RABINDRA BHARATI JOURNAL OF PHILOSOPHY
ISSN : 0973-0087

AI- A NEW METHOD OF INNOVATION

Dr. Ambar Beharay Associate Professor Department of Management
Dr. Pranati Tilak Dean Department of Management

Abstract:
Artificial intelligence appears to be the focus of this decade. Without a question, AI plays a significant role in the current economy around the world. However, pursuing innovation or research within a business requires a fresh approach, and Artificial Intelligence can undoubtedly help. Application-oriented learning research has grown in popularity since 2009. When we refer to automation-oriented applications like robotics, the potential for current advances in "deep learning" as a general-purpose method of invention may be substituted. This can be described as a paradigm shift away from labor-intensive, systematic research and toward research that incorporates passively generated huge datasets and improved prediction algorithms. It will not only assist organisations in mastering this form of study, but it will also provide potential commercial advantages. This strategy can assist in the acquisition and control of big datasets and application-specific algorithms. We believe that organisations should adopt rules that foster transparency and sharing of essential datasets across public and private players, since they will be critical instruments for boosting research productivity and innovation-driven competition in the future.

Introduction:

Rapid advancements in artificial intelligence have far-reaching ramifications for the business and society as a whole. These advancements have the potential to have a direct impact on the production and qualities of a wide range of goods and services, with significant implications for productivity, employment, and competition. But, as significant as these benefits are expected to be, artificial intelligence also has the ability to alter the innovation process itself, with equally significant ramifications that may eventually outweigh the direct influence.

Consider the case of Atomwise, a startup that is developing revolutionary technology for identifying possible medication candidates (and insecticides) by predicting the bioactivity of candidate molecules using neural networks. Deep convolutional neural networks, according to the business, "far outperform" traditional "docking" algorithms. The company's AtomNet software is characterised as being able to "recognise" foundational building blocks of organic chemistry after sufficient training on huge amounts of data, and is capable of giving very accurate predictions of the outcomes of real-world physical tests (Wallach et al., 2015). Such innovations carry the promise of a significant increase in the efficiency of early-stage drug screening.

Of fact, Atomwise's technology (and those of other businesses employing artificial intelligence to assist drug discovery or medical diagnosis) is still in its infancy: despite encouraging preliminary results, no new medications have yet been released using these new approaches. Whether Atomwise's technology delivers on its promise or not, it is representative of the ongoing effort to develop a new innovation "playbook," one that uses large datasets and machine learning algorithms to engage in precise prediction of biological phenomena in order to guide the design of effective interventions. Atomwise, for example, is now using this approach to discover and develop new insecticides and agents for crop disease control.

While some AI applications will undoubtedly provide lower-cost or higher-quality inputs into many existing production processes (raising concerns about the potential for large job displacements), others, such as deep learning, hold the promise of not only increased productivity across a wide range of industries but also changes in the nature of the innovation process within those industries. The "creation

of a process of invention," as notably described by Griliches (1957), has the potential to have a significantly larger economic impact than the production of any single new product by enabling creativity across multiple applications.
We suggest that recent breakthroughs in machine learning and neural networks, by improving both the performance of end-use technologies and the nature of the data, have improved the nature of the data.

Recent advancements in machine learning and neural networks, we suggest, are likely to have a particularly big influence on innovation and growth because of their ability to improve both the performance of end-use technologies and the structure of the invention process. As a result, the incentives and roadblocks that may influence the development and diffusion of these technologies are a hot topic in economic research, and gaining a better understanding of the conditions under which different potential innovators can gain access to these tools and use them in a pro-competitive manner is a major policy concern.

