
Renata Katz
rkatz@cambridgeinnovationinstitute.com

CONTACT INFORMATION
info@aitrends.com

250 First Avenue, Suite 300
Needham, MA 02494

CONNECT WITH US:

#1 TOP AI BLOG
TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

ARTIFICIAL INTELLIGENCE IN CANADA

CANADA'S AI INITIATIVE

Bringing Together Government, Academia, Industry in a Quest to Expand the National Economy

BY JOHN P. DESMOND
AI TRENDS EDITOR

JUNE 26, 2018

Canada has emerged as a front-runner in AI development. The country has declared a national initiative around AI, pouring hundreds of millions of dollars into it and relaxing immigration rules to attract the best engineers.

Canada has a number of strengths that position the country well in AI, including a growing supply of Ph.D-level AI engineers in an era characterized by an acute global shortage of AI talent. This is largely due to leading universities in Toronto, Waterloo, Montreal and Edmonton that have invested heavily in AI research for years. "Canada now has one of the most significant concentrations of AI talent anywhere," stated Salim Teja, venture capitalist with MaRS, an innovation hub in Toronto, writing recently in Forbes.

CONTINUED

"CANADA NOW HAS ONE OF THE MOST SIGNIFICANT CONCENTRATIONS OF AI TALENT ANYWHERE."

SALIM TEJA, VENTURE CAPITALIST

#1 TOP AI BLOG
TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

ARTIFICIAL INTELLIGENCE IN CANADA

Graduates from Canada have gone on to lead labs in the world's top technology companies. They include: Geoffrey Hinton, VP and Engineering Fellow at the AI research team Google Brain, which also includes many other Canadians as senior researchers; Craig Boutilier, principal scientist at Google Research, which was recently renamed to Google AI; Alex Graves and Yujia Li, senior researchers at Google DeepMind, the AI research unit based in London, acquired by Google in 2014; Ilya Sutskever, co-founder and research director at Open AI, the non-profit AI research company that aims to promote "friendly AI;" Rusian Salakhutdinov, director of AI Research at Apple and a computer science professor at Carnegie Mellon University; Yann LeCun, VP and Chief AI Scientist at Facebook; and Nebojsa Jojic, Principal Researcher at Microsoft Research.

Here are some recent developments:

- **VECTOR INSTITUTE**, a $170 million research center headed by Geoffrey Hinton (who some call the "godfather of AI") recently opened in Toronto. Hinton stated, "Increasingly, the world's most promising researchers in deep learning and other AI subfields are looking at Canada as a hub with many opportunities to collaborate."
- **GEOFFREY GORDON** , a researcher at Carnegie Mellon University's machine learning department, recently announced he is moving to Montreal to lead its Microsoft Research Lab;
- **GOOGLE'S DEEPMIND**, which developed the first software to defeat a human player in the strategy board game Go, has opened its first lab outside of the UK in Edmonton.
- Microsoft in early 2017 acquired **MALUUBA**, a Montreal and Waterloo-based startup specializing in machine language understanding. Microsoft plans to double the size of the company's Montreal office by 2019.
- **TECH MAHINDRA**, an IT services firm based in India, announced in February 2018 it will invest $100 million in a Centre of Excellence in Toronto. The center plans to jointly develop business solutions in emerging technologies including AI with academic institutes, innovators and accelerators in the Canadian startup ecosystem. Canada's Prime Minister Justin Trudeau recently traveled to India to promote enhance bilateral collaboration between the two countries.

Google DeepMind

Geoffrey Hinton, left, heads the recently opened Vector Institute in Toronto; Geoffrey Gordon, center, will lead the Microsoft Research Lab in Montreal. Google's DeepMind has opened a lab in Edmonton.

Canada also has vast amounts of data, by virtue of its publicly-funded systems in services including health, energy and transit. "These large datasets are goldminds for AI developers, not only because of their size but also because of their diversity," states Teja. Canadian diversity is a strength when it comes to address bias in AI datasets. In Toronto, more than 50% of the population is foreign-born.

Rather than trying to poach the best AI academics, startups and corporations in Canada are creating partnerships with universities that enable the researchers to stay involved in academia as well as industry. For example, the Royal Bank of Canada has created Borealis AI, a subsidiary that gives its staff latitude to collaborate with universities on research and jointly publish results.

In Montreal, Element AI is a venture founded by neural network pioneer

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

ARTIFICIAL INTELLIGENCE IN CANADA

Yoshua Bengio, working with University of Montreal (see recent AI Trends interview here) and McGill University researchers to jointly advance algorithms for companies looking to incorporate AI solutions into their business.

CANADIAN GOVERNMENT INVESTING, STARTUPS RAISING MONEY

PWC estimates that AI will add $19.4 trillion to the global economy by 2030. Canada sees AI as its biggest opportunity to expand its economy, alleviate its dependence on natural resources, evolve to new technology-driven industries and create a prosperous future, writes Steve Irvine recently in the Toronto Globe and Mail. Irvine in early 2017 left an executive role at Facebook to return to Canada and build integrate.ai, an AI-focused startup where he is the CEO.

In 2017, Canada invested more than $300 million in new funding for research and Canadian startups raised $260 million. In addition to DeepMind, global tech giants Facebook, Samsung and Uber have established AI research labs in Canada.

GLOBAL COMPETITION

Canada faces competition for AI leadership. The US remains home to the world's leading AI companies including Google, Facebook, Apple, Amazon, many well-funded startups and strong universities including MIT, Stanford and Carnegie Mellon.

China plans to be a world AI leader by 2030; China recently committed $2.5 billion to a national AI research park in Beijing with a goal of supporting 400 companies producing revenue of $9.8 billion annually. With its own technology giants Alibaba, Tencent and Baidu, and its large population, China is well-positioned to capture a large share of the global AI market opportunity.

