

Introduction to the Economics of Artificial Intelligence

September 2018

Avi Goldfarb
University of Toronto and NBER

Based on the book “The Economics of Artificial Intelligence: An Agenda”
Ajay Agrawal, Joshua Gans, and Avi Goldfarb, Editors

What is Artificial Intelligence?



Defining Artificial Intelligence

- Oxford English Dictionary definition: “the theory and development of computer systems able to perform tasks normally requiring human intelligence.”
- A moving target!

Daniel Kahneman

Nobel prize winner for behavioral economics



A well-known novelist wrote me some time ago that he's planning a novel. The novel is about a love triangle between two humans and a robot and what he wanted to know is how would the robot be different from the people. I proposed three main differences.

One is obvious: the robot will be much better at statistical reasoning and less enamored with stories and narratives than people are.

The other is that the robot would have much higher emotional intelligence.

The third is that the robot would be wiser... *I do not think that there is very much that we can do that computers will not eventually [learn] to do.*"

Automation

- A process is automatic if it is performed without human assistance.
- Automation occurs when a process previously performed by humans becomes automatic.
- Fear of massive job losses largely driven by the potential of AI to automate processes.

Microsoft's System of Intelligence

AI = Domain Structure + Data Generation + General Purpose ML

Business Expertise
Structural Economics
Relaxations and Heuristics

Reinforcement Learning
Big Data Assets
Sensor/Video Tracking

Deep Neural Nets
Video/Audio/Text
OOS + SGD + GPUs


Figure 1: AI systems are self-training structures of ML predictors that automate and accelerate human tasks.

- Domain structure “allows you to break a complex task into composite tasks that can be solved with ML”
- Data generation in AI systems “requires an active strategy to keep a steady stream of new and useful information flowing into the composite learning algorithms”
- **“The current instance of AI is ML-driven. ML algorithms are implanted in every aspect of AI.”**

Why now?

FIGURE 3-1

Image classification error over time



Geoff Hinton

Computer scientist. “Godfather” of deep learning.



“Take any old problem where you have to predict something and you have a lot of data, and deep learning is probably going to make it work better than the existing techniques.”

<https://www.youtube.com/watch?v=2HMPRXstSvQ>

**PREDICTION:
USING INFORMATION THAT YOU DO
HAVE TO GENERATE INFORMATION
THAT YOU DON'T HAVE**

A policy vacuum



“I’ve talked to Mark about this. His understanding of the subject is limited.”

Elon Musk

“...people who drum up doomsday scenarios... I actually think it is pretty irresponsible.”

Mark Zuckerberg

Without a better understanding of how AI might impact the economy, we cannot design good policy.

On Fri, Jan 6, 2017 at 3:12 PM, Joshua Gans <joshua.gans@gmail.com> wrote:

Dear Paul

Ajay Agrawal, Avi Goldfarb, and I are organizing the NBER's first conference on the Economics of AI: 9am Wednesday September 13 - 12noon Thursday September 14.

The small group will include some of the top economic scholars in the world. The economists will be focused on laying out a research agenda for the economics of AI. The context: imagine back to 1995 when the Internet was about to become a thing. What would have happened to economic research into that revolution had the leading economists gathered to scope out a research agenda at that time?

Today, we are facing the same opportunity with regard to AI. This time around we are organizing a group of leading economists to scope out the research agenda for the next 20 years into the economics of AI. Our intention is to have talks on various aspects of the agenda from historical context to short and long term impacts.