We also start to look into some of our analysis' organisational, institutional, and policy implications. We consider machine learning to be the "creation of a method of invention," with each application requiring access to massive, granular datasets on physical and social behaviour in addition to the underlying algorithms. Even if the underlying scientific methodologies (i.e., the basic multi-layered neural networks algorithms) are open, terms of access to complementary data are expected to have a substantial impact on chances for future progress in this field—and commercial uses thereof.
If there are increasing returns to scale or scope in data acquisition (i.e., there is more learning to be had from a "larger" dataset), it is possible that early or aggressive entrants into a particular application area may be able to create a substantial and long-lasting competitive advantage over potential rivals simply by controlling data rather than formal intellectual property or demand-side network effects. Strong incentives to keep data private have the potential drawback of preventing data from being shared among academics, limiting everyone's capacity to access an even larger set of data that would result from public aggregation. As incumbents' competitive advantage grows, new entrants' ability to push technical change grows. Though this is a substantial potential, it is also true that, at least so far, most main application sectors appear to have a significant degree of entrance and experimentation.

Machine learning and neural networks look to offer a lot of promise as a research tool for categorization and prediction challenges. These are also significant limiting issues in a number of research activities, and the use of "learning" techniques to AI, as demonstrated by the Atomwise example, holds the promise of drastically cheaper costs and enhanced performance in R&D projects where these are big challenges.

AI-based learning, like hybrid corn, may be better viewed as an IMI rather than a narrowly confined solution to a specific problem. On the one hand, AI-based learning may be able to "automate" much of the "finding" process in a variety of disciplines where categorization and prediction tasks are crucial. On the other hand, they may "extend the playbook" in the sense of broadening the range of problems that may be addressed and dramatically altering the conceptual methods and framing of problems in scientific and technological communities. Optical lenses were invented in the 17th century and had a significant direct economic influence on applications like spectacles. However, optical lenses such as microscopes and telescopes had massive and long-lasting indirect effects.

Machine learning, for example, is described by Leung et al. (2016) as a technique to "learn to read the genome" in ways that human cognition and perception cannot.
Many research tools, of course, are neither IMIs nor GPTs, and their primary purpose is to lower the cost or improve the quality of an existing innovation process. New materials, for example, have the potential to improve the efficiency of various research procedures in the pharmaceutical business. Advances in AI

present a difficulty in that they appear to be research tools that have the ability to change not just the process of invention, but also have ramifications in a wide range of sectors.

Another key property of research instruments from a policy standpoint is that it may be particularly difficult to appropriate their benefits. Providing appropriate incentives for an upstream innovator who develops only the first "stage" of an innovation (such as a research tool) can be especially difficult when contracting is imperfect and the ultimate application of the new products whose development is enabled by the upstream innovation is uncertain, as Scotchmer (1990) points out. When the ultimate innovation that creates value requires multiple steps, Scotchmer and her co-authors emphasised a key point about a multi-stage research process: providing appropriate innovation incentives is not only a question of whether and how to provide property rights in general, but also of how best to distribute property rights.

The Evolution of Artificial Intelligence: Robotics, Symbolic Systems, and NeuralNetworks

Nilsson (2010) defines AI as "that effort committed to making machines intelligent, and intelligence is that attribute that enables an entity to function effectively and with foresight in its environment" in his comprehensive historical history of AI research. His account covers a wide range of disciplines, including biology, linguistics, psychology, and cognitive sciences, neuroscience, mathematics, philosophy, and logic, engineering, and computer science, to name a few. And, regardless of their individual approaches, artificial intelligence research has been linked since its inception by its engagement with Turing (1950) and his exploration of the prospect of automating intelligence.

A second influential AI trajectory has been in the field of robotics in general. While the concept of "robots" as machines that can perform human tasks has been around since the 1940s, the field of robotics began to take off in the 1980s as a result of advances in numerically controlled machine tools and the development of more adaptive but still rules-based robotics that rely on active sensing of a known environment. The largest-scale deployment of "industrial robots" in manufacturing applications has been perhaps the most economically impactful application of AI to date.