China seemed to mimic the Obama-era roadmap laid out in 2016 for AI leadership. As the New York Times reported in February 2018, the Trump White House has been silent on a national AI strategy. 'We are still waiting on the White House to provide some direction" on how to respond to the competition, stated Tim Hwang, who worked at AI policy at Google and is now the director of the Ethics and Governance of AI Initiative, an organization created by LinkedIn founder Reid Hoffman and others to fund ethical research in AI.

Canada on the other hand is marshalling its national leadership to lead the charge. "It is essential that Canada shifts its focus more aggressively to commercialization from academic prowess," states Irvine. Canada today lacks the homegrown, global-leading technology companies that can capture the economic growth associated with the opportunity. "We need at least a handful of $10-billion-plus companies to help create the commercial density and global relevance needed to attract and retain elite talent, create new high-value jobs and capture economic gain," Irvine stated.

"WE NEED AT LEAST A HANDFUL OF $10-BILLION-PLUS COMPANIES TO HELP ... CREATE NEW HIGH-VALUE JOBS AND CAPTURE ECONOMIC GAIN"

STEVE IRVINE, TORONTO GLOBE AND MAIL

He encouraged Canadian businesses to commit to AI projects engaging startups to tackle meaningful business problems, to put strong people on it, and not tie them up for months with legal and procurement issues. "There is no better learning agenda for your team than rolling up their sleeve and seeing this new technology in practice delivering results for you," he states.

AI TRENDS SERIES ON AI INNOVATION IN CANADA

The Canadian government recently hosted journalists including Eliot Weinman, Executive Editor of AI Trends, on a tour of Canada to learn about the country's accomplishments and opportunities in AI. The group visited universities and companies, and heard from experts in AI development and research, from innovators at AI startups and executives of global companies. As a follow up, AI Trends is producing a series of articles on AI initiatives in Canada centered around Toronto, Montreal, Edmonton and Vancouver.

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW

Dr. Ajay Agrawal

is the Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management.

CREATIVE DESTRUCTION LAB

Mission of Creative Destruction Lab Founder is to Create A Canadian AI Ecosystem

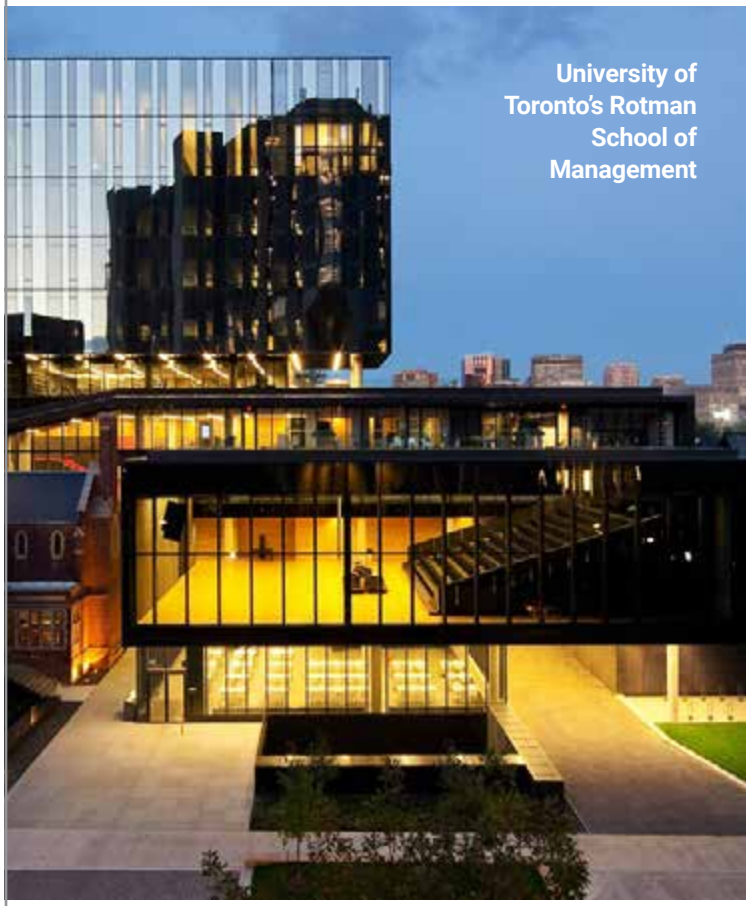

University of Toronto's Rotman School of Management

Describe what happens at the Creative Destruction Lab (CDL).

The CDL is a seed-stage program for massively scalable science-based companies. Some start-ups come from the University of Toronto community, but we now also receive applications from Europe, the U.S. (including Silicon Valley), Israel and Asia.

We launched the program in September 2012, and each autumn since, we've admitted a new cohort of start-ups into the program. Most companies that we admit have developed a working prototype or proof of concept. The most common type of founder is a recently graduated PhD in Engineering or Computer Science who has spent several years working on a problem and has invented something at the frontier of their field.

The program does not guarantee financing, but the majority of companies that succeed raise capital from the CDL's Fellows and

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: DR. AJAY AGRAWAL

Associates — a carefully- selected group of individuals who themselves are serial entrepreneurs and early-stage investors. Throughout the year, our MBA students work with the start-up founders as part of a second-year elective course, helping them develop financial models, evaluate potential markets, and fine-tune their scaling strategies.

To date, more than 100 start-ups have come through the Lab. When we launched, we set a goal of generating $50 million in equity value created in terms of the aggregate value created by companies that went through the Lab. When we finished our fifth year in June 2017, we had exceeded $1.4 billion in equity value created.

What exactly does the Lab provide to entrepreneurs?