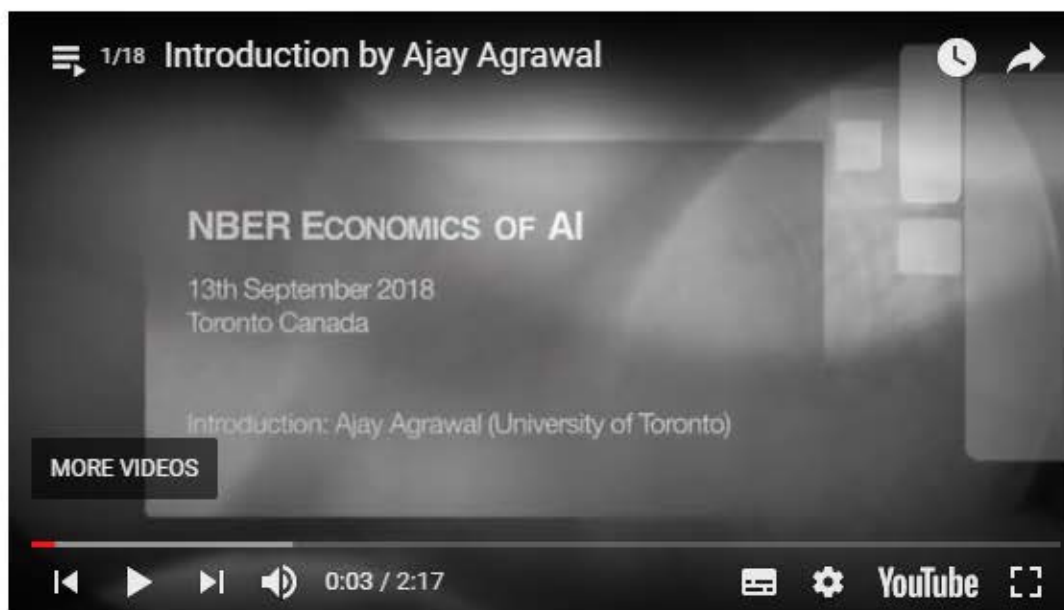
Please let me know whether you are keen to join us.

Best

Joshua

On 13-14 September, 2017, we held our inaugural conference in Toronto to set the research agenda on The Economics of AI. The conference was made possible by generous support from the [NBER](#), [Creative Destruction Lab](#) and [CIFAR](#). Ajay Agrawal, Joshua Gans and Avi Goldfarb of the University of Toronto organised the conference.

Presentation Videos



Forthcoming Conference volume:
Agrawal, Gans, and Goldfarb eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press

Themes

- A.I. as a General Purpose Technology
- Growth, jobs, and inequality
- Regulation
- Impact on the practice of economics