These advancements are significant, and when the term AI is mentioned, the most advanced robots continue to captivate the public imagination. In general, however, robotics advancements are not IMIs. Although rising laboratory automation boosts research efficiency, robotics breakthroughs are not (yet) intimately linked to the fundamental manner in which researchers themselves could build methodologies to conduct innovation across various disciplines. Of course, there are counterexamples to this assertion: robotic space probes have been a crucial research tool in planetary science, and the capacity of automated remote sensing devices to collect data at very large scales or in difficult situations may alter several fields of research. However, robots are still mostly used in specific end-use "production" applications.

Finally, a "learning" approach can be broadly described as a third stream of research that has been a major feature of AI since its inception. The learning approach, rather than focusing on symbolic logic or exact sense-and-react systems, aims to develop dependable and accurate methods for predicting specific outcomes (physical or logical) in the presence of specific inputs. In this field, the concept of a neural network has been extremely essential. A neural network is a software that converts a collection of inputs into a set of outputs using a mix of weights and thresholds, measures the "closeness" of the outputs to reality, and then adjusts the weights to shrink the gap between the outputs and reality.

As more inputs are fed into neural networks, they can learn (Rosenblatt, 1958; 1963). Hinton and his co-authors improved the conceptual foundation on which neural networks are founded in the 1980s by developing "back-propagating multi-layer" approaches that further improve their capacity for supervised learning.

Deep Learning as a General-Purpose Invention in the Method of Invention:Considerations for Organizations, Institutions and Policy

With these findings in mind, we can now think about the consequences for innovation and policy whether deep learning is a general-purpose technology (GPT) and/or a general-purpose invention in the process of invention (IMI). If deep learning is simply a GPT, it will certainly create innovation in a variety of applications (with potential spillovers both back to the learning GPT and to other application sectors), but it will not change the structure of the innovation production function. If it's also a general-purpose IMI, we may anticipate it to have a considerably greater impact on the economy's overall innovation, growth, and productivity as the dynamics unfold—and to cause far more severe short-term disruptions in labour markets and the internal structure of businesses.

Deep learning's widespread use as a research tool signals a trend toward investigative methodologies that employ massive datasets to provide predictions for physical and logical phenomena that have hitherto eluded systematic empirical inspection. Earlier knowledge (like in IBM's Watson's "learning" of prior literatures), online transactions (e.g., search or online purchase activity), and physical events are possible sources of these data (e.g., the output from various types of sensors or geolocation data) What does this mean for the proper structure of innovation, the institutions we have for teaching and performing long-term research, and legislation, particularly as we consider private incentives to retain proprietary datasets and application-specific algorithms?

The Management and Organization of Innovation

Perhaps most quickly, the advent of general-purpose predictive analytics based on massive datasets appears to be leading to a shift in the research production process from labour to capital. Many sorts of R&D and, more broadly, innovation are effectively labor-intensive search issues with high marginal cost per search (Evenson and Kislev, 1975, among others). Deep learning's progress promises dramatically lower marginal search costs, causing R&D businesses to shift away from highly skilled workers and toward fixed-cost AI expenditures. Deep learning's progress promises dramatically lower marginal search costs, causing R&D businesses to shift away from highly skilled workers and toward fixed-cost AI expenditures. These investments are anticipated to increase performance in existing "search demanding" research initiatives, as well as open up new avenues for investigating social and physical phenomena previously thought to be intractable or even beyond the scope of systematic scientific and empirical inquiry.

Deep learning's arrival has major ramifications for the patent system. Though there has been relatively little patenting of deep learning innovations to date, historical episodes such as the discovery and attempted wholesale patenting of express sequence tags and other types of genetic data suggest that breakthroughs in research tools—often combined with a lack of capacity at patent offices and conflicting court decisions—can result in long periods of uncertainty, which has hampered the issuing of new patents and, as a result, has resulted in lowe patentees.

Deep learning also raises tough legal doctrine concerns for patent systems that are based on the concept of creative writers and inventors. In patent law, for example, the term "inventorship" has a highly particular meaning, with significant implications for ownership and control of the claimed invention. Is it possible for an AI system to be an inventor in the sense that the US Constitution's drafters intended?

