artificial intelligence index

2019 annual report

Stanford
Human-Centered
Artificial Intelligence
HAI

Steering Committee

Raymond Perrault (report coordinator)
SRI International

Yoav Shoham (chair)
Stanford University

Erik Brynjolfsson
MIT

Jack Clark
OpenAI

John Etchemendy
Stanford University

Barbara Grosz
Harvard University

Terah Lyons
Partnership On AI

James Manyika
McKinsey Global Institute

Juan Carlos Niebles
Stanford University

Project Manager and Report Editor-in-Chief

Saurabh Mishra
Stanford University

Table of Contents

How to cite this Report:

Raymond Perrault, Yoav Shoham, Erik Brynjolfsson, Jack Clark, John Etchemendy, Barbara Grosz, Terah Lyons, James Manyika, Saurabh Mishra, and Juan Carlos Niebles, "The AI Index 2019 Annual Report", AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA, December 2019.

The AI Index is as an independent initiative at Stanford University's Human-Centered Artificial Intelligence Institute (HAI).

Stanford
Human-Centered
Artificial Intelligence

The AI Index was conceived within the One Hundred Year Study on AI (AI100).

We thank your supporting partners

We welcome feedback and new ideas for next year. Contact us at AI-Index-Report@stanford.edu.

Introduction to the AI Index 2019 Report

The AI Index Report tracks, collates, distills, and visualizes data relating to artificial intelligence. Its mission is to provide unbiased, rigorously-vetted data for policymakers, researchers, executives, journalists, and the general public to develop intuitions about the complex field of AI. Expanding annually, the Report endeavors to include data on AI development from communities around the globe.

Before diving into the data, it is worth noting the following about the 2019 edition of the AI Index Report:

1. This edition tracks three times as many data sets as the 2018 edition. It includes an update of previous measures, as well as numerous new ones, across all aspects of AI: technical performance, the economy, societal issues, and more.

2. This volume of data is challenging to navigate. To help, we've produced a tool that provides a high-level global perspective on the data. The **Global AI Vibrancy Tool** (vibrancy.aiindex.org) compares countries' global activities, including both a cross-country perspective, as well as a country-specific drill down. Though it is tempting to provide a single ranking of countries, such comparisons are notoriously tricky. Instead, we've provided a tool for the reader to set the parameters and obtain the perspective they find most relevant when comparing countries. This tool helps dispel the common impression that AI development is largely a tussle between the US and China. Reality is much more nuanced. Our data shows that local centers of AI excellence are emerging across the globe. For example, Finland excels in AI education, India demonstrates great AI skill penetration, Singapore has well-organized government support for AI, and Israel shows a lot of private investment in AI startups per capita.

3. We are also releasing the **AI Index arXiv Monitor** (arxiv.aiindex.org), a tool to support research on current technological progress in AI via full-text searches of papers published on the pre-print repository.

Given that measurement and evaluation in complex domains remain fraught with subtleties, the AI Index has worked hard to avoid bias and seek input from many communities. As part of this effort, on October 30, 2019, the Stanford HAI-AI Index Workshop: **Measurement in AI Policy: Opportunities and Challenges (https://hai.stanford.edu/ai-index/workshops)** convened over 150 industry and academic experts from a variety of disciplines related to AI to discuss the many pressing issues that arise from data measurement of AI. The Workshop Proceedings will be available shortly here.

AI Index 2019 Report Highlights

Each of the nine chapters presents well-vetted data on important dimensions related to the activity in, and technical progress of artificial intelligence. Here is a sample of the findings.

1. Research and Development

- Between 1998 and 2018, the volume of peer-reviewed AI papers has grown by more than 300%, accounting for 3% of peer-reviewed journal publications and 9% of published conference papers.
- China now publishes as many AI journal and conference papers per year as Europe, having passed the US in 2006. The Field-Weighted Citation Impact of US publications is still about 50% higher than China's.
- Singapore, Switzerland, Australia, Israel, Netherlands, and Luxembourg have relatively high numbers of Deep Learning papers published on arXiv in per capita terms.
- Over 32% of world AI journal citations are attributed to East Asia. Over 40% of world AI conference paper citations are attributed to North America.
- North America accounts for over 60% of global AI patent citation activity between 2014-18.
- Many Western European countries, especially the Netherlands and Denmark, as well as Argentina, Canada, and Iran show relatively high presence of women in AI research.

2. Conferences

- Attendance at AI conferences continues to increase significantly. In 2019, the largest, NeurIPS, expects 13,500 attendees, up 41% over 2018 and over 800% relative to 2012. Even conferences such as AAAI and CVPR are seeing annual attendance growth around 30%.
- The WiML workshop has eight times more participants than it had in 2014 and AI4ALL has 20 times more alumni than it had in 2015. These increases reflect a continued effort to include women and underrepresented groups in the AI field.

3. Technical Performance

- In a year and a half, the time required to train a large image classification system on cloud infrastructure has fallen from about three hours in October 2017 to about 88 seconds in July, 2019. During the same period, the cost to train such a system has fallen similarly.
- Progress on some broad sets of natural-language processing classification tasks, as captured in the SuperGLUE and SQuAD2.0 benchmarks, has been remarkably rapid; performance is still lower on some NLP tasks requiring reasoning, such as the AI2 Reasoning Challenge, or human-level concept learning task, such as the Omniglot Challenge.
- Prior to 2012, AI results closely tracked Moore's Law, with compute doubling every two years. Post-2012, compute has been doubling every 3.4 months.

4. Economy

- Singapore, Brazil, Australia, Canada and India experienced the fastest growth in AI hiring from 2015 to 2019.

AI Index 2019 Report Highlights

- In the US, the share of jobs in AI-related topics increased from 0.26% of total jobs posted in 2010 to 1.32% in October 2019, with the highest share in Machine Learning (0.51% of total jobs). AI labor demand is growing especially in high-tech services and the manufacturing sector.

- The state of Washington has the highest relative AI labor demand. Almost 1.4% of total jobs posted are AI jobs. California has 1.3%, Massachusetts 1.3%, New York 1.2%, the District of Columbia (DC) 1.1%, and Virginia has 1% online jobs posted in AI.

- In the US, the share of AI jobs grew from 0.3% in 2012 to 0.8% of total jobs posted in 2019. AI labor demand is growing especially in high-tech services and the manufacturing sector.

- In 2019, global private AI investment was over $70B, with AI-related startup investments over $37B, M&A $34B, IPOs $5B, and Minority Stake valued around $2B.

- Globally, investment in AI startups continues its steady ascent. From a total of $1.3B raised in 2010 to over $40.4B in 2018 (with $37.4B in 2019 as of November 4th), funding has increased at an average annual growth rate of over 48%.

- Autonomous Vehicles (AVs) received the largest share of global investment over the last year with $7.7B (9.9% of the total), followed by Drug, Cancer and Therapy ($4.7B, 6.1%), Facial Recognition ($4.7B, 6.0%), Video Content ($3.6B, 4.5%), and Fraud Detection and Finance ($3.1B, 3.9%).

- 58% of large companies surveyed report adopting AI in at least one function or business unit in 2019, up from 47% in 2018.

- Only 19% of large companies surveyed say their organizations are taking steps to mitigate risks associated with explainability of their algorithms, and 13% are mitigating risks to equity and fairness, such as algorithmic bias and discrimination

5. Education

- Enrollment continues to grow rapidly in AI and related subjects, both at traditional universities in the US and internationally, and in online offerings.

- At the graduate level, AI has rapidly become the most popular specialization among computer science PhD students in North America, with over twice as many students as the second most popular specialization (security/information assurance). In 2018, over 21% of graduating Computer Science PhDs specialize in Artificial Intelligence/Machine Learning.

- In the US and Canada, the number of international PhD students graduating in AI continues to grow, and currently exceeds 60% of the PhDs produced from these programs (up from less than 40% in 2010).

- Industry has become, by far, the largest consumer of AI talent. In 2018, over 60% of AI PhD graduates went to industry, up from 20% in 2004. In 2018, over twice as many AI PhD graduates went to industry as took academic jobs in the US.

- In the US, AI faculty leaving academia for industry continues to accelerate, with over 40 departures in 2018, up from 15 in 2012 and none in 2004.

- Diversifying AI faculty along gender lines has not shown great progress, with women comprising less than 20% of the new faculty hires in 2018. Similarly, the share of female AI PhD recipients has remained virtually constant at 20% since 2010 in the US.[1]

[1] Studies on participation of under-represented minorities coming in 2020

AI Index 2019 Report Highlights

6. Autonomous Systems

- The total number of miles driven and total number of companies testing autonomous vehicles (AVs) in California has grown over seven-fold between 2015-2018. In 2018, the State of California licensed testing for over 50 companies and more than 500 AVs, which drove over 2 million miles.

7. Public Perception

- Global central bank communications demonstrate a keen interest in AI, especially from the Bank of England, Bank of Japan, and the Federal Reserve.

- There is a significant increase in AI related legislation in congressional records, committee reports, and legislative transcripts around the world.

8. Societal Considerations

- Fairness, Interpretability and Explainability are identified as the most frequently mentioned ethical challenges across 59 Ethical AI principle documents.

- In over 3600 global news articles on ethics and AI identified between mid-2018 and mid-2019, the dominant topics are framework and guidelines on the ethical use of AI, data privacy, the use of face recognition, algorithm bias and the role of big tech.

- AI can contribute to each of the 17 United Nations (UN) Sustainable Development Goals (SDGs) through use cases identified to-date that address about half of the 169 UN SDG targets, but bottlenecks still need to be overcome to deploy AI for sustainable development at scale.

PUBLIC DATA AND TOOLS

The AI Index 2019 Report supplements the main report with three additional resources: The raw data underlying the report, and two interactive tools, detailed below. We invite each member of the AI community to use these tools and data in a way most relevant to their work and interests.

Public Data

The public data is available on Google Drive. The Graphics folder provides hi-res images for all the charts.

The **Technical Appendix** contains sources, methodologies, and nuances.

Tools

- For those who want to focus on the extensive global data included in the report, we offer for the first time the Global AI Vibrancy Tool - vibrancy.aiindex.org - an interactive tool that compares countries across 34 indicators, including both a cross-country perspective and an intra-country drill down.