9:10 am	<p>Iain M. Cockburn, Boston University and NBER Rebecca Henderson, Harvard University and NBER Scott Stern, Massachusetts Institute of Technology and NBER <i>The Impact of Artificial Intelligence on Innovation</i></p> <p>Discussant: Matthew Mitchell, University of Toronto</p>		<p>Dinner Panel: Geoffrey Hinton, University of Toronto and Google Yann LeCun, New York University and Facebook Ruslan Salakhutdinov, Carnegie Mellon University and Apple</p>	
9:55 am	<p>Erik Brynjolfsson, Massachusetts Institute of Technology and NBER Daniel Rock, Massachusetts Institute of Technology Chad Syverson, University of Chicago and NBER <i>Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics</i></p> <p>Discussant: Rebecca Henderson, Harvard University and NBER</p>	8:30 am	<p>Colin Camerer, California Institute of Technology <i>Behavioural Economics (slides)</i></p> <p>Discussant: Daniel Kahneman, Princeton University</p>	
10:40 am	Break		9:15 am	<p>Jeffrey D. Sachs, Columbia University and NBER <i>Income Distribution</i></p> <p>Discussant: Susan Dynarski, University of Michigan and NBER</p>
11:00 am	<p>Paul Milgrom, Stanford University Steven Tadelis, University of California at Berkeley and NBER <i>Market Design</i></p> <p>Discussant: Matt Taddy, University of Chicago</p>	10:10 am	Break	
11:45 am	<p>Susan Athey, Stanford University and NBER <i>Impact on Economics (slides)</i></p> <p>Discussant: Mara Lederman, University of Toronto</p>	10:30 am	<p>Philippe Aghion, College de France Benjamin Jones, Northwestern University and NBER Charles I. Jones, Stanford University and NBER <i>Artificial Intelligence and Economic Growth</i></p> <p>Discussant: Patrick Francois, University of British Columbia</p>	
12:30 pm	Lunch - Portman Room, Second Floor Presentation by Vinod Khosla, Khosla Ventures	11:15 am	<p>Joel Mokyr, Northwestern University <i>Historical Context and the Long Run</i></p> <p>Discussant: Manuel Trajtenberg, Tel Aviv University and NBER</p>	
1:45 pm	<p>Ajay K. Agrawal, University of Toronto and NBER Joshua Gans, University of Toronto and NBER Avi Goldfarb, University of Toronto and NBER <i>Prediction, Judgment and Complexity</i></p> <p>Discussant: Andrea Prat, Columbia University</p>	12:00 noon	Lunch - Portman Room, Second Floor	
2:30 pm	Break	1:00 pm	<p>Carl Shapiro, University of California at Berkeley and NBER Hal Varian, University of California at Berkeley <i>Machine Learning, Market Structure and Competition</i></p> <p>Discussant: Judith A. Chevalier, Yale University and NBER</p>	
2:45 pm	<p>Catherine Tucker, Massachusetts Institute of Technology and NBER <i>Privacy</i></p> <p>Discussant: Ginger Zhe Jin, University of Maryland and NBER</p>	1:45 pm	Break	
3:30 pm	<p>Daniel Trefler, University of Toronto and NBER Avi Goldfarb, University of Toronto and NBER <i>Trade</i></p> <p>Discussant: Dave Donaldson, Massachusetts Institute of Technology and NBER</p>	2:20 pm	<p>Joseph E. Stiglitz, Columbia University and NBER Anton Korinek, Johns Hopkins University and NBER <i>Artificial Intelligence, Worker-Replacing Technological Change, and Income Distribution (slides)</i></p> <p>Discussant: Tyler Cowen, George Mason University</p>	
4:15 pm	Break	3:05 pm	<p>David Autor, Massachusetts Institute of Technology and NBER <i>Robocalypse Now: Does Productivity Growth Threaten Employment?</i></p> <p>Discussant: Betsey Stevenson, University of Michigan and NBER</p>	
4:30 pm	<p>Panel: <i>Impact on Policy Making</i> Susan Athey, Stanford University and NBER Austan Goolsbee, University of Chicago and NBER Lawrence H. Summers, Harvard University and NBER</p>	3:50 pm	Wrap up and closing remarks	

General Purpose Technologies (Bresnahan & Trajtenberg 1995)

- GPTs are enabling technologies that open up new opportunities.
- The steam engine, electrification, the internet, etc.
- Long run impact on productivity and society is large.
- The productivity impact of GPTs takes time. Complementary “co-invention” needed (Bresnahan & Greenstein 1996)




Figure 1: Labor Productivity Growth in the Portable Power and IT Eras.
Source: Brynjolfsson, Rock, and Syverson (2018). Figure 8.

A.I. as a General Purpose Technology for Innovation




Figure 2: Publications in Computer Science versus Application Journals, by AI Field
Source: Cockburn, Henderson, and Stern (2018). Figure 4

Could lead to exponential growth in ideas and reverse the trend that scientific ideas are getting harder to find.

Aghion, Jones, and Jones (2018)
Agrawal, McHale, and Oettl (2018)
Cockburn, Henderson, and Stern (2018)

An invention of a method of inventing

- Griliches (1957) highlighted that some new research tools are inventions that constitute a new way of creating new products.
- Hybrid corn represented a widely applicable method for breeding many new varieties. Previously, a primary focus of agricultural innovation was increased specialization of natural varieties through self-fertilization.
- The discovery of double-cross hybridization “was the invention of a method of inventing”, generating a large impact on agricultural productivity.

