

The Impact of Artificial Intelligence on Innovation



September 2017

Iain M. Cockburn, BU and NBER

Rebecca Henderson, Harvard and NBER


Scott Stern, MIT and NBER

The Impact of Optical Lenses





Atomwise
Better medicines faster.



Outline

- The Evolution of Artificial Intelligence
- The Economics of Research Tools: Generality and Inventions for the Method of Invention
- Deep Learning as a GPT
- Deep Learning as an Invention for the Method of Invention
- Implications



Symbolic Systems

Capturing the
Logical Flow of
Human
Intelligence
through Symbolic
Logic



Robotics

Performing Key
Human Tasks in
Response to
Sensory Stimuli
(*Elephant Don't
Play Chess*)



Neural Networks & Learning

Reliable and
accurate
predictions of
output in
relation to
complex inputs

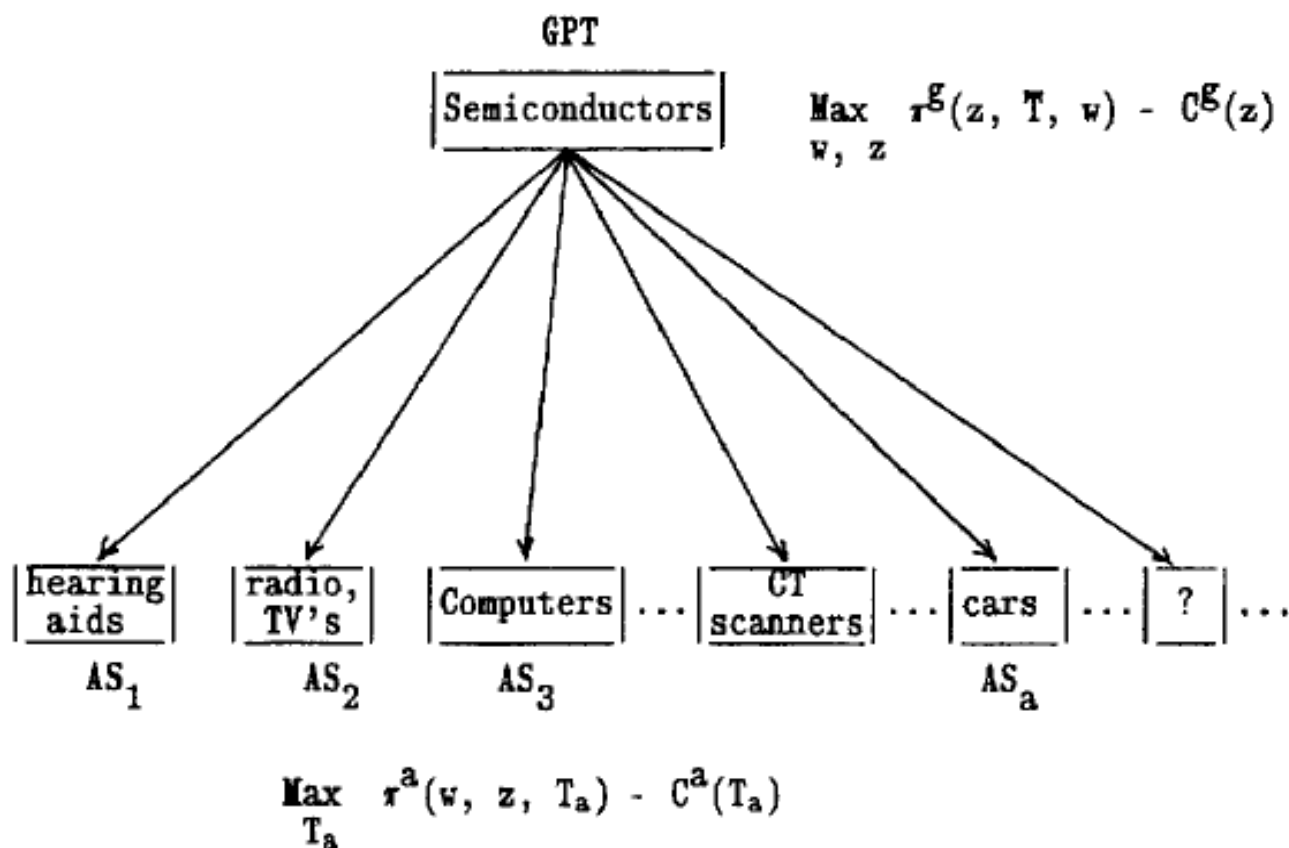
2009....



IMAGENET

How might this purported shift in the science of artificial intelligence influence the rate and direction of innovation?

General Purpose Technologies (Bresnahan and Trajtenber, 1995)