Similarly, the size of the inventive step required to acquire a patent is determined by considering whether the claimed invention would or would not be evident to a "person of ordinary competence in the art." Who this "person" is, and what defines "ordinary skill" in an era of deep learning algorithms trained on proprietary data, are concerns that are much beyond the scope of this essay.

In addition to these classic innovation policy concerns, the advent of deep learning presents a slew of new ones, including privacy concerns, the risk of prejudice (deep learning has been shown to reinforce stereotypes already prevalent in society), and consumer protection concerns (related to areas such as search, advertising, and consumer targeting and monitoring). The idea is that, to the extent that deep learning is general-purpose, the difficulties that arise in each of these domains (and more) will be played out across a wide range of industries and situations, and on a global rather than local scale.

Little research has been done to aid in the creation of institutions that will be responsive at the application sector level, as well as absorb the potential challenges that may occur as a result of deep learning's anticipated status as a GPT.
A crucial concern in the future will be ensuring that deep learning does not create monopolisation and entry barriers across a variety of industries.

Concluding Thoughts:
This exploratory essay's objective is not limited to presenting a systematic analysis or prediction of AI's anticipated impact on innovation. It does not appear to be a guide to innovation policymaking. Instead, we argue that deep learning can be used to create new things.
Our preliminary study identifies a few significant concepts that have received little attention in the economics and policy debates thus far. First, it's important to distinguish between significant and important advances in fields like robotics and the potential of a general-purpose method of invention based on the application of multi-layered neural networks to large amounts of digital data to be a "invention in the method of invention," at least from the standpoint of innovation. This idea is supported by both existing qualitative data and our preliminary empirical analysis, which show a significant movement toward deep learning-based application-oriented research since 2009.
Second, the idea of a shift in the innovation process presents important policy and management challenges, ranging from how to evaluate this new sort of science to the potential for prediction tools to create new barriers to entry in a variety of industries. Future study should focus on proactive examination of the proper commercial and public policy responses to these discoveries.

REFERENCES

Aghion, P. and P. Howitt (1992) “A Model of Growth Through Creative Destruction,” *Econometrica*, 60(2), 323-251.

Bresnahan, T., E. Brynjolfsson, and L. Hitt (2002) “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence,” *The Quarterly Journal of Economics*, 117(1), 339-376.

Bresnahan, T. and S. Greenstein (1999) “Technological Competition and the Structure of the Computer Industry,” *Journal of Industrial Economics*, 47(1), 1-40.

Bresnahan, T. and M. Trajtenberg (1995) “General Purpose Technologies ‘Engines of Growth’?” *Journalof Econometrics*, 65 (1995) 83-108.

Brooks, R. (1990) “Elephants Don’t Play Chess,” *Robotics and Autonomous Systems*, 6, 3-15.Brooks, R. (1991) “Intelligence Without Representation,” *Artificial Intelligence*, 47, 139-159.

Brynjolfsson, E. and K. McElheran (2017) “The Rapid Adoption of Data-Driven Decision-Making,” *American Economic Review*, 106(5), 133-139

Griliches, Z. (1957) “Hybrid Corn: An Exploration in the Economics of Technological Change,” *Econometrica*, 25(4), 501-522.

Henderson, R. and K. Clark (1990) “Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms,” *Administrative Science Quarterly*, 35(1), 9-30.