Start-up founders benefit from a structured, objectives-oriented process that increases their probability of success. The process

"WHEN WE FINISHED OUR FIFTH YEAR IN JUNE 2017, WE HAD EXCEEDED $1.4 BILLION IN EQUITY VALUE CREATED."

is orchestrated by the CDL team, while CDL Fellows and Associates generate the objectives. Objective-setting is a cornerstone of the process. Every eight weeks the Fellows and Associates set three objectives for the start-ups to achieve, at the exclusion of everything else. In other words, they define clear goals for an eight week 'sprint'. Objectives can be business, technology or HR-oriented. Our Fellows and Associates—all volunteers—are critical to the CDL's success.

Tell us more about the CDL Fellows and Associates.

We have designed a marketplace — a community that functions under a set of rules and norms — that facilitates efficient transactions between first-time founders and experienced entrepreneurs, many of whom are also investors. Often, the two sides don't know each other until the rookie founder seeks out the experienced entrepreneur/investor when raising capital. Knowing very little about the entrepreneur, the investor usually says no, but occasionally says yes, at which point they are very committed. By the time we hit the end of the academic year, the Fellows and Associates have met with the ventures many times, and they've gotten to know each other. Furthermore, the entrepreneurs have demonstrated their ability (or inability) to deliver against an aggressive set of objectives through several cycles. We don't require Fellows and Associates to invest, but they can — and many do.

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW

Dr. Russell Greiner

Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

amii

Machine Learning Tools, Metabolics, Probablilty and

SURVIVAL ANALYSIS

After earning a PhD from Stanford, Russ Greiner worked in both academic and industrial research before settling at the University of Alberta, where he is now a Professor in Computing Science and the founding Scientific Director of the Alberta Innovates Centre for Machine Learning (now Alberta Machine Intelligence Institute), which won the ASTech Award for "Outstanding Leadership in Technology" in 2006. He has been Program Chair for the 2004 "Int'l Conf. on Machine Learning", Conference Chair for 2006 "Int'l Conf. on Machine Learning", Editor-in-Chief for "Computational Intelligence", and is serving on the editorial boards of a number of other journals. He was elected a Fellow of the AAAI (Association for the Advancement of Artificial Intelligence) in 2007, and was awarded a McCalla Professorship in 2005-06 and a Killam Annual Professorship in 2007. He has published over 200 refereed papers and patents, most in the areas of machine learning and knowledge representation, including 4 that have been awarded Best Paper prizes. The main foci of his current work are (1) bioinformatics and medical informatics; (2)

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: DR. RUSSELL GREINER

learning and using effective probabilistic models and (3) formal foundations of learnability. He recently spoke with AI Trends.

Q: Who do you collaborate with in your work?

I work with many very talented medical researchers and clinicians, on projects that range from psychiatric disorders, to stroke diagnosis, to diabetes management, to transplantation, to oncology, everything from breast cancer to brain tumors. And others — I get many cold-calls from yet other researchers who have heard about this "Artificial Intelligence" field, and want to explore whether this technology can help them on their task.

Q: How do you see AI playing a role in the fields of oncology, metabolic disease, and neuroscience?

There's a lot of excitement right now for machine learning (a subfield of Artificial Intelligence) in general, and especially in medicine, largely due to its many recent successes. These wins are partly because we now have large data sets, including lots of patients — in some cases, thousands, or even millions of individuals, each described using clinical features, and perhaps genomics and metabolomics data, or even neurological information and imaging data. As these are historical patients, we know which of these patients did well with a specific treatment and which ones did not.

I'm very interested in applying supervised machine learning techniques to find patterns in such datasets, to produce models that can make accurate predictions about future patients. This is very general — this approach can produce models that can be used to diagnose, or screen novel subjects, or to identify the best treatment — across a wide range of diseases.

It's important to contrast this approach with other ways to analyze such data sets. The field of biostatistics includes many interesting techniques to find "biomarkers" — single features that are correlated with the outcomes — as a way to try to understand the etiology, trying to find the causes of the disease. This is very interesting, very relevant, very useful. But it does not directly lead to models that can decide how to treat Mr. Smith when he comes in with his particular symptoms.

At a high level: I'm exploring ways to find personalized treatments — identifying the treatment that is best for each individual. These treatment decisions are based on evidence-based models, as they are learned from historical cases — that is, where there is evidence that the model will work effectively.

In more detail, our team has found patterns in neurological imaging, such as functional MRI scans, to determine who has a psychiatric disorder — here, for ADHD, or autism, or schizophrenia, or depression, or Alzheimer's disease.

Another body of work has looked at how brain tumors will grow by looking at brain scans of people, using standard structural MRI imaging. Other projects learn screening models that determine which people have adenoma (from urine metabolites), or models that predict which liver patients will most benefit from a liver transplant (from clinical features), or which cancer patients will have cachexia, etc.

Q: How can machine learning be useful in the field of Metabolomics?

Machine learning can be very useful here. Metabolomics has relied on technologies like mass spec and NMR spectroscopy to identify and quantify small molecules in a biofluid (like blood or urine); this previously was done in a very labor-intensive way, by skilled spectroscopists.

My collaborator, Dr. Dave Wishart (here at the University of Alberta) and some of our students, have designed tools to automate this process — that can effectively find the molecules present in say blood. This means metabolic profiling is now high-throughput and automated, making it relatively easy to produce datasets that include the metabolic profiles from a set of patients, along with their outcome. Machine learning tools can then use this labeled dataset to produce models for predicting who has a disease, for screening or for diagnosis. This has led to models that can detect cachexia (muscle wasting) and adenoma (with a local company, MTI).

Q: Can you go in to some detail on the work you have done designing algorithms to predict patient-specific survival times?

This is my current passion; I'm very excited about it.

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

The challenge is building models that can predict the time until an event will happen — for example, given a description of a patient with some specific disease, predict the time until his death (that is, how long he will live). This seems very similar to the task of regression, which also tries to predict a real value for each instance –for example, predicting the price of a house based on its location, the number of rooms, and their sizes, etc.. Or given a description of a kidney patient (age, height, BMI, urine metabolic profile, etc.), predict the glomerular filtration rate of that patient, a day later.