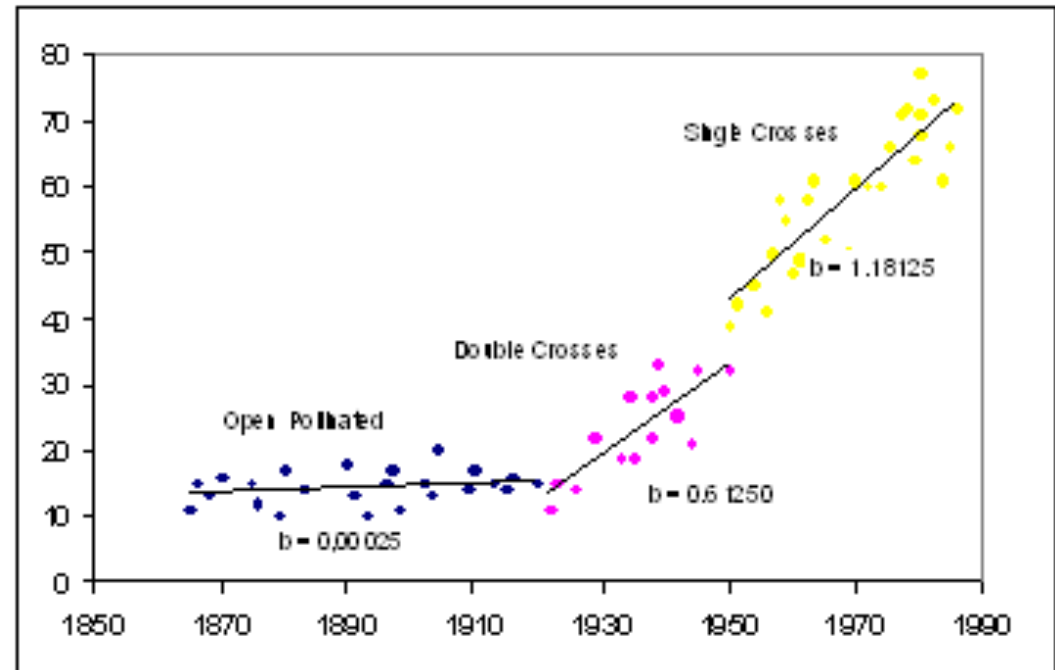




The Economics of GPTs: Underprovision?

- Profound impact
 - Up and down the value chain
 - Across multiple sectors
 - And over time

An Invention for Method of Invention

Griliches 1957



The Interplay between GPTs and IMIs

		<u>General Purpose Technology</u>	
		NO	YES
<u>Invention in the Method of Invention</u>	NO	INDUSTRIAL ROBOTS	PICK AND PLACE ROBOTS
	YES	fMRI ALGORITHM	DEEP LEARNING

Key Hypothesis:

***Relative to other areas of artificial intelligence,
deep learning may represent a new general-
purpose invention for the method of invention***


What evidence exists to suggest that deep learning is actually a GPT, an IMI, or both?

Empirical Approach and Data

- We collected a new dataset of all publications (Web of Science) and patent (USPTO) data from 1990-2015 (2014 for patents).
- To investigate the relative evolution of different aspects of the field, we classify each paper and patent based on detailed keywords into three mutually exclusive areas:
 - Symbol Processing Methods (Symbolic reasoning, pattern analysis)
 - Robotics (robot, sensor networks, etc)
 - Learning (machine learning, neural networks)
- Out of an initial sample of 98124 publications, 91446 are able to be classified uniquely into one of the three fields


AI research over time

- Finding 1: Rapid growth overall in the field of AI, largely driven by Learning Systems




Geography of Learning Research

- Finding 2a: U.S. lower/slower in this field of AI




AI research over geography

- Finding 2b: much lower share of Learning in total AI publications in the U.S. 2000-2010 as compared to ROW




AI research in computer science journals vs. other application sectors.

Publications in Computer Science vs. Application Journals



AI research in computer science journals vs. other application sectors by AI field.


Publications in CS vs. Application Journals, by AI Field



Diffusion of AI into multiple sectors


AI publications in subjects other than computer science

Publications by year (subjects)



Learning systems research in computer science journals vs. other application sectors by geography.

- Finding 5: Upswing in Learning publications since 2009 is driven by international publications (though U.S. researchers are catching up).



Implications for Innovation

- Possible shift from monocausal to predictive reasoning
- Reduction in the costs of “incremental” innovation
 - Likely to impact high marginal cost repetitive R&D search functions such as routinized testing of samples or detailed calculation exercises against a known objective function
- Lower barriers to entry in many scientific fields and to innovation in operations and commercial practice?
- A complement to more fundamental innovation – “exploring unknown unknowns”?

The Management of Innovation

- A substitution away from skilled repetitive labor towards AI capital
 - Likely to impact high marginal cost repetitive R&D search functions such as routinized testing of samples or detailed calculation exercises against a known objective function
 - But, this type of work is an important labor input to R&D (Evenson and Kislev, 1975, etc)
- Potential for lower barriers to entry in scientific fields through reduction in need for expertise at detailed routinized tasks
- But, increasing barriers to entry based on data or AI capital access
 - Limited “market” for AI services
- *A shift in the nature of science from monocausal to predictive reasoning*

Implications for Innovation Policy

- Classical “GPT” concerns: under-provision of innovation up and down the value chain, across sectors and over time
- Potential for significant gap between the private and social returns to transparency and data sharing
- Potential for deepening of replicability crisis, as well as balkanization of scientific knowledge
- Proactive policies encouraging transparency and data sharing among communities likely to yield higher innovation productivity

Implications for Competition Policy: Data

- Potential for increasing returns to data as ever-larger or more granular datasets allow for significant performance advantage at prediction of disparate application-specific phenomena
 - See Google versus Bing versus Field
- This prospect may result in racing, with the potential for duplication and near-term overinvestment
 - And private sector incentives for data exclusivity to lock in advantage
- Implications for incumbent/entrant dynamics
- Policy question
 - Should the data be an essential facility (Shapiro and Varian)?
 - Who should “own” data on private social behavior (Miller and Tucker)?

Concluding Thoughts

- ***A Tentative Hypothesis: Not only a general-purpose technology likely to diffuse across the economy, but also an invention in the method of invention.***
- If true (a BIG if), the potential impact of deep learning may be as or more important in generating research and innovation as in transforming existing practice
- If so, a potential mechanism for keeping humans “busy” (in an actually productive way) as an increasing number of routinized tasks are automated through AI capital.

Thank you!

Iain M. Cockburn, BU and NBER

Rebecca Henderson, Harvard and NBER

Scott Stern, MIT and NBER