Krizhevsky, A., I. Sutskever, G. Hinton (2012) “ImageNet Classification with Deep Convolutional Neural

Networks," *Advances in Neural Information Processing*, 25, *MIT Press*.
Leung, M.K.K., A. Delong, B. Alipanahi, and B.J. Frey (2016) "Machine Learning in Genomic Medicine: A Review of Computational Problems and Data Sets," *Proceedings of the IEEE*, 104(1): 176-197.
Marco, A., A. Myers, S. Graham, P. D'Agostino, and K. Apple (2015) "The USPTO Patent Assignment Dataset: Descriptions and Analysis," *USPTO Working Paper No. 2015-02*, 1-53.
Marco, A., M. Carley, S. Jackson and A. Myers (2015) "The USPTO Historical Patent Data Files," *USPTO Working Paper No. 2015-01*, 1-57.
Minsky, M. (1961) "Steps Toward Artificial Intelligence," *Proceedings of the IRE*, 8-30.Mokyr, J (2002) *Gifts of Athena*, Princeton University Press.
Nilsson, N. (2010), *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge University Press.
Romer, P. (1990) "Endogenous Technological Change," *Journal of Political Economy*, 98(5), S71-S102.
Rosenberg, N. and M. Trajtenberg (2004) "A General Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States," *Journal of Economic History*, 61(1), 61-99.
Rumelhart, D., G. Hinton, and R. Williams (1986) "Learning Internal Representations by Error Propagation," in J. McClelland and D. Rumelhart (editors), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 2: Psychological and Biological Models*, MIT Press, 7-57.
Scotchmer, S. (1991) "Standing on the Shoulders of Giants: Cumulative Research and the Patent Law," *Journal of Economic Perspectives*, 5(1), 29-41.
Turing, A. (1950) "Computing Machinery and Intelligence," *Mind*, 59, 433-460.
Wallach, I. Dzamba, M. and Heifels, A. "AtomNet: A Deep Convolutional Neural Network for Bioactivity Prediction in Structure-based Drug Discovery." arXiv:1510.02855 [cs.LG]
Williams, H. .(2013) "Intellectual Property Rights and Innovation: Evidence from the Human Genome,", *Journal of Political Economy*, 121(1): 1-27

Table 1A: Publication Data Summary Statistics

	Mean	Std. Dev.	Min	Max
Publication Year	2007	6.15	1990	2015
Symbolic Systems	.12	.33	0	1
Learning Systems	.61	.48	0	1
Robotics	.21	.41	0	1
Artificial Intelligence	.06	.23	0	1
Computer Science	.44	.50	0	1
Other Applications	.56	.50	0	1
US Domestic	.25	.43	0	1
International	.75	.43	0	1
Observations	95840			

Table 1B: Patent Data Summary Statistics

	Mean	Std. Dev.	Min	Max
Application Year	2003	6.68	1982	2014
Patent Year	2007	6.98	1990	2014
Symbolic Systems	.29	.45	0	1
Learning Systems	.28	.45	0	1

Robotics	.41	.49	0	1
Artificial Intelligence	.04	.19	0	1
Computer Science	.77	.42	0	1
Other Applications	.23	.42	0	1
US Domestic Firms	.59	.49	0	1
International Firms	.41	.49	0	1
Org Type Academic	.07	.26	0	1
Org Type Private	.91	.29	0	1
Observations	13615			

Table 2A: Distribution of Publications across Subjects

	Mean	Std. Dev.
Biology	.034	.18
Economics	.028	.16
Physics	.034	.18
Medicine	.032	.18
Chemistry	.038	.19
Mathematics	.042	.20
Materials Science	.029	.17
Neurology	.038	.19
Energy	.015	.12
Radiology	.015	.12
Telecommunications	.055	.23
Computer Science	.44	.50
Observations	95840	

Table 2B: Distribution of Patents across Application Sectors

	Mean	Std. Dev.
Chemicals	.007	.08
Communications	.044	.20
Computer Hardware and Software	.710	.45
Computer Peripherals	.004	.06
Data and Storage	.008	.09
Business software	.007	.09
All Computer Science	.773	.42
Medical	.020	.14
Electronics	.073	.26
Automotive	.023	.15
Mechanical	.075	.26
Other	.029	.16
Observations	13615	