Survival prediction looks very similar because both try to predict a number for each instance. For example, I describe a patient by his age, gender, height, and weight, and his genetic information, and metabolic information, and now I want to predict how long until his death — which is a real number.

The survival analysis task is more challenging due to "censoring". To explain, consider a 5 year study that began in 1990. Over these five years, many patients have passed away, including some who lived for three years, others for 2.7 years, or 4.9 years. But many patients didn't pass away during these 5 years –which is a good thing… I'm delighted these people haven't died! But this makes the analysis much harder: for the many patients alive at the end of the study, we know only that they lived at least 5 years, but we don't know if they lived 5 years and a day or lived 30 years — we don't know and never will know.

This makes the problem completely different from the standard regression tasks. The tools that work for predicting glomerular filtration rate or for predicting the price of a house just don't apply here. You have to find other techniques. Fortunately, the field of survival analysis provides many relevant tools. Some tools predict something called "risk", which gives a number to each patient, with the understanding that this tool is predicting that patients with higher risks will die before those with lower risk. So if Mr A's risk for cancer is 7.2 and Mr B's is 6.3 — that is, Mr A has a higher risk — this model predicts that Mr A will die of cancer before Mr B will. But does this mean that Mr A will die 3 days before Mr B, or 10 years — the risk score doesn't say.

Let me give a slightly different way to use this. Recall that Mr A's risk of dying of cancer is 7.2. There are many websites that can do "what if" analysis: perhaps if he stops smoking, his risk reduces to 5.1. This is better, but by how much? Will this add 2 more months to his life, or 20 years? Is this change worth the challenge of not smoking?

Other survival analysis tools predict probabilities — perhaps Ms C's chance of 5-year disease-free survival, is currently is 65%. but if she changes her diet in certain way, this chance goes up to 78%. Of course, she wants to increase her five-year survival. But again, this is not as tangible as learning, "If I continue my current lifestyle then this tool predicts I will develop cancer in 12 years, but if I stop smoking, it goes from 12 to 30 years". I think this is much more tangible, and hence will be more effective in motivating people to change their lifestyle, versus changing their risk, or their 5-year survival probability.

THE SURVIVAL ANALYSIS TASK IS MORE CHALLENGING DUE TO "CENSORING."

So my team and I have provided a tool that do exactly that, by giving each person his or her individualized survival curve, which shows that person's expected time to event. I think that will help motivate people to change their lifestyle. In addition, my colleagues and I also applied this to a liver transplant dataset, to produce a model that can determine which patient with end-stage liver failure, will benefit the most from a new liver, and so should be added to the waitlist.

Those examples all deal with time to death, but in general, survival analysis can deal with time to event, for any event. So it can be used to model a patient's expected time to re-admission. Here, we can seek a model that, given a description of a patient being discharged from a hospital, can predict when that patient will be readmitted — eg, if she will return to the hospital, for the same problem, soon or not.

Imagine this tool predicted that, given Ms. Jones' current status, if she leaves the hospital today, she will return within a week.

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

But if we keep her one more day and give some specific medications, we then predict her readmission time is 3 years. Here, it's probably better to keep her that one more day and give one more medication. It will help the patient, and will also reduce costs.

Q: What do you see are the challenges ahead for the healthcare space in adopting machine learning and AI?

There are two questions: what machine learning can do effectively, and what it should do.

The second involves a wide range of topics, including social, political, and legal issues. Can any diagnostician — human or machine — be perfect? If not, what are the tradeoffs? How to verify the quality of a computer's predictions? If it makes a mistake, who is accountable? The learning system? Its designer? The data on which it was trained? Under what conditions should a learned system be accepted? … and eventually incorporated into standard of care? Does the program need to be "convincing", in the sense of being able to explain its reasoning — that is, explain why it asked for some specific bit of information? … or why it reached a particular conclusion? While I do think about these topics, I am not an expert here.

My interest is more in figuring what these systems can do — how accurate and comprehensive can they be? This requires getting bigger data sets — which is happening as we speak. And defining the tasks precisely — is the goal to produce a treatment policy that works in Alberta,

MY INTEREST IS MORE IN FIGURING WHAT THESE SYSTEMS CAN DO – HOW ACCURATE AND COMPREHENSIVE CAN THEY BE? .. TO PRODUCE A TREATMENT POLICY THAT WORKS IN ALBERTA, OR THAT WORKS FOR ANY PATIENT, ANYWHERE IN THE WORLD?

or that works for any patient, anywhere in the world? This helps determine the diversity of training data that is required, as well as the number of instances. (Hint: building an Alberta-only model is much easier than a universal one.) A related issue is defining exactly what the learned tool should do: In general, the learned performance system will return a "label" for each patient — which might be a diagnosis (eg, does the patient have ADHD), or a specific treatment (eg, give a SSRI [that is, a selective serotonin reuptake inhibitor]). Many clinicians assume the goal is a tool that does what they do. That would be great if there was an objective answer, and the doctor was perfect, but this is rarely the case. First, in many situations, there is significantly disagreement between clinicians (eg, some doctors may think that a

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:

Dr. Ajay Agrawal

Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:

Dr. Russell Greiner

Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:

Dr. Foteini Agrafioti

Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:

Yoshua Bengio

MILA, University of Montreal

aitrends The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: DR. RUSSELL GREINER

specific patient has ADHD, while others may disagree) — if so, which clinician should the tool attempt to emulate? It would be better if the label instead was some objective outcome — such as "3 year disease-free survival", or "progression within 1 year" (where there is an objective measure for "progression", etc.)