- The AI Index arXiv Monitor - arxiv.aiindex.org - is another tool that enables search of the full text of papers published to this pre-print repository, providing the most up-to-date snapshot of technical progress in AI.

Acknowledgements

We appreciate the individuals who provided data, advice, and expert commentary for inclusion in the AI Index 2019 report (in alphabetic organization order):

arXiv
Paul Ginsparg, Joe Halpern

BloombergGOV
Chris Cornillie

BurningGlass Technologies
Bledi Taska, Layla O'Kane

Campaign to Stop Killer Robots
Marta Kosmyna

Computing Research Association (CRA)
Andrew Bernat, Susan Davidson, Betsy Bizot

Coursera
Vinod Bakthavachalam, Eva Nierenberg

Elsevier
Maria de Kleijn, Clive Bastin, Sarah Huggett, Mark Siebert, Jörg Hellwig

GDELT Project
Kalev Leetaru

Indeed
Carrie Engel

Intento
Konstantin Savenkov, Grigory Sapunov

International Federation of Robotics
Susanne Bieller

Joint Research Center, European Commision
Alessandro Annoni, Giuditta DePrato

LinkedIn
Guy Berger, Di Mo, Mar Carpanelli, Virginia Ramsey

Metaculus
Ben Goldhaber, Jacob Lagerros

McKinsey Global Institute
Monique Tuin, Jake Silberg

Microsoft Academic Graph (MAG)
Kuansan Wang, Iris Shen, Yuxiao Dong

NESTA
Juan Mateos-Garcia, Kostas Stathoulopoulos, Joel Klinger

Paperwithcode
Robert Stojnic

PricewaterhouseCoopers (PwC)
Anand Rao, Ilana Golbin, Vidhi Tembhurnikar

Quid
Cara Connors, Julie Kim, Daria Mehra, Dan Buczaczer, Maggie Mazzetti

RightsCon
Nikki Gladstone, Fanny Hidvégi, Sarah Harper

Stanford University
Percy Liang, Mehran Sahami, Dorsa Sadigh, James Landay

Stockholm International Peace Research Institute (SIPRI)
Vincent Boulanin, Maaike Verbruggen

Udacity
Leah Wiedenmann and Rachel Kim

All the conference and university contributors; AI4All board; Bloomberg Philanthropies; WiML board;

Pedro Avelar (UFRGS), Dhruv Batra (Georgia Tech / FAIR); Zoe Bauer; Sam Bowman (NYU); Cody Coleman (Stanford); Casey Fiesler (University of Colorado Boulder); Brenden Lake (NYU); Calvin LeGassick; Natalie Garrett (University of Colorado Boulder); Bernard Ghanem (King Abdullah University of Science and Technology); Carol Hamilton (AAAI); Arthur Jago (University of Washington Tacoma); Zhao Jin (University of Rochester); Lars Kotthoff (University of Wyoming); Luis Lamb (Federal University of Rio Grande do Sul); Fanghzhen Lin (Hong Kong University of Science and Technology); Don Moore (UC Berkeley Haas School of Business), Avneesh Saluja (Netflix), Marcelo Prates (UFRGS), Michael Gofman (University of Rochester); Roger McCarthy (McCarthy Engineering); Devi Parikh (Georgia Tech / FAIR); Lynne Parker (White House Office of Science and Technology Policy); Daniel Rock (MIT); Ayush Shrivastava (Georgia Tech College of Computing); Cees Snoek (University of Amsterdam); Fabro Steibel (ITS-RIO); Prasanna Tambe (Wharton); Susan Woodward (Sandhill Econometrics); Chenggang Xu (Cheung Kong Graduate School of Business).

Matthew Kenney (Duke University) built the arXiv search engine tool. Tamara Pristc (Stanford) and Agata Foryciarz (Stanford) provided research contributions. Biswas Shrestha (Stanford) supported editing. Special thank you to Monique Tuin (McKinsey Global Institute) for invaluable comments and feedback.

Michael Chang (graphic design and cover art), Kevin Litman-Navarro (data visualization), Ruth Starkman (editor) and Biswas Shrestha (Stanford) were integral to the daily production of the report.

SYMBOLS

Pages appear with following symbols that denote global, sectoral, sub-regional, or other attributes for a given chapter.

Beginning: The first section of each chapter generally corresponds to either global, national, or regional metrics.

Middle: The middle section of each chapter corresponds to sectoral, cross country comparisons, or deep dives specific to each chapter.

End: The end section of each chapter offers sub-regional and state level analyses, results from cities, and data relevant to societal considerations of AI such as ethics and applications to the UN Sustainable Development Goals (SDG's) metrics.

Measurement Questions: Each chapter concludes with a short discussion on measurement questions related data and metrics presented in the chapter.

Chapter Preview

Chapter 1: Research and Development

[Table_of_Contents]
[Research_Development_Technical_Appendix]

Introduction

This chapter presents bibliometrics data, including volume of journal, conference and patent publications and their citation impacts by world regions. The chapter also presents Github Stars for key AI software libraries followed by societal considerations and gender diversity of AI researchers based on arXiv.

The Report has used different datasets to comprehensively assess the state of AI R&D activities around the world. The MAG dataset covers more publications than Elsevier's Scopus, which is mostly limited to peer-reviewed publications, but there are also publications on Scopus that are not in MAG.[2] arXiv, an online repository of electronic preprints, reflects the growing tendency of certain parts of the field of AI, particularly those depending on machine learning, to post papers before peer review, so reflects recent work more quickly than the other sources. Our arXiv Monitor tool uses full-text papers to quickly identify new results.

[2]see these studies by Anne-Wil Harzing and Martijn Visser.

Published Papers: AI Papers in All Publications

Elsevier's Scopus is the world's largest abstract and citation database of peer-reviewed literature with over 22,800 titles from more than 5,000 international publishers. The graph below (Figure 1.1) shows the percentage of AI publications in peer-reviewed publications (conferences, reviews, and articles) between 1998-2018. Here, AI papers correspond to all publications in AI, including journal publications and conference publications in the Scopus database. In the late 1990's AI papers accounted for less than 1% of articles and around 3% of conference publications. By 2018, the share of published AI papers in total papers has grown three-fold in 20 years, accounting for 3% of peer reviewed journal publications and 9% of published conference papers (see Appendix Graph).

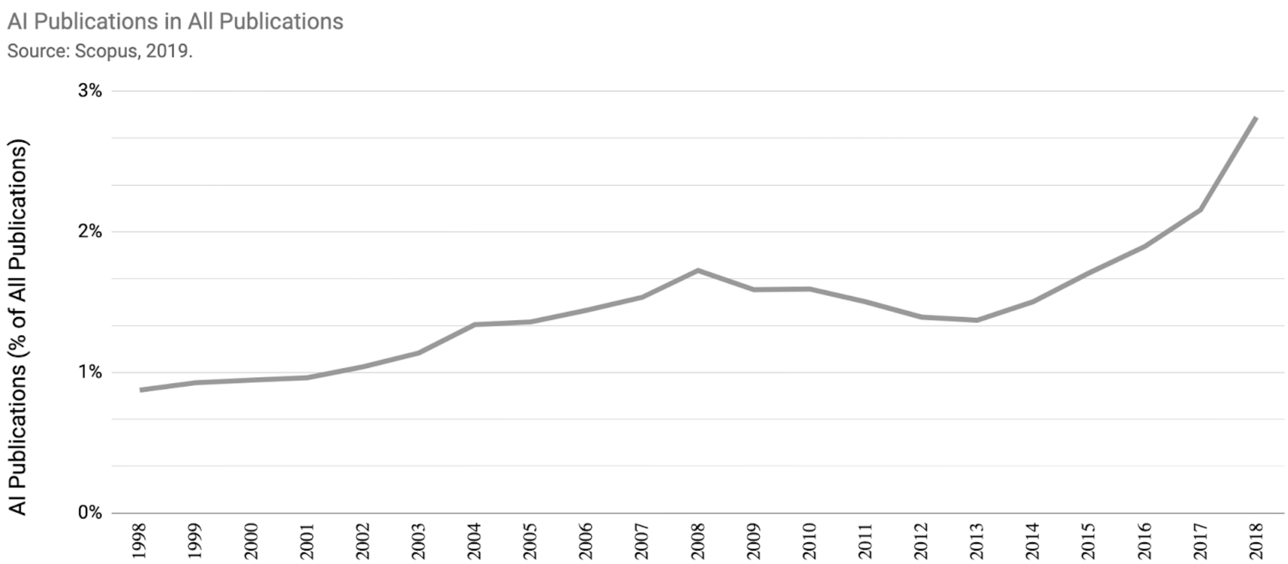

Fig. 1.1.

Between 1998 and 2018, the share of AI papers among all papers published worldwide has grown three-fold, now accounting for 3% of peer reviewed journal publications and 9% of published conference papers.

Published Papers: AI Papers By Region

Which regions witnessed the fastest growth in peer-reviewed AI publications? The graphs below show the number of AI papers published annually by region (Figure 1.2a), and the growth in AI papers published by region (Figure 1.2b). Europe has consistently been the largest publisher of AI papers — rising to over 27% of AI publications tracked by Scopus in 2018. Papers published from Chinese entities increased from 10% of global AI publications in 2000 to 28% in 2018 (see Appendix Graph). See Technical Appendix for data and methodology.