Table 3: Publications Across Sectors, by AI Field, 2004-2006 versus 2013-2015

		Biology	***Economics***	***Physics***	***Medicine***	***Chemistry***	***Math***	***Materials***	***Neuro.***	***Energy***	***Radiology***	***Telecom.***	***CompSci***
Learning Systems	2004-2006	258	292	343	231	325	417	209	271	172	94	291	3889
	2013-2015	600	423	388	516	490	414	429	970	272	186	404	4582
	% growth	133%	45%	13%	123%	51%	-1%	105%	258%	58%	98%	39%	18%
Robotics	2004-2006	33	10	52	69	24	45	36	31	6	47	653	1431
	2013-2015	65	12	122	83	92	80	225	139	18	25	401	1322
	% growth	97%	20%	135%	20%	283%	78%	525%	348%	200%	-47%	-39%	-8%
Symbol Systems	2004-2006	93	8	68	96	139	54	32	35	15	82	51	827
	2013-2015	105	10	125	84	149	60	101	73	22	56	88	1125
	% growth	13%	25%	84%	-13%	7%	11%	216%	109%	47%	-32%	73%	36%

Table 4: Herfindahl-Hirschman Index for Application Sectors

Application	**H=** $\sum \boldsymbol{PatShare^2}$
Chemical Applications	153.09
Communications	140.87
Hardware and Software	86.99
Computer Science Peripherals	296
Data and Storage	366.71
Computer Science Business Models	222
Medical Applications	290.51
Electronic Applications	114.64
Automotive Applications	197.03
Mechanical Applications	77.51
Other	129.20

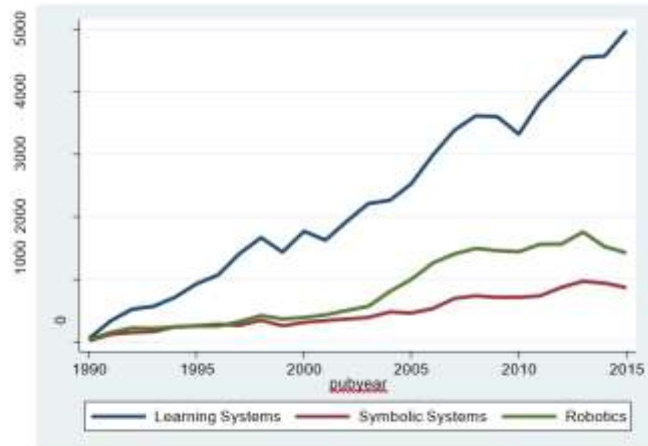

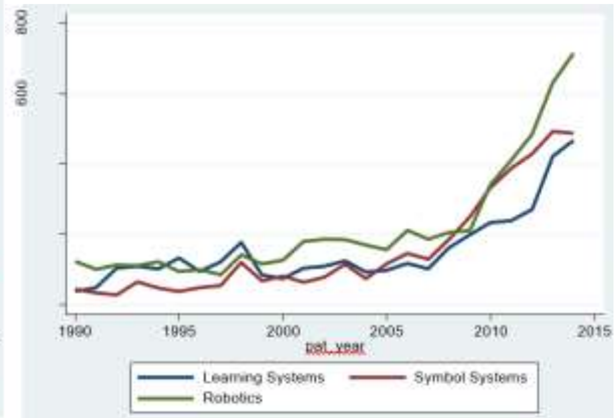

Figure 1A: Publications by AI field over Time Figure 1B: Patents by AI field over Time