This can get more complicated when the label is the best treatment — for example, given a description of the patient, determine whether that patient should get drug-A or drug-B. (That is, the task is prognostic, not diagnostic.) While it is relatively easy to ask the clinician what she would do, for each patient, recall that clinicians may have different treatment preferences… and those preferences might not lead to the best outcome. This is why we advocate, instead, first defining what "best" means, by having a well-defined objective score for evaluating a patient's status, post treatment. We then define the goal of the learned performance system as finding the treatment, for each patient, that optimizes that score.

One issue here is articulating this difference, between "doing what I do" versus optimizing an objective function. A follow-up challenge is determining this objective scoring function, as it may involve trading off, say, treatment efficacy with side-effects, etc. Fortunately, clinicians are very smart, and typically get it! We are making in-roads.

Of course, after understanding and defining this objective scoring function, there are other challenges — including collecting data from a sufficient number of patients and possibly controls, from the appropriate distributions, then building a model from that data, and validating it, perhaps on another dataset. Fortunately, there are an increasing number of available datasets, covering a wide variety of diseases, with subjects (cases and controls) described with a many different types of features (clinical, omics, imaging, etc etc etc). Finally comes the standard machine learning challenge of producing a model from that labeled data. Here, too, the future is bright: There are faster machines, and more importantly, I have many brilliant colleagues developing ingenious new algorithms, to deal with many different types of information.

All told, this is a great time to be in this important field! I'm excited to be a part of it.

Thank you Dr. Greiner!

Learn more at the Alberta Machine Intelligence Institute. https://aitrends.com/ai-in-canada/executive-interview-dr-foteini-agrafioti-head-of-borealis-ai-and-chief-science-officer-rbc/

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW

Dr. Foteini Agrafioti

Head of Borealis AI and Chief Science Officer, RBC

BOREALIS AI
RBC Institute for Research

Combining Fundamental and Applied AI Research

to Create Opportunity for Canada

Dr. Agrafioti is the Chief Science Officer at RBC and Head of Borealis AI. She is responsible for RBC's intellectual property portfolio in the fields of artificial intelligence and machine learning. Prior to joining Borealis AI, Foteini was the Chief Innovation Officer at Architech, where she led Research and Innovation. She also founded and served as Chief Technology Officer at Nymi, a biometrics security company and maker of the Nymi wristband.

Foteini is the inventor of HeartID, the first biometric technology to authenticate users based on their unique cardiac rhythms. She is a TED speaker and serves on the editorial review boards of several scientific journals. Foteini was named "Inventor of the Year" in 2012 at the University of Toronto where she received a Doctorate in Electrical and Computer Engineering and was named one of Canada's "Top 40 Under 40" for 2017. She recently spent a few minutes talking to AI Trends Editor John P. Desmond

Q. What is the background of Borealis AI?

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:

Dr. Ajay Agrawal

Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:

Dr. Russell Greiner

Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:

Dr. Foteini Agrafioti

Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:

Yoshua Bengio

MILA, University of Montreal

aitrends The Business and Technology of AI

Cambridge Innovation Institute

250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: DR. FOTEINI AGRAFIOTI

A. Borealis AI is a new group within RBC that started two-and-a-half years ago. The objective was to create an institute where researchers can do fundamental machine learning research after graduating from academia. There are two reasons for this. One is that Canada has done well technically in the AI space, having developed deep learning and reinforcement learning foundational techniques, largely out of the University of Toronto, the University of Montreal and the University of Alberta. However, we have found that many of the graduates who had been on the forefront of publishing and creating the research were deciding to leave the country and pursue opportunities elsewhere. There were not enough interesting jobs in Canada.

We thought that AI could be a niche for Canada, so we wanted to create jobs that would be interesting enough to retain these people and to have them continue to do AI research in Canada.

We are a data-driven business, from our retail banks to the capital markets. We use data to drive decision-making; we use machine learning algorithms in our applications. So Borealis is the place where we build our intellectual property in this space.

We started off in Toronto and now have around 70 people with five research centers: Toronto, Montreal, Edmonton and in the fall, we are opening centers in Vancouver and Waterloo. Each is led by a professor who wants to continue to teach and to help commercialize the technology.

Each center will divide its time between fundamental and applied research. For fundamental research, we are looking to advance the state-of-the-art in machine learning and in supervised learning with a human-in-the-loop approach. These are hot academic research topics. This group is dedicated to pushing the boundaries of science and publishing the results of our efforts.

The applied effort is from a combination of researchers and software engineers. This is where we use the algorithms to develop software. We apply them to challenges in the financial services industry and beyond. We don't have a sole mandate of delivering banking applications; we're interested also in climate change and in healthcare, for instance. But the primary focus is on RBC applications.

Q. What is the role of the Edmonton center?

A. Edmonton is the first research center we opened outside of Toronto. The reason was twofold. First, the group there at the University of Alberta and the Alberta Machine Intelligence Institute had done extremely well in investing in reinforcement learning and in proving it can work. They have built a strong practice on that. We wanted to strengthen that community. They did not have many surrounding technology companies. We are a Canadian company and we wanted to see our community prosper from coast to coast, and we wanted to create jobs right in Edmonton to support the university.

Second, reinforcement learning is extremely interesting to us. It's good for analyzing data. A lot of the data we deal with in banking requires pattern analysis. We hired Prof. Matt Taylor out of Washington State University to lead the Edmonton lab. We also work with private companies. We have an applied research practice in the financial sector.

The reason we expanded to Vancouver, which is a much larger city, was to create more opportunities for graduates. There is a strong computer vision specialty there, encompassing graphic and special effects, too. The University of British Columbia and Simon Fraser University have students advancing the state-of-the-art in video applications. We are expanding the data sources that we look to and are taking advantage of visual data.