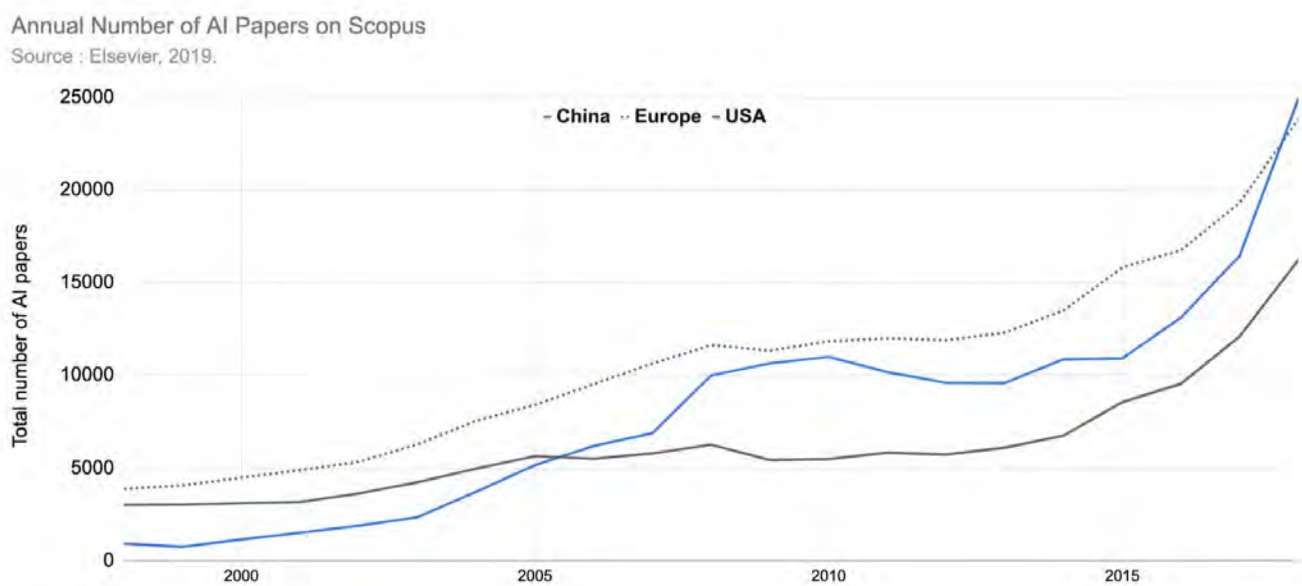

Fig. 1.2a.

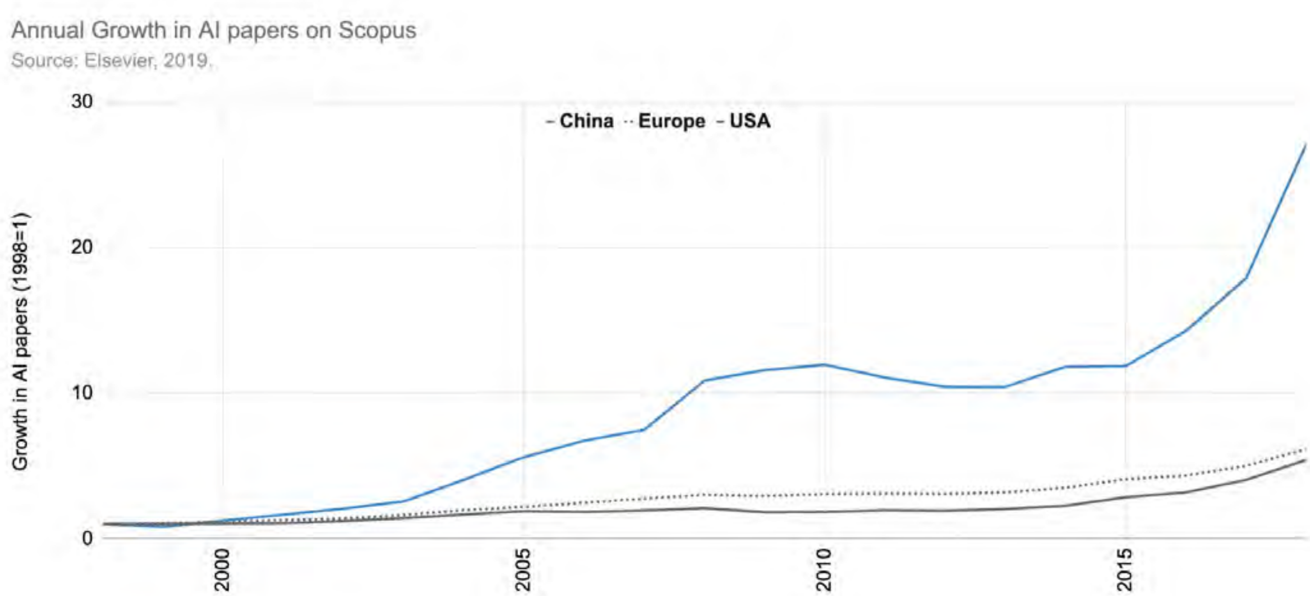

Fig. 1.2b.

Published Papers: Institutional Affiliation

The following graphs show the number of Scopus papers affiliated with government, corporate, medical, and other organizations in China (Figure 1.3a), the United States (Figure 1.3b), and Europe (Figure 1.3c). Excluding academia, the graphs show that government-affiliated institutions contribute the highest number of AI publications in China and Europe, whereas, corporate-affiliated AI papers make up a higher proportion in the US

In 2018, Chinese government institutions produced nearly three times more AI papers than Chinese corporations. China has also seen a 300-fold increase in government-affiliated AI papers since 1998, while corporate AI papers increased by 66-fold in the same period.

In the US., a relatively large proportion of AI papers are affiliated with corporations. In 2018, the number of corporate-affiliated AI papers in the US was over seven times the proportion of corporate AI papers in China, and almost twice that of Europe.

Note that in all three regions, academic papers (not shown) outweigh government, corporate, and medical papers by a large margin, making up 92% of AI publications from China, 90% from Europe, and 85% from the US Growth trends of institutional affiliation dynamics are available in the Appendix.

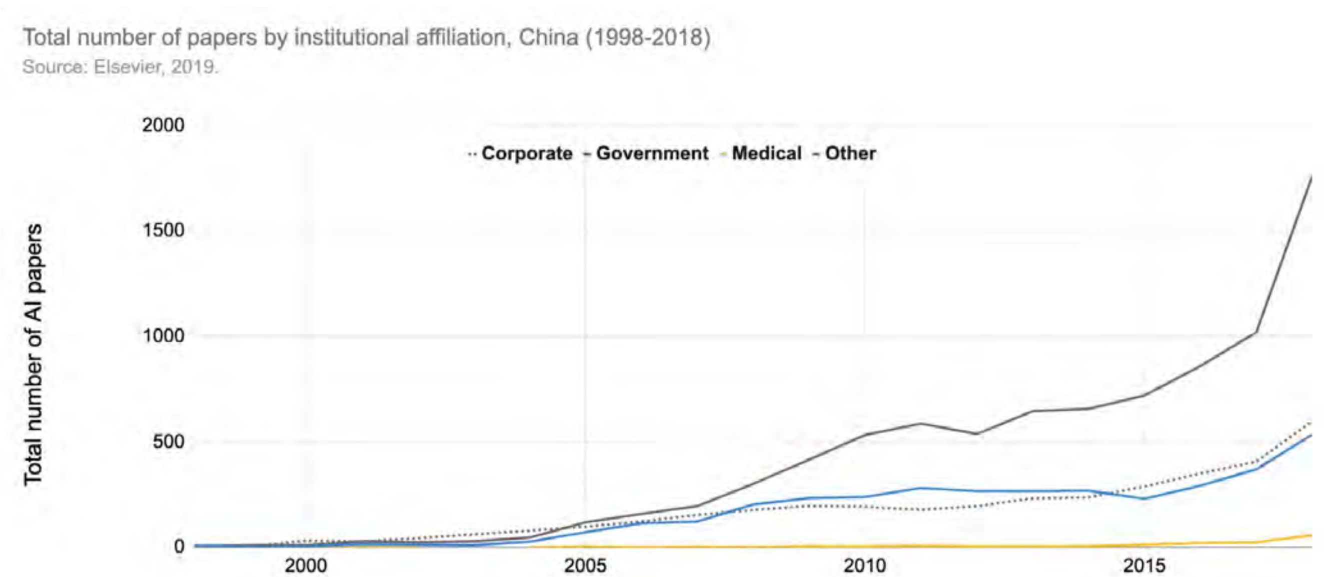

Fig. 1.3a.

[Research_Development_Technical_Appendix]
[Access_Data]

Published Papers: Institutional Affiliation

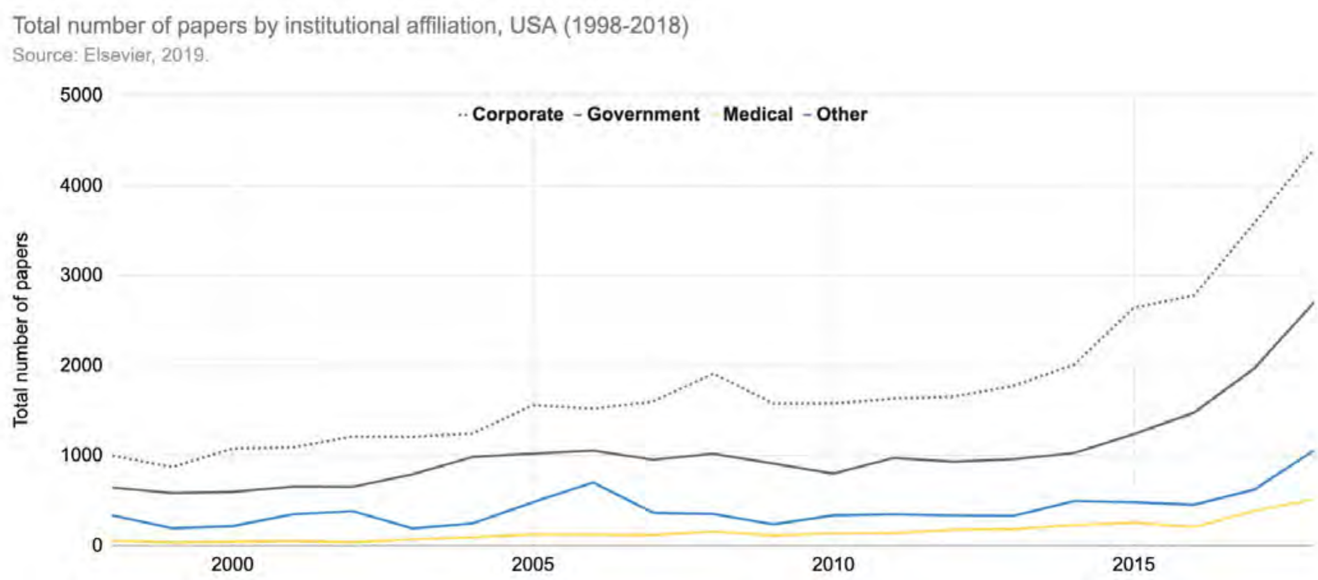

Fig. 1.3b.