A.I. as an invention of a method of inventing

- “One of the important insights to be gained from thinking about IMIs, therefore, is that the economic impact of some types of research tools is not limited to their ability to reduce the costs of specific innovation activities—perhaps even more consequentially they enable a new approach to innovation itself, by altering the “playbook” for innovation in the domains where the new tool is applied.”
- A.I. is already widely used in research and innovation across many fields.



Self driving chemical laboratories

Orchestrating artificial intelligence and robotic solutions in “self-driving laboratories” to accelerate autonomous material discovery.

How Artificial Intelligence and Machine Learning Can Impact Market Design

Paul Milgrom, Stanford University
Steven Tadelis,
University of California-Berkeley

Cockburn,
Henderson,
and Stern (2018)

GROWTH, JOBS, AND INEQUALITY

The pessimistic view

1. AI will replace jobs leaving little for humans to do.
2. AI will not be as influential as the technologies that diffused between 1870 and 1970.

“The good news is that those pessimistic predictions cannot both be right. The even better news is that they can both be wrong.”—Joel Mokyr




Assuming A.I. is productivity-enhancing...



“There are two separate questions: There’s an employment question, in which the fundamental question is can we find fulfilling ways to spend our time if robots take our jobs? And there’s an income question, can we find a stable and fair distribution of income.”

Is this the end of jobs?

- No!
- Though if the change is fast, it could lead to an ugly couple of years.
- The long run can be a very long time!



Will inequality get worse?

Yes, if the future mimics the past.





Figure 3: Probability of automation by Occupation Median Hourly Wage. Source: Furman (2018). Figure 4



- Skill-biased technology.

- Ownership of capital.

Addressing inequality and the changing nature of work

- The social safety net: With the familiar tradeoffs.
- Bill Gates called for taxation of robots.
 - Standard argument: less investment, slower productivity growth
 - Stiglitz and Korinek (2018): A combination of finely balanced IP rights and capital taxation can limit distortions and enable distribution.
- Universal basic income
 - Goolsbee (2018), Furman (2018), Furman and Seamans (2018), and others are quite critical of this idea.
- The challenge of finding meaning in leisure.
 - Francois (2018), Stevenson (2018)
- Political economy
 - Arrow's impossibility theorem (Francois 2018)
 - Inequality, mass displacement, and threats to democracy (Trajtenberg 2018).

REGULATION

Antitrust: Will a few companies dominate?



- Economies of scale in data.
- Economies of scope in talent.
 - Endogenous Sunk Costs (Sutton 1998)