Q. How many jobs do you anticipate being able to create in Edmonton?

A. Our goal is to create 160 jobs across the country by the end of 2019. We are at 70 today and we have 15 people in Edmonton at the moment. It's a matter of finding the right people in each location. We don't have a hiring number goal for each lab.

Q. How did you get into this work? What is your background?

A. I studied electrical engineering. It was at the University of Toronto that I got introduced to machine learning – biometric authentication in particular. We developed a way to authenticate people's identities using the human heartbeat. Nymi was a company I built that commercialized that research. It was a wearable device. So, academia plus startups is my background.

I decided to join RBC because I had seen that path to commercialization and felt we could be doing it better in Canada. The US is very good at commercialization. I saw

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: DR. FOTEINI AGRAFIOTI

a lot of responsibility at RBC for nurturing the well-being of the country. Our interests aligned very well in that respect. I feel very passionate about it. It's not straightforward to build groups like this within known technology companies. It's very unusual.

"We do not yet have practitioners of AI."

Q. What are the primary challenges?

A. I see two big ones. The first is recruiting talent, which is a very scarce resource, primarily because everyone who does AI today is coming from research, PhD programs and academia. We do not yet have practitioners of AI. That is really limiting when you are trying to build large teams, especially when you face the competition we do. There might be 20,000 people in the world who can do this job, and maybe 3,000 coming out of the universities each year. The competition is fierce.

The other challenge is that we are extremely regulated, for good reason. We have rules for how we can use AI to bring products to life. As you know, AI comes with concerns around issues including bias and privacy. These are the kinds of problems where there is active research and people are trying to solve these ethical problems. When we try to use these technologies in our products, we can hit walls because of how our industry functions. This is why we are so heavily invested in fundamental research. If bias in AI can be solved, we will be the first in line to make use of that.

We are enjoying the challenge. AI is moving so fast, and companies are adopting it so quickly. I don't know if it's good to always get into a race to bring new products to market. Sometimes it might make sense to step back and ask if the technology can be compromised. We have no other option to do business than the right way. So, it's an engineering problem you have to solve. It constrains you, but it's a good challenge.

Q. Do you have any advice for students interested in pursuing a career in AI for what they should study?

A. I would recommend computer science and machine learning. A career in AI can be very fulfilling. These technologies can be applied to many disciplines. Machine learning can be applied to finance, healthcare, consumer electronics, even the fashion industry. It's very diverse. It's exciting to be in this space right now. It's challenging and you don't have to limit yourself to one industry or field. This field will be pervasive as we move forward. We need more people who understand the core technology.

The other thing I would say is that students should consider coming to Canada to study. There is no better place in the world to learn. I have seen the support of the country for long-term investments in AI, such as in reinforcement learning, which has led to the creation of these new institutes [like Borealis AI]. We have attracted world-class talent from around the world. It's a very good place.

Q. Where are you from originally?

A. I am Greek. I came to Canada for grad school. I found it cold, then I got used to it.

Q. Can you talk about one of your AI projects?

"WE DO NOT YET HAVE PRACTITIONERS OF AI."

A. We have a project that uses AI to analyze news. Real-time information is important to our businesses. We make investment decisions based on it. If you ask a financial analyst what their job is, they will tell you, for instance, that their job is to know everything about healthcare in North America. But there's too much information and every human is limited in what they can absorb, and what languages they speak. So you need to invest in technology to help.

Apollo is the product we built as a result of these challenges. It is an agent that reads the news for you and alerts you to information it thinks will be of interest. It can look at individual companies, see historical patterns and predict a likely change. It is using natural language processing, deep learning and graphic rendering.

To learn more, go to borealisai.com

#1 TOP AI BLOG TOP 50 BLOGS

aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW

Yoshua Bengio

MILA, University of Montreal

MILA

Combining AI Research and Business Collaboration

Thoughts on the Impact of AI on Society

Yoshua Bengio is among the most cited Canadian computer scientists. He is the author of two books and more than 200 publications, the most cited being in the areas of deep learning, recurrent neural networks, probabilistic learning algorithms, natural language processing and manifold learning.

He earned a PhD in Computer Science from McGill University in 1991 and worked at the Canadian Institute for Advanced Research (CIFAR) alongside with Yann LeCun (now at Facebook) and Geoffrey Hinton (now at Google). He has collaborated with IBM in work on the Watson supercomputer.

His current interests are centered around a quest for AI through machine learning, and include fundamental questions on deep learning and representation learning, the geometry of generalization in high-dimensional spaces, manifold learning, biologically inspired learning algorithms, and challenging applications of statistical machine learning. He recently participated in an interview with journalists learning about Canada's AI initiative that included Eliot Weinman, Executive Editor of AI Trends.

Q. Why is it important for the Canadian government to engage in this AI initiative?

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: YOSHUA BENGIO

A. AI is not just another technology. It will have a big impact on our societies, and there are many ethical and social questions associated with how AI is being deployed and how it will be deployed. If we don't think about these considerations, the public will eventually reject advanced technologies that they see as threatening and against their well-being. So governments have to really care about these questions, whether for moral reasons or for practical reasons.

Q. What would be your AI horror scenario?

A. I am most concerned about the use of AI in the military and security arenas. I'm sure you've heard about killer robots, and you may have also heard of how the technology can be used to recognize people from their facial images. So there are Big Brother scenarios that could be upon us if we're not careful. I also have concerns related to privacy issues when we are dealing with private data. Then we have economic issues. Automation will be accelerated with AI; that may create more inequality than we already suffer. And that is at the level of people, companies and countries. To have more countries involved will create a healthier playing field.

Q. What is the role of universities in the evolution of AI?

A. I am a professor at University of Montreal; we have created MILA (Montreal Institute for Learning Algorithms), which is similar to the Vector Institute (collaboration of government and business in partnership with University of Toronto) which has similar goals and were both funded by the federal government and the provincial government. These institutes – there's another in Alberta (Alberta Machine Intelligence Institute) – have been set up so they will have more agility than universities have, but they're still academic research organizations. They also have a mandate to help the ecosystem through the startups and companies that are creating value with AI.