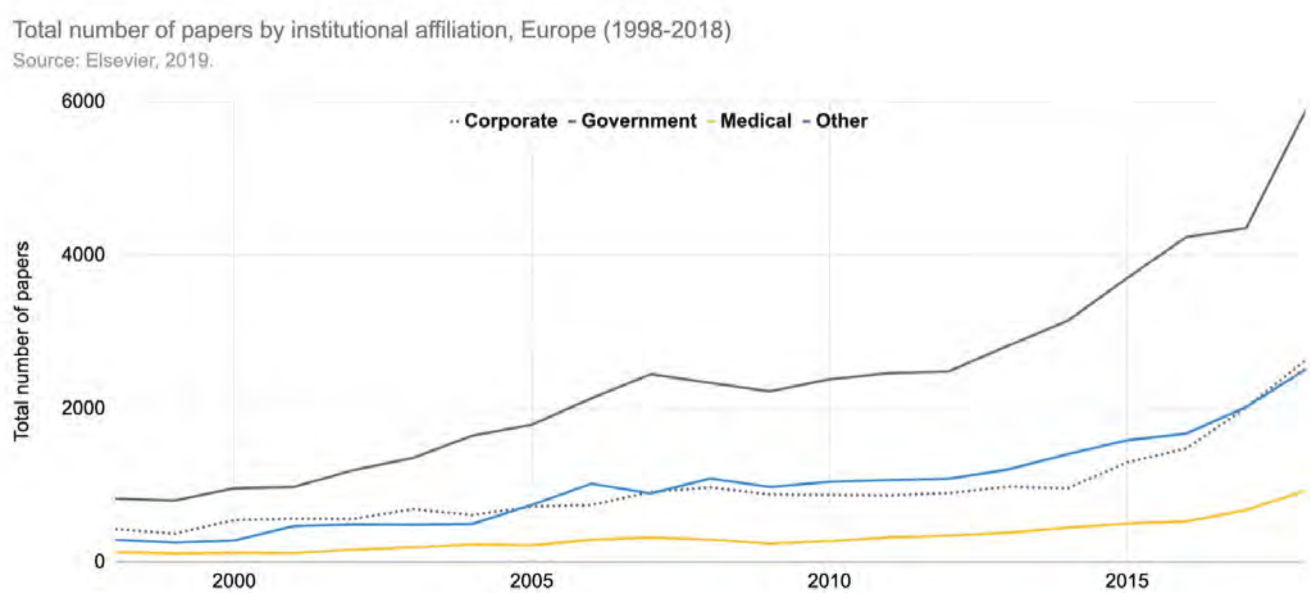

Fig. 1.3c.
Note: Europe refers to EU44.

Published Papers: Citation impact by region

The graph below (Figure 1.4) shows the average field-weighted citation impact of AI authors by region. A region's ***Field-Weighted Citation Impact (FWCI)*** is the average number of citations received by AI publications originating from that region divided by the average number of citations by all AI publications worldwide in the same publication year, subject area, and document type.

In this visual, the citation impacts are shown relative to the world average for AI, whose FWCI is normalized at 1. A re-based FWCI of 1 indicates that the publications have been cited on par with the world average for AI. A re-based FWCI of 0.85 indicates that the papers are 15% less cited than the world average for AI.

While Europe has the largest number of annually published AI papers in Scopus, Europe's FWCI has remained relatively flat and on-par with the world average. In contrast, China has increased its FWCI considerably. Still, the US outperforms other regions in total citations. Authors from the US are cited 40% more than the global average. See Technical Appendix for data and definitions. Both the US and China are gaining in prominence in Field-Weighted Download Impact (FWDI) of AI publications (see Appendix Graph).

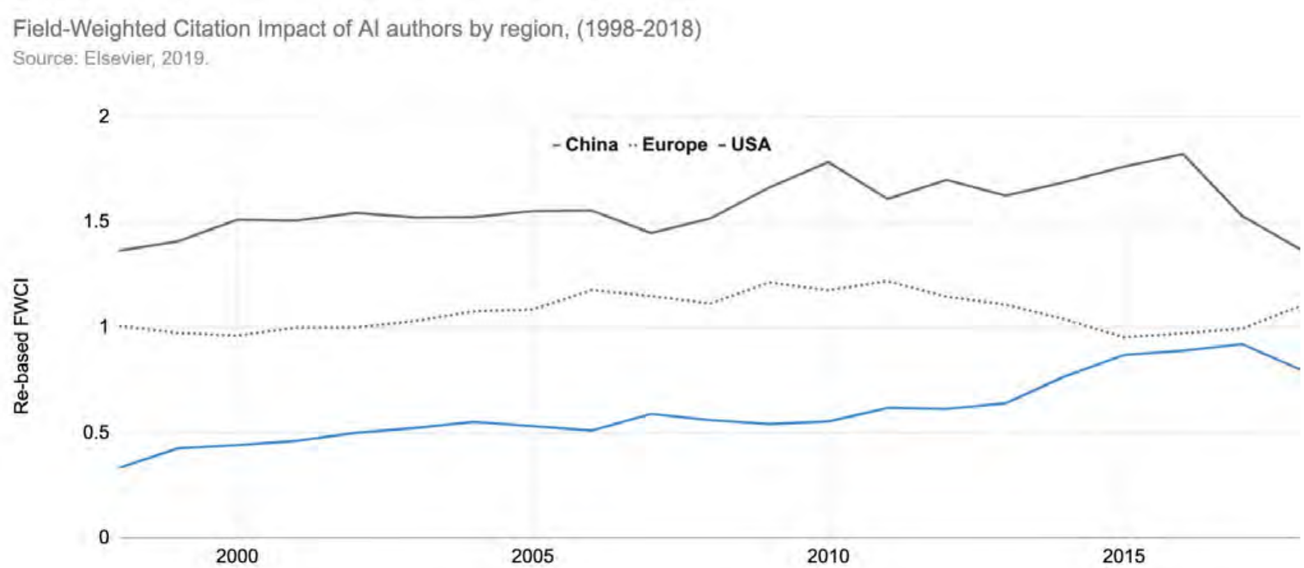

Fig. 1.4.

"China has consistently increased its footprint in AI research, both in terms of volume and quality. Their advance is truly remarkable."
Maria de Kleijn, SVP Analytical Services, Elsevier

Cross Country Trends in Impact and Academic-Corporate Collaboration

In recent years it's increasingly common for AI-focused companies to conduct research in partnership with colleagues in academia. This map (Figure 1.5a) shows the quantity of academic-corporate collaborations in different countries around the world. Academic-corporate collaborations are identified through publications with at least one author with an academic affiliation and at least one author with a corporate affiliation. Academic-corporate AI collaborations are largely prevalent in the US, China, Japan, France, Germany, and the UK.

World Map of Academic-Corporate Collaboration: Total Number of AI papers

Source: Scopus, 2019.

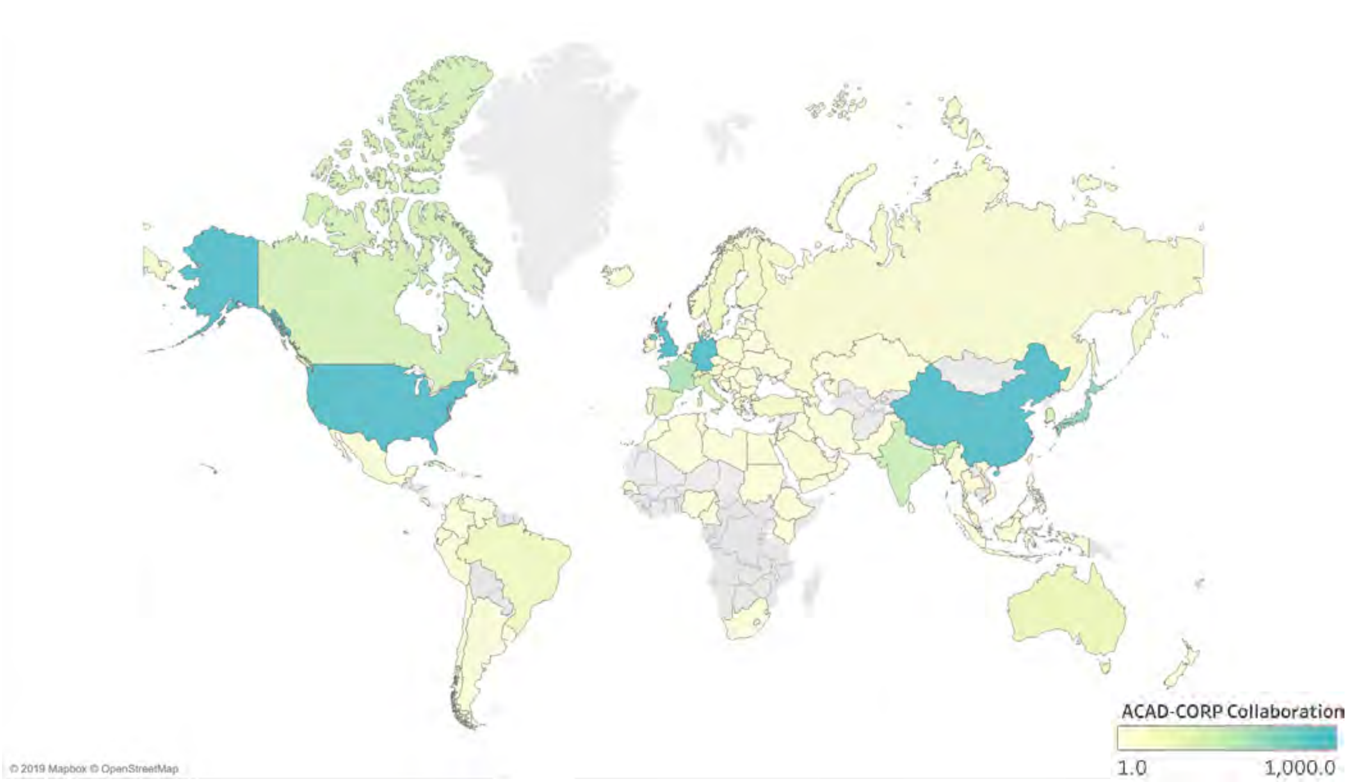

Fig. 1.5a.