Privacy

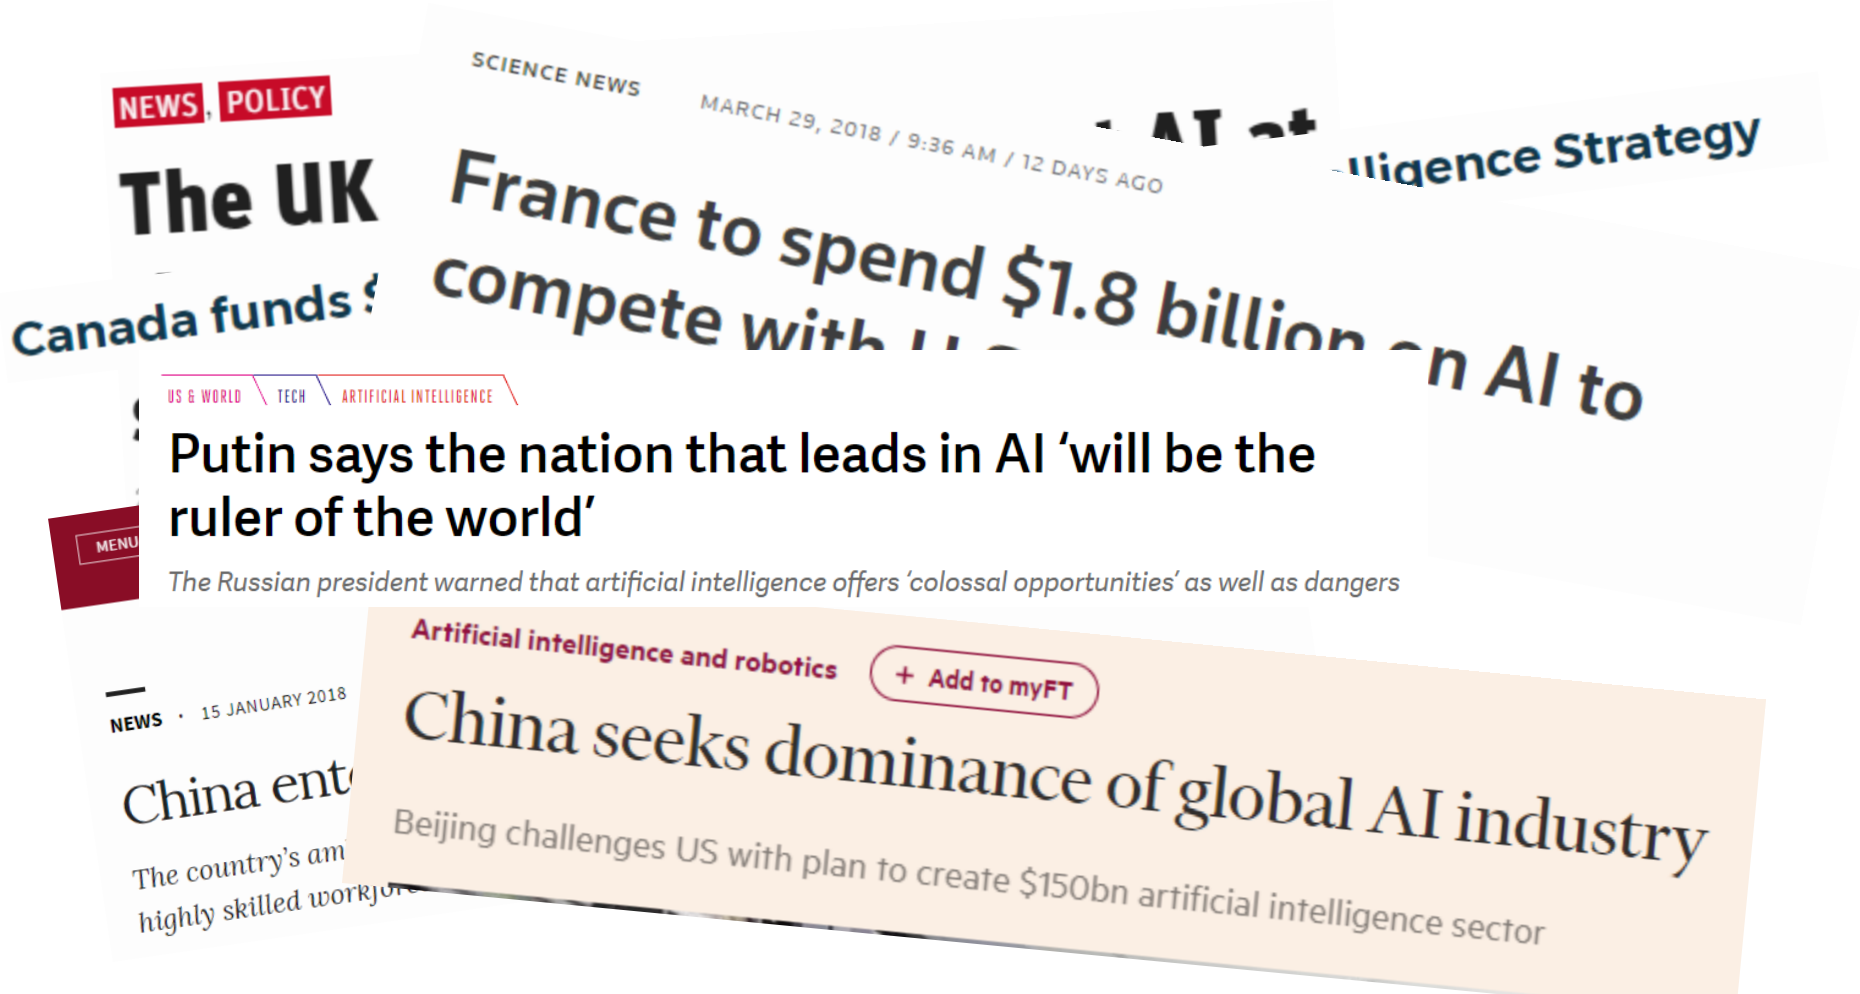
- In the context of A.I., privacy refers to limitations on the collection and use of data by firms and governments.
- Data is a key input into A.I.
- Data is non-rival, meaning that it can be repurposed.
- Data has externalities: Data created by one individual may contain data about others.