These institutes are in a better position to be neutral about how AI will be used and keep in mind the well-being of people, and to orient research in directions that will be good for people, and engage in the public dialogue in a credible way. I think it's good that companies like Facebook and Google participate in that dialogue, but I'm not sure if they are neutral agents in those discussions. Universities, which care, first and foremost, about the public good, are really important agents in the discussions and in the kind of research that can be done.

Q. What steps can government take to foster this dialogue?

A. Here in Montreal, we are creating an organization that will be focused on the social, economic and ethical questions around AI. It will sponsor research in the social sciences and humanities around AI, but also will participate in the public debate. I think we don't have all the answers to how to do this right. Scholars and scientists need to really think through this and engage the public. We did something like this in the last six months in Montreal and in Quebec, and also in Ontario. After a forum of experts, we brought in ordinary people. We went to public libraries and places where people could comment and discuss the questions. We're coming up with something that will be initiated by scholars and experts, and

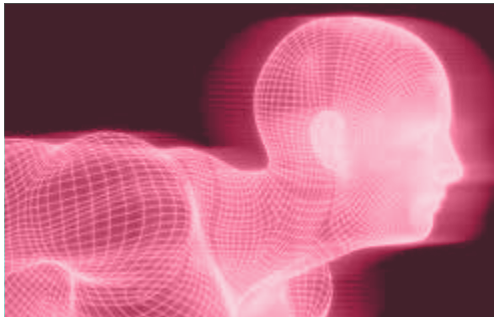

also have feedback and contributions from ordinary people. I think we have to continue in that direction.

This observatory on AI will be in a good position to make recommendations to governments, which will be part of the mission both locally and in different countries. The questions are pretty much the same in most countries. I think there should be a global coordination about these questions. There are issues like military use which will obviously need to be international, and even questions about regulating companies, which are multinationals. It would be much better if we can agree on principles globally.

Q. What do you see as the next evolution of the core technology that enables what we know of as AI today?

A. I'm a scientist. I don't have a crystal ball. And I can make educated guesses like many people. But one thing for sure is that there are obstacles on our way towards smarter machines, and it's always been like this when we make progress. We've achieved something important, and now we see that there are other challenges. We've made huge progress in industry using supervised

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: YOSHUA BENGIO

learning where humans have to really teach machines by telling them what to do. A lot of the current emphasis in basic research is on supervised learning or reinforcement learning, where the machines have to learn in a more autonomous way. And we haven't solved that in a satisfactory way yet. It will probably take years, or decades to really make big breakthroughs there. But given the exponential growth of research in these areas, I'm very optimistic that things will move very swiftly.

Q. Are you concerned that the massive investments in AI today are too risky?

A. One reason why companies are investing so much, and are so optimistic, is that a lot of future wealth growth from AI doesn't depend on new discoveries. In other words, we take what we have already scientifically, and we just make a lot of progress in the hardware. That's going to happen. It's moving. We will make progress in bringing together the right data. Like medical data, we don't do a good job yet. In lots of industries and sectors, the ingredients for applying that science are not there yet, but they will be there soon.

We have at least a decade to just reap the benefits of the science we already have. On top of that, there's so much money being poured into research, both in industry and in academia, that it would be surprising if the science doesn't move forward over the next decade. So it's almost a sure gain. Now, of course, you know, commercial enterprises can fail for all kinds of reason. But at a high level, I think it's a very safe bet.