Academic-corporate AI collaborations are prevalent in the US, China, France, Hong Kong, Switzerland, Canada, Japan, Germany, and the UK. US Authors are cited 40% more than the global average.

Cross Country Trends in Impact and Academic-Corporate Collaboration

How do academic-corporate collaborations impact the overall FWCI of AI research publications from different countries? This graph (Figure 1.5b) shows the FWCI (for all AI papers) on the y-axis and the total number of AI papers based on academic-corporate collaborations on the x-axis. The chart can be split into four quadrants: high degree of collaboration and high degree of impact (*top right quadrant*); low degree of collaboration but high impact (*top left quadrant*); low degree of collaboration and low impact (*bottom left quadrant*); high degree of collaboration but low impact (*top left quadrant*); Chart for countries across scholarly output metrics is available in the Appendix.

Four Quadrants for Overall AI Citation Impact (vertical axis) and the Total number of Academic-Corporate AI Papers (horizontal axisSource)

Source: Scopus, 2019.

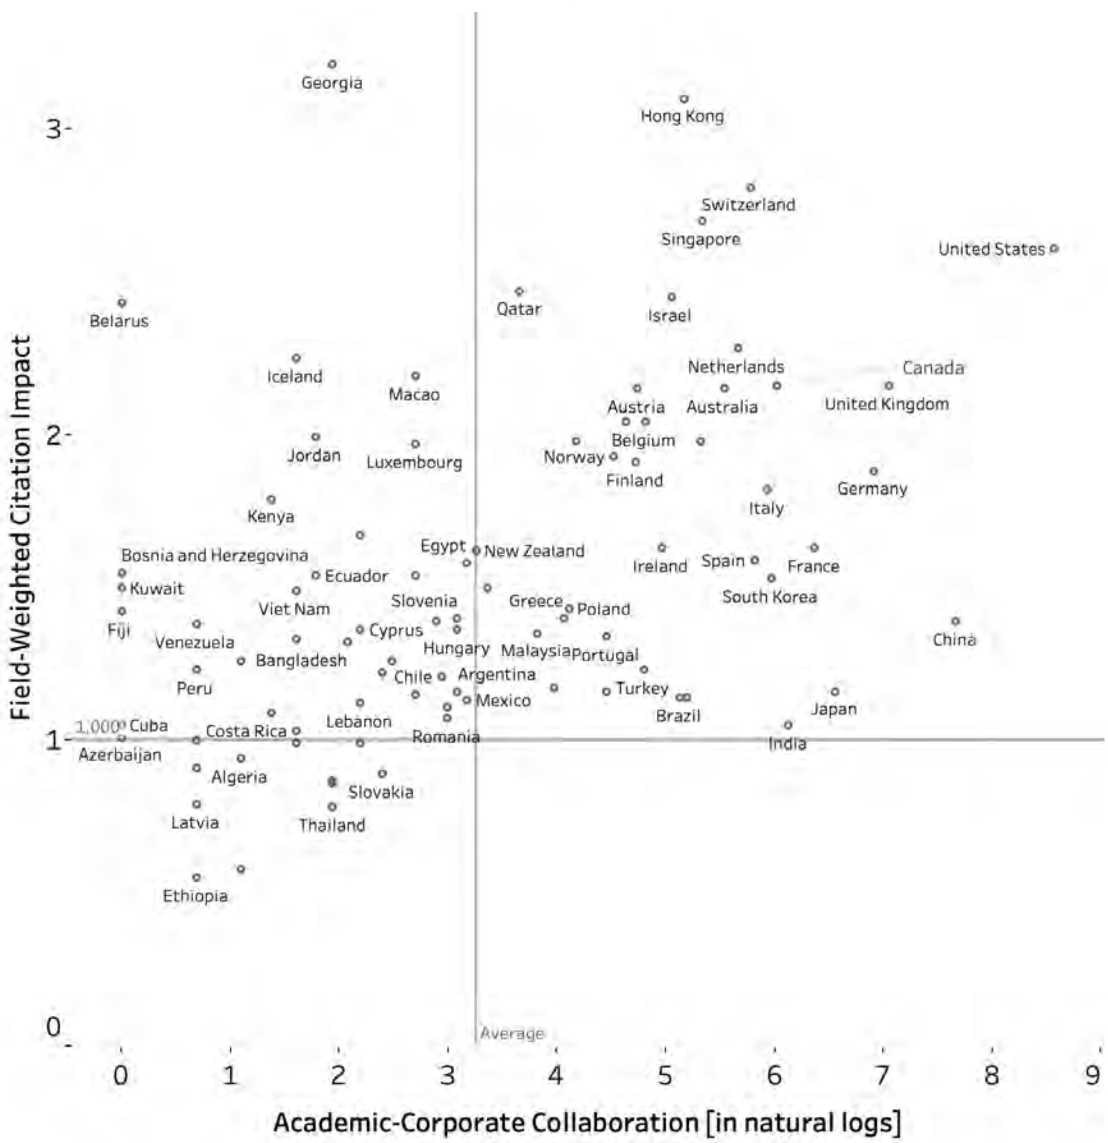

Fig. 1.5b.

"Counter to common assumptions, working together with Corporate institutions is beneficial to the academic impact of universities."
Maria de Kleijn, SVP Analytical Services, Elsevier

AI papers on arXiv

In recent years, AI researchers have adopted the practice of publishing paper pre-prints (frequently before peer-review) on arXiv, an online repository of electronic preprints. The graph below shows the number of AI papers on arXiv by each paper's primary subcategory (Figure 1.6).

The number of AI papers on arXiv is increasing overall and in a number of subcategories, reflecting a broader growth in AI researchers publishing preprints of their research. Between 2010 and 2019, the total number of AI papers on arXiv increased over twenty-fold. Submissions to the Computation & Language arXiv sub-category have grown almost sixty-fold since 2010.

In terms of volume, *Computer Vision (CV)* and *Pattern Recognition* had been the largest AI subcategory on arXiv since 2014 but *Machine Learning* has become the largest category of AI papers in 2019. In addition to showing a growing interest in *Computer Vision* and *Machine Learning* (and its general applied applications), this chart also indicates growth in other AI application areas, such as *Robotics* growing over thirty-fold between 2010 and 2019. See Technical Appendix for data and methodology.

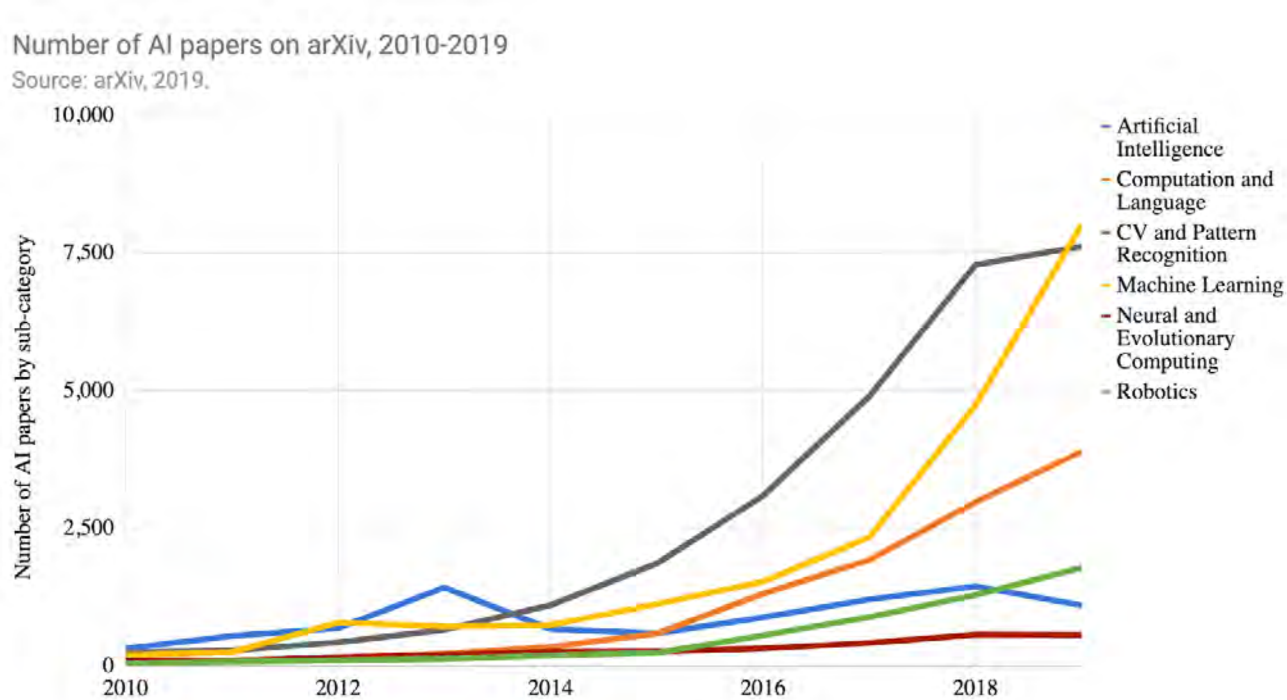

Fig. 1.6.

Deep Learning Papers on arXiv

Machine Learning (ML) is one of the most active research areas in AI. Within ML, Deep Learning (DL) approaches have become increasingly popular in recent years. The number of deep learning (DL) papers published on arXiv is increasing across regions. The first chart (Figure 1.7a) shows that North America published the largest volume of DL papers, followed by Europe in 2018. The volume of DL papers from East Asia reached the same level as Europe in 2018.