Privacy regulation

- Too much privacy regulation means companies can't use data in order to build and implement A.I.
- Too little privacy regulation means consumers will be unwilling to share data with companies.
- The empirical evidence so far is that most attempts at privacy regulation of the digital economy have resulted in less innovation.
- How that plays out given recent scandals and attention is an open question.

Trade

- Many countries view A.I. investments as strategic.



A.I. and trade policy

- Only worth these investments if there are large anticipated rents from leading in A.I.
- Trade agreements currently include provisions for environmental and labor standards to avoid a race-to-the-bottom.
- Could include A.I. industry subsidy provisions, including rules on access to government data.
- On privacy, the E.U. might want U.S. and China to have stricter privacy results so that E.U. companies can succeed despite the E.U.'s stricter take on the right to privacy.

Liability

- The purpose of the tort system is to deter people and companies from injuring others, and to compensate injured parties.
- Tort risk can increase or decrease innovation, depending on whether the risk is driven by new or existing products.
- Need clear liability rules. Those rules need to be strict enough for consumers to trust the technology (and for the technology to be safe!) but not so strict that companies bear too much risk.

Liability and bias

- Algorithms can generate biased results, even by accident.
- In a STEM education ad campaign on social media, ads were shown more to men than women.
- The reason was innocuous: It is more expensive to advertise to women. So the algorithm showed it to men more.
- Should the social media network be found liable? What if human decision makers are more biased?

IMPACT ON THE PRACTICE OF ECONOMICS

Recall: A.I. as a General Purpose Technology for Innovation




Figure 2: Publications in Computer Science versus Application Journals, by AI Field
Source: Cockburn, Henderson, and Stern (2018). Figure 4

Could lead to exponential growth in ideas and reverse the trend that scientific ideas are getting harder to find.

Aghion, Jones, and Jones (2018)
Agrawal, McHale, and Oettl (2018)
Cockburn, Henderson, and Stern (2018)

Econometrics and ML

- Outsourcing of model selection to the machine
- Policy problems related to prediction
- Creation of new variables for economic analysis
- Prediction to verify theory
- ML and causal inference

WHAT'S NEXT?

Economics of Artificial Intelligence

Ajay K. Agrawal, Joshua S. Gans, Avi Goldfarb, and Catherine Tucker, Organizers

September 13-14, 2018

Supported by the Alfred P. Sloan Foundation, CIFAR, and the Creative Destruction Lab

Intercontinental Hotel

220 Bloor Street West

Toronto, ON Canada M5S 1T8

We have questions. We need answers.

- At this point, almost all A.I. research is speculative. We cannot measure the impact of AI and test the theories until it diffuses.
- Measurement is key. Empirical research opportunities driven by what is measured and by creativity in identifying ways to measure the impact.
- We organized this workshop to bring many smart young economists together to start to provide answers.
- The research you do will guide economists, businesses, and (especially) policymakers for decades.

THANK YOU