Q. Is China ahead in the race to be the leading AI country?

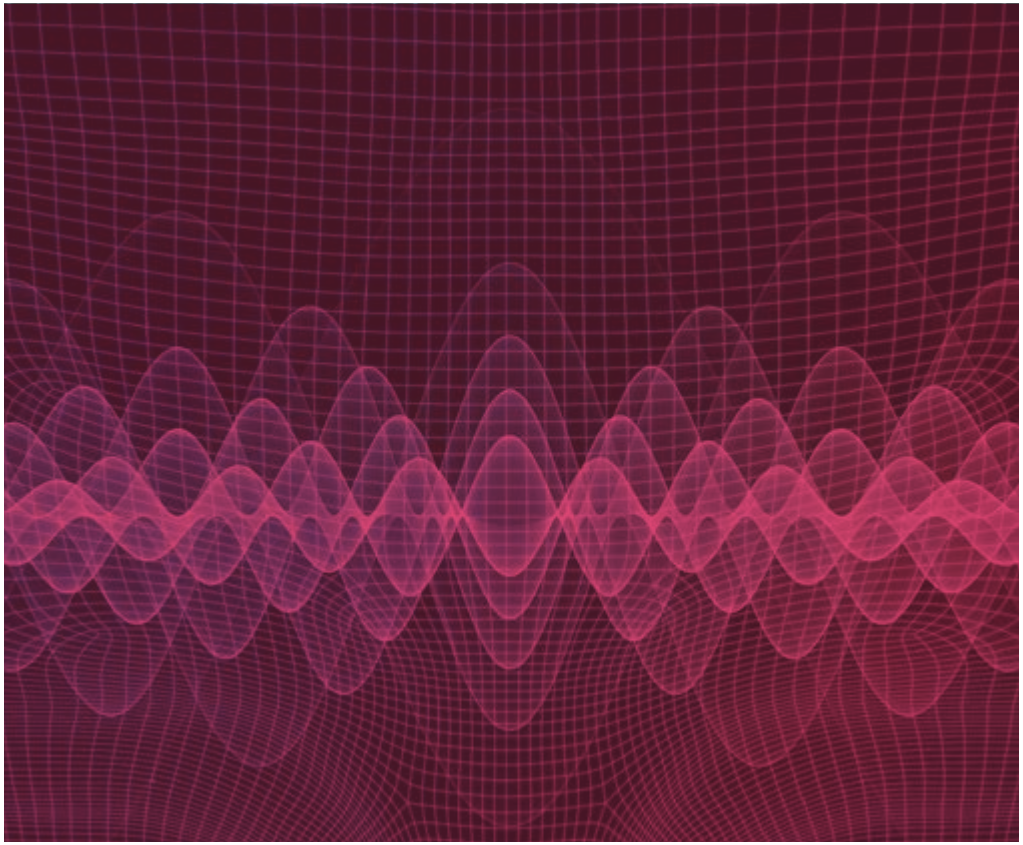

... WE TAKE WHAT WE HAVE ALREADY SCIENTIFICALLY, AND WE JUST MAKE A LOT OF PROGRESS IN THE HARDWARE. THAT'S GOING TO HAPPEN. IT'S MOVING.

A. I don't like to make like these kinds of comparisons. Silicon Valley is a very small place. The progress can come from anywhere in the world. China does have huge advantages in this race. One of the most important ones is that it's the biggest market in the world, and has the volumes of data that go with that. So from the point of view of investing, this is a very appealing place to do AI. And in addition, there's a huge enthusiasm for AI in China from all quarters. And lots and lots of students are jumping

CONTINUED

#1 TOP AI BLOG TOP 50 BLOGS
aiworld CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:
Dr. Ajay Agrawal
Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:
Dr. Russell Greiner
Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:
Dr. Foteini Agrafioti
Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:
Yoshua Bengio
MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

EXECUTIVE INTERVIEW: YOSHUA BENGIO

into this. It's a worldwide phenomenon, but I think with all the enthusiasm behind it, China probably wins the race for now.

Q. Do you envision big companies and startups and small companies collaborating to advance AI?

A. There is room for many kinds of business models in this new world. Large companies have leadership strong enough to make the fast turns that are needed, and companies like Element AI can help with that that. And big companies will be in competition with up and running small companies, building new products and new services which may not even exist now. New markets will be created. I'm also a big believer in the collaboration between startups and large companies. They have complimentary advantages. This is important from the point of view of a country with a national strategy because the startups are more agile. They can more easily recruit people who are excited about the fast pace of development. They can recruit talent more easily.

But the large companies have the larger market where they can deploy. They have lots of cash to invest, and they have lots of data. Ideally companies, a little bit like researchers, learn to cooperate better with their strengths and weaknesses to build something stronger.

Q. Are you concerned about the risk of jobs lost to AI automation?

A. Absolutely. The potential impact on the job market is very serious. It's not going to happen in one day, but it will happen way too fast for our ability to handle those changes. Many people are likely to lose their jobs in the middle of their careers.

We have to rethink our social safety net. Most developed countries have a social safety net, but it's been designed for a particular kind of economy. We will need to look into things like a universal basic income, and do more pilots. We may have to forget about our traditional values around work, such as if you don't work, you don't get money. And that's only one aspect of it. We need to rethink the education system so people can be rescaled in the middle of their career, while they are at a job.

The education system will need to train people in a way that is more appropriate for a fast-changing world, where human skills are going to be more important than they were in the past. Of course, we want to train more scientists and engineers; that's a no-brainer. But we have to train people not for one job that's very, very specialized, but rather how to think for themselves about how to be good citizens, and to rapidly learn the skills they need.

And we have to ask what is going to be the impact on society? Will AI be beneficial for the whole society or just a few people? I don't have the answers but I think it's important to ask the questions not and let the market by themselves figure out the answers. Those answers might not be in favor of ordinary people. Governments need to think about this and if necessary, find the right regulations.

Q. How is MILA progressing and can you describe your typical day?

A. MILA is the Montreal Institute for Learning Algorithms, a machine learning research lab with business collaboration as part of the mission. It's growing very fast. It already has the highest concentration of deep learning researchers in academia in the world. We're going to be doubling the number of professors over the next few years, thanks to the Canadian government.

MILA is mostly academic in nature, a non-profit, but with the mandate to help companies, to guide them in their development of AI.

I love working at the university. It allows me to be a more neutral agent in the changes that are coming, and gives me a voice that can have an impact as we adapt to this changing world of AI. Also, I'm in a position to steer research in directions that I think are important, and to contribute to the training of the next generation. I think this is something really, really important. I just enjoy the research with all of my students, which I would lose if I went to private industry.

Learn more at MILA.

#1 TOP AI BLOG
TOP 50 BLOGS

aiworld
CONFERENCE & EXPO

EDITOR'S LETTER

FEATURED ARTICLES

Canada's AI Initiative to Expand National Economy

EXECUTIVE INTERVIEW:

Dr. Ajay Agrawal

Peter Munk Professor of Entrepreneurship, Professor of Strategic Management, and founder of the Creative Destruction Lab at the University of Toronto's Rotman School of Management

EXECUTIVE INTERVIEW:

Dr. Russell Greiner

Professor CS and founding Scientific Director of the Alberta Machine Intelligence Institute

EXECUTIVE INTERVIEW:

Dr. Foteini Agrafioti

Head of Borealis AI and Chief Science Officer, RBC

EXECUTIVE INTERVIEW:

Yoshua Bengio

MILA, University of Montreal

aitrends
The Business and Technology of AI

Cambridge Innovation Institute
250 First Avenue, Suite 300
Needham, MA 02494
cambridgeinnovationinststute.com

aitrends
The Business and Technology of Enterprise AI

#1 TOP AI BLOG
TOP 50 BLOGS

News, resources, AI webinars, conference discounts, research and more!

AITrends.com

SUBSCRIBE TODAY

aiworld
CONFERENCE & EXPO
Accelerating Innovation in the Enterprise

AIWorld.com