The following graphs show the ranking of countries with the largest volume of DL papers (Figure 1.7b) as well as the associated per capita DL papers (Figure 1.7c). Singapore, Switzerland, Australia, Israel, Netherlands, and Luxembourg have relatively high per capita DL papers published on arXiv. More details on methodology (see Technical Appendix) and detailed country chart (see Appendix Graph).

Number of Deep Learning Papers on arXiv

Source: arXiv, NESTA, 2019.

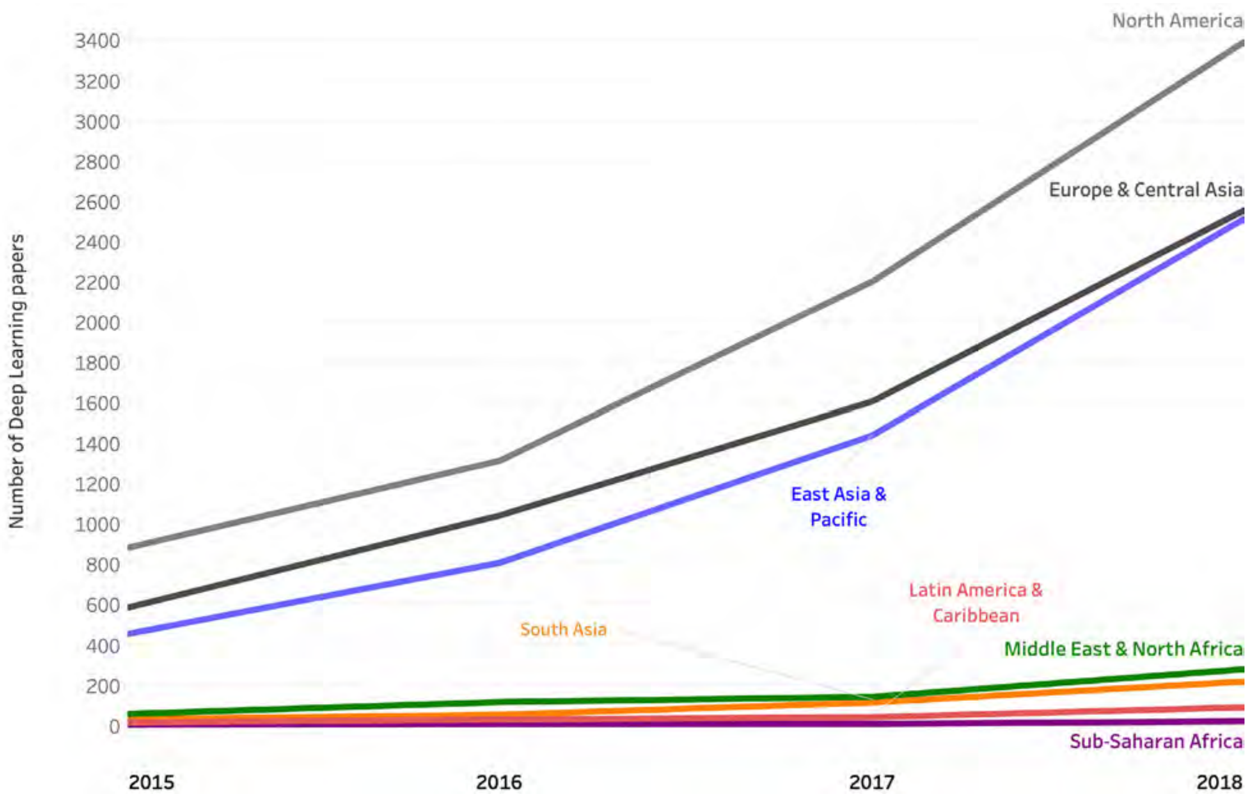

Fig. 1.7a.

Notes on World Regions:

In the following sections, cross-country bibliometrics analysis may correspond to World Bank region codes where explicitly stated. The regions include: East Asia & Pacific, Europe & Central Asia, Latin America & Caribbean, Middle East & North Africa, North America, South Asia, and Sub-Saharan Africa. "East Asia" can be referred to East Asia & Pacific and "Europe" to Europe & Central Asia. The country codes and API are available.

Singapore, Switzerland, Australia, Israel, Netherlands, and Luxembourg have relatively high per capita DL papers published on arXiv.

Deep Learning Papers on arXiv

Ranking Countries based on Total Number of Deep Learning Papers on arXiv, 2015-18

Source: arXiv, NESTA, 2019.

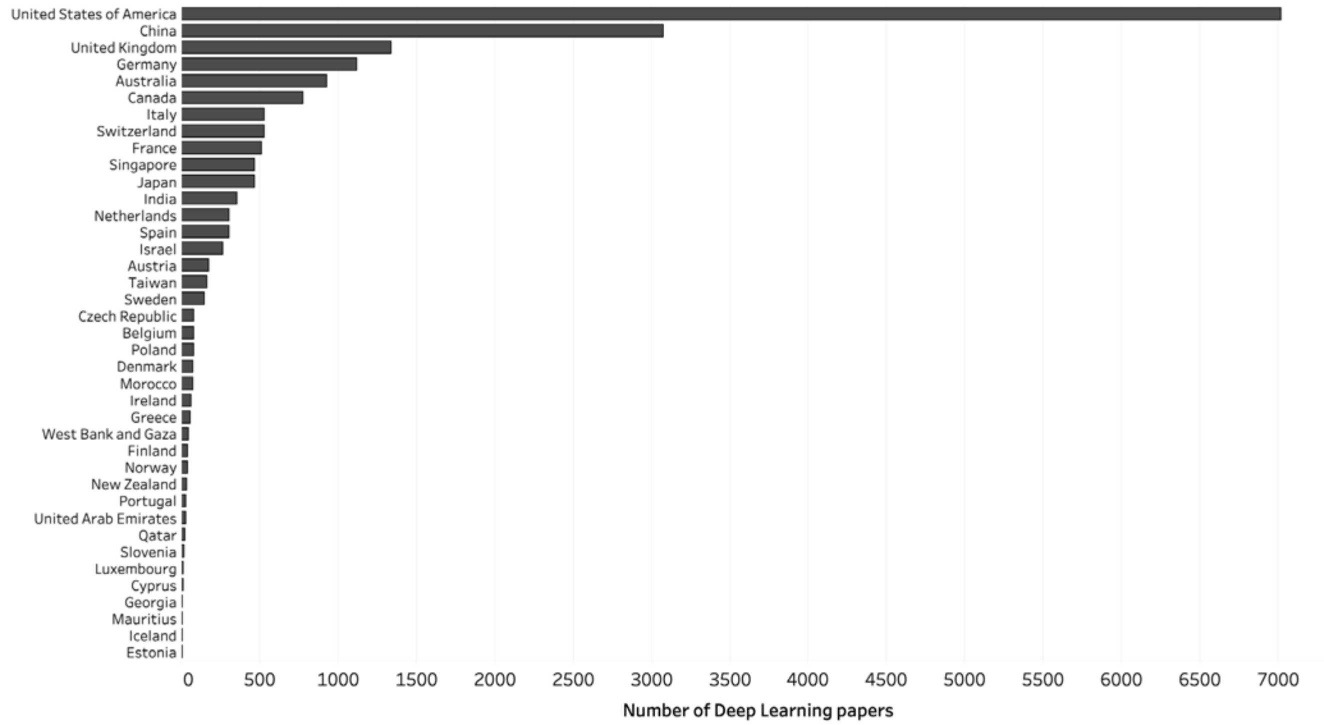

Fig. 1.7b.

Ranking Countries based on Number of Deep Learning Papers per capita on arXiv, 2015-18

Source: arXiv, NESTA, 2019.

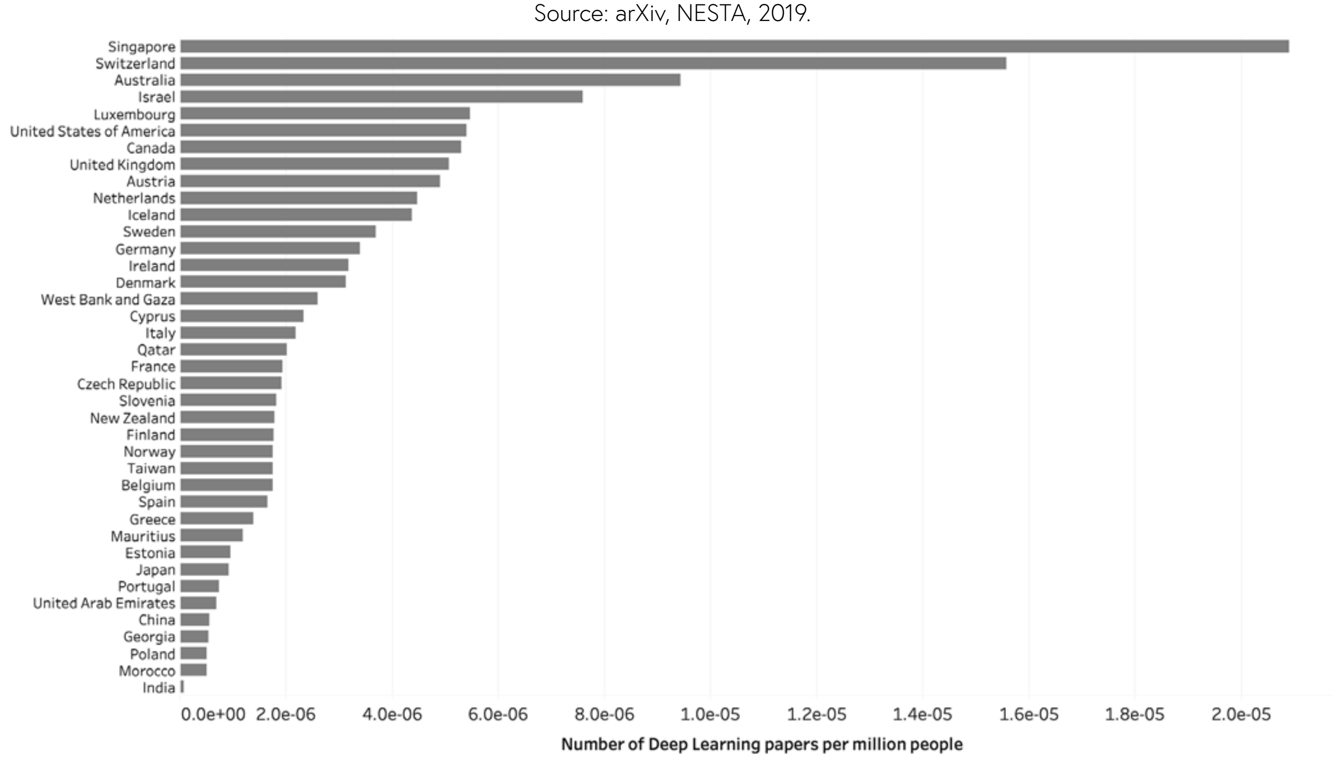

Fig. 1.7c.