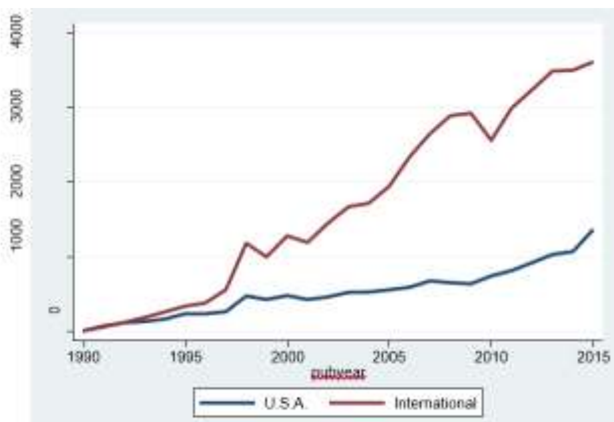

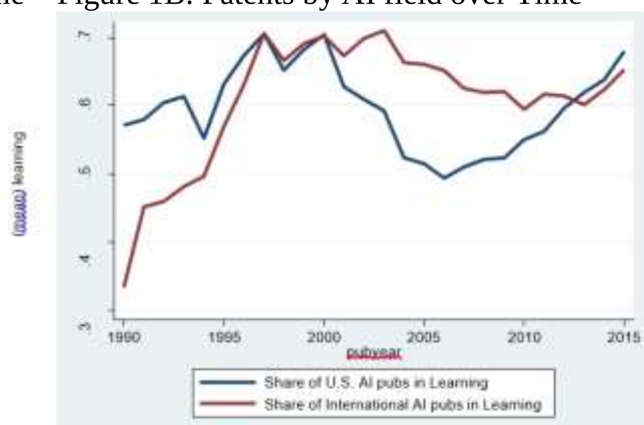

Figure 2A: Academic Institution Publication Fraction by AI Field
Figure 2B: Fraction of Learning Publications by US versus World

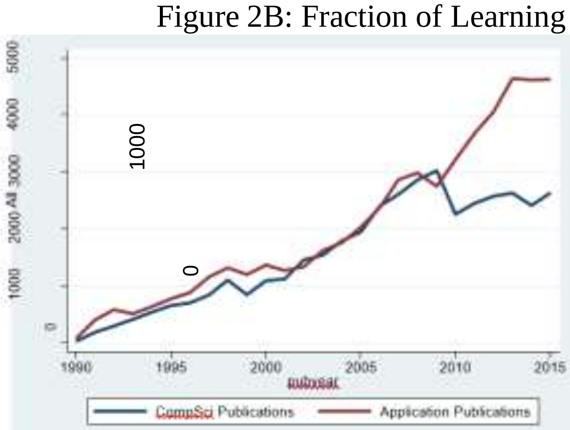

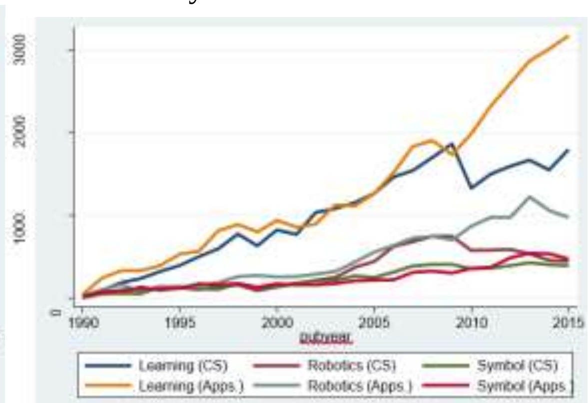

Figure 3: Publications in Computer Science versus Application Journals
Figure 4: Publications in Computer Science versus Application Journals, by AI Field

Figure 5: Learning Publications in Computer Science versus Applications.

Appendix A

Appendix Table 1: Artificial Intelligence Keyword Allocation

Symbols	Learning	Robotics
natural language processing	machine learning	computer vision
image grammars	neural networks	robot
pattern recognition	reinforcement learning	robots
image matching	logic theorist	robot systems
symbolic reasoning	bayesian belief networks	robotics
symbolic error analysis	unsupervised learning	robotic
pattern analysis	deep learning	collaborative systems
symbol processing	knowledge representation and reasoning	humanoid robotics
physical symbol system	crowdsourcing and human computation	sensor network
natural languages	neuromorphic computing	sensor networks
pattern analysis	decision making	sensor data fusion
image alignment	machine intelligence	systems and control theory
optimal search	neural network	layered control systems
symbolic reasoning		
symbolic error analysis		