Published Papers: AI Journal Publications

The graph below (Figure 1.8a) shows the share of AI journal papers on Microsoft Academic Graph (MAG) by world regions between 1990-2018. 37% of published journal papers are attributed to East Asia and Pacific (herein referred to as East Asia), 24% to Europe and Central Asia (herein referred to as Europe), and 22% to the North America in 2018. The share of South Asia in world AI journal publications has risen steadily to almost 8% in 2018.

The following graph (Figure 1.8b) shows the total number of journal publications and average journal publications per million people between 2015-18. China had the highest volume of AI papers, followed by the US, India, UK, and Germany. East Asia has the highest volume of AI journal papers on MAG (see Appendix Graph).

Share of World AI Journal Publications (%), 1990-2018

Source: MAG, 2019.

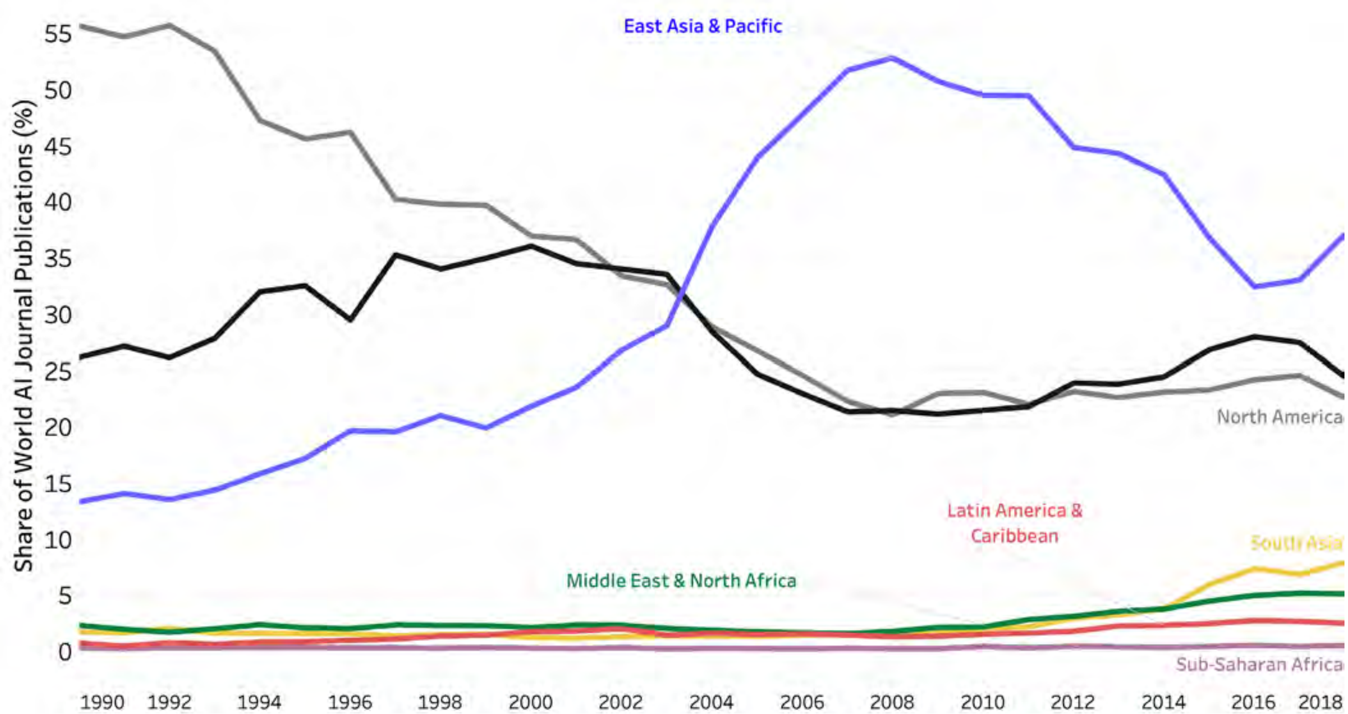

Fig. 1.8a.

Published Papers: AI Journal Publications

Total Volume and average annual per capita AI Journal Publications, 2015-2018

Source: MAG, 2019.

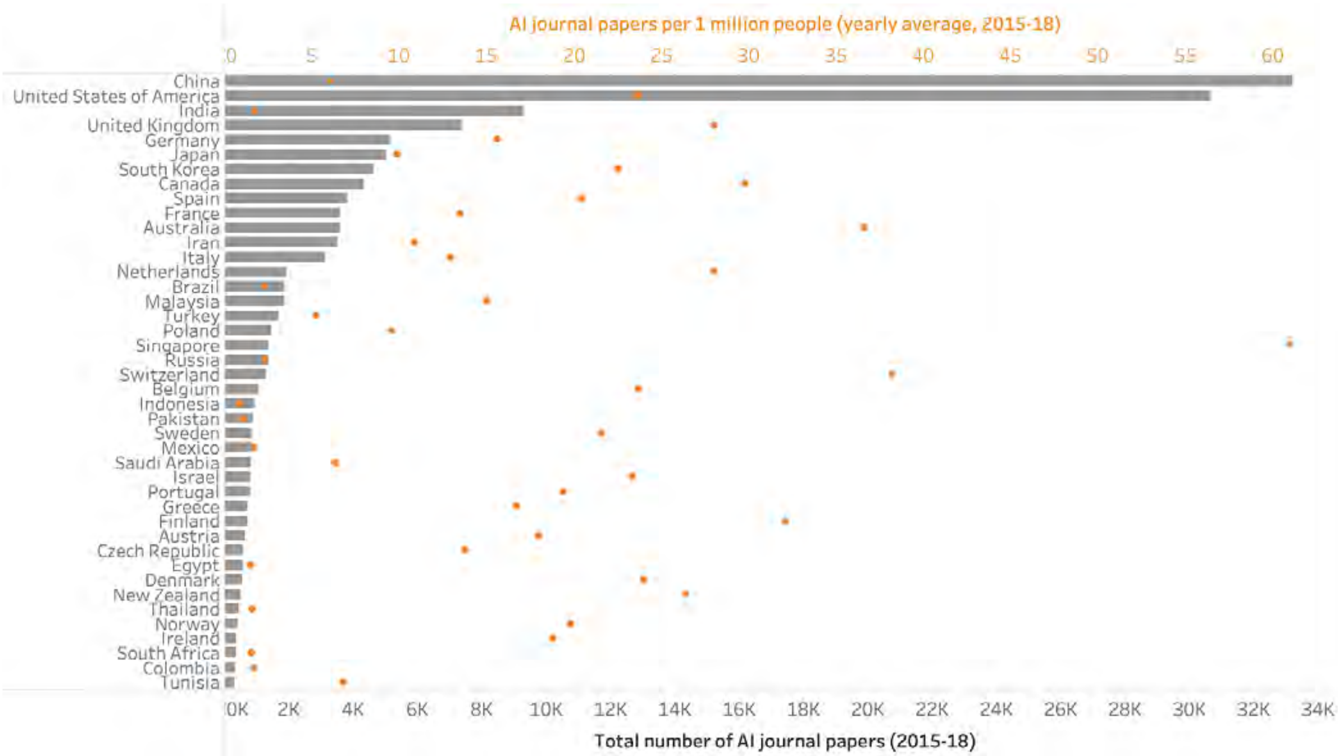

Fig. 1.8b.

In 2018, China had the highest volume of AI journal papers, followed by US, India, UK, and Germany.

Published Papers: AI Journal Citation

AI journal citation provides a signal for AI R&D impact. The share of world AI journal citation from **all** journal papers in MAG data is presented (see Box 1.1). North American papers were most cited by East Asian authors over 220k times, followed by European authors over 191k times. The interactive graphs are available on the web. Methodology paper A Century of Science: Globalization of Scientific Collaborations, Citations, and Innovations.

AI journal citations to East Asia journal papers account for over 32% of world citations; followed by Europe accounting for over 31%, and North America over 27% of world AI journal citations (Figure 1.9).

Box 1.1
•Between 2014-18, 17% of world citation was self-citation with East Asia; 15% was self-citation within Europe; 9% was self-citation within North America.

•Between regions, 8% of world citations were East Asian journals papers citing North American journal papers and 7% papers citing North American papers.

•7% of world citations were East Asian journal papers citing European papers. The share of European and North American journal papers citing East Asian journals was 5% of world citation each.

Note: Percentage of journal citations to unknown country is 19.1%. Self-citation in these sections is referred to citation from one region to the same, not the more conventional author-cites-self interpretation.

AI Journal Citation Attributed to Region (% of world journal citations), 2014-18
Source: MAG, 2019.

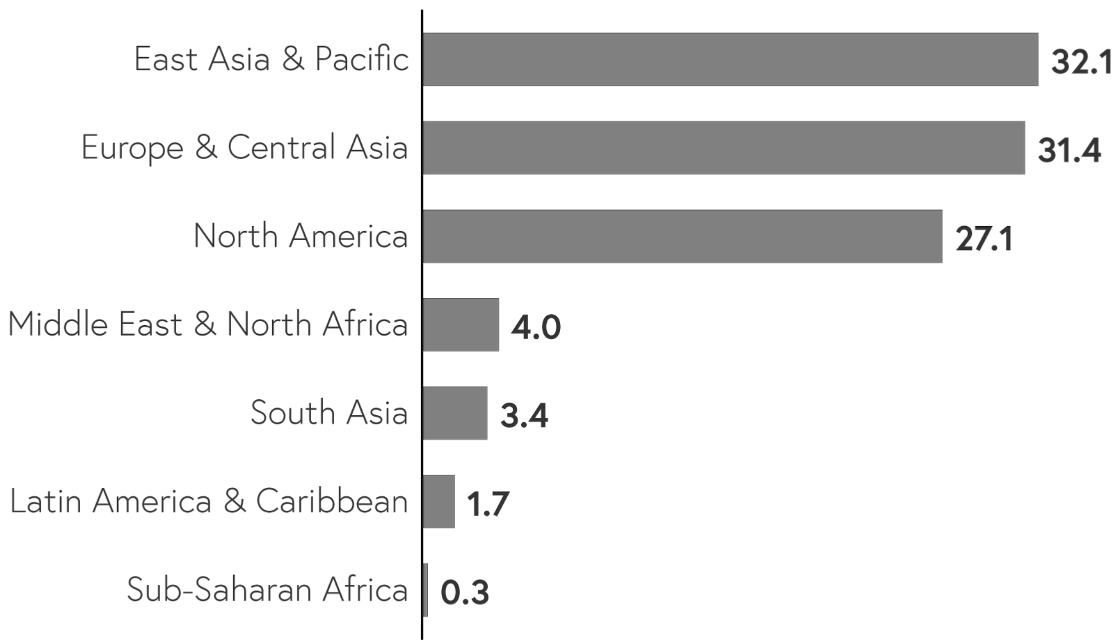

Fig. 1.9.
Note: Percentage of journal citations to unknown country is 19.1%.

AI journal citations to East Asia journal papers account for over 32% of world citations; followed by Europe accounting for over 31%, and North America over 27%

Published Papers: AI Conference Publications

The graph below (Figure 1.10a) shows the share of AI conference papers on MAG by world regions between 1990 and 2018. 33% of published AI conference papers are attributed to East Asia, 27% to North America, 26% to Europe in 2018. The share of South Asia in world AI conference publications has risen steadily to almost 6% in 2018.

The following graph (Figure 1.10b) shows the total number of AI conference publications and number of AI conference publications per million people between 2015-18. The US followed by China, India, Japan, and Germany had the highest volume of published AI conference papers. See Technical Appendix for data and methodology.

Share of World AI Conference Publications (%), 1990-2018

Source: MAG, 2019.

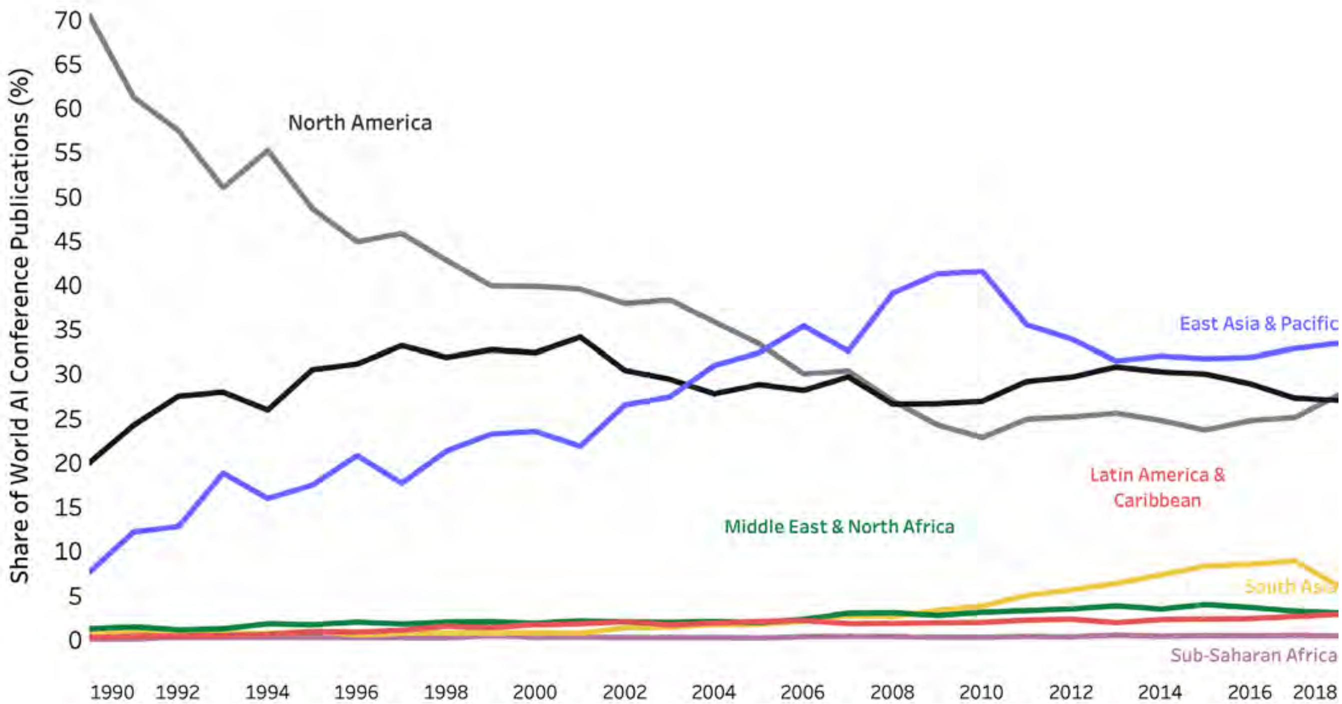

Fig. 1.10a.

Published Papers: AI Conference Publications

Total Volume and average annual per capita AI Conference Publications, 2015-2018

Source: MAG, 2019.

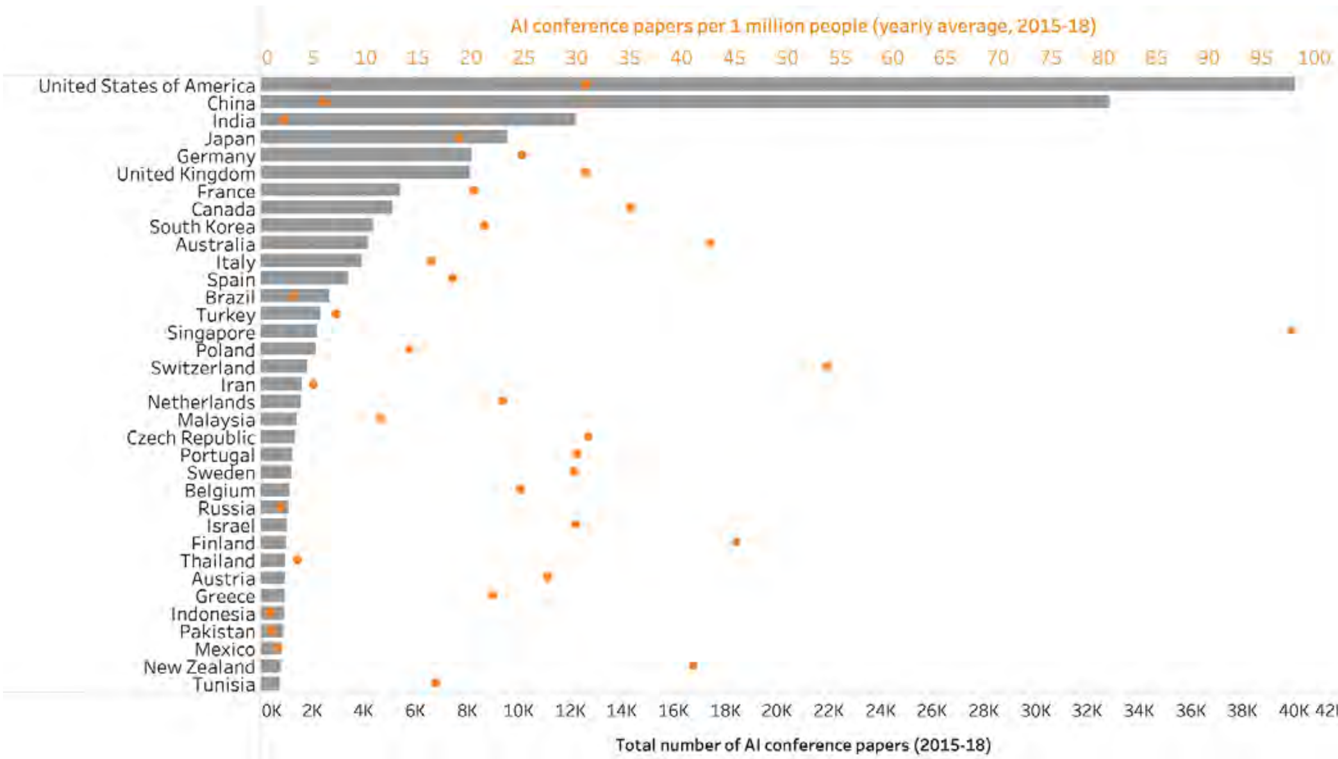

Fig. 1.10b.

Published Papers: AI Conference Citation

Over 40% of world conference paper citations are attributed to North America (self citation - 17%, East Asia - 13%, Europe - 10% of world citation). Self-citation in Europe accounted for 13% and self-citation in East Asia accounted for 11% of world conference publication citation. Box 1.2. presentes the highlights for conference citation and the interactive graphs are available on the web.

Almost 43% of world conference citations in AI papers is attributed to North American conference papers. The share of world citation in AI conference papers to European papers was over 28%, and to East Asian papers was over 22% of world AI conference citation activity (Figure 1.11).

Box 1.2.
•Citations to European conference papers by North America and East Asia accounted for 7% and 6% respectively of world conference citation.

•Citation to East Asian papers by North America and Europe accounted for 6% and 4% respectively of world conference citation

Note: Percentage of conference citations to unknown country is 12.7%. Self-citation in these sections is referred to citation from one region to the same, not the more conventional author-cites-self interpretation.

AI Conference Citation Attributed to Region (% of world journal citations), 2014-18

Source: MAG, 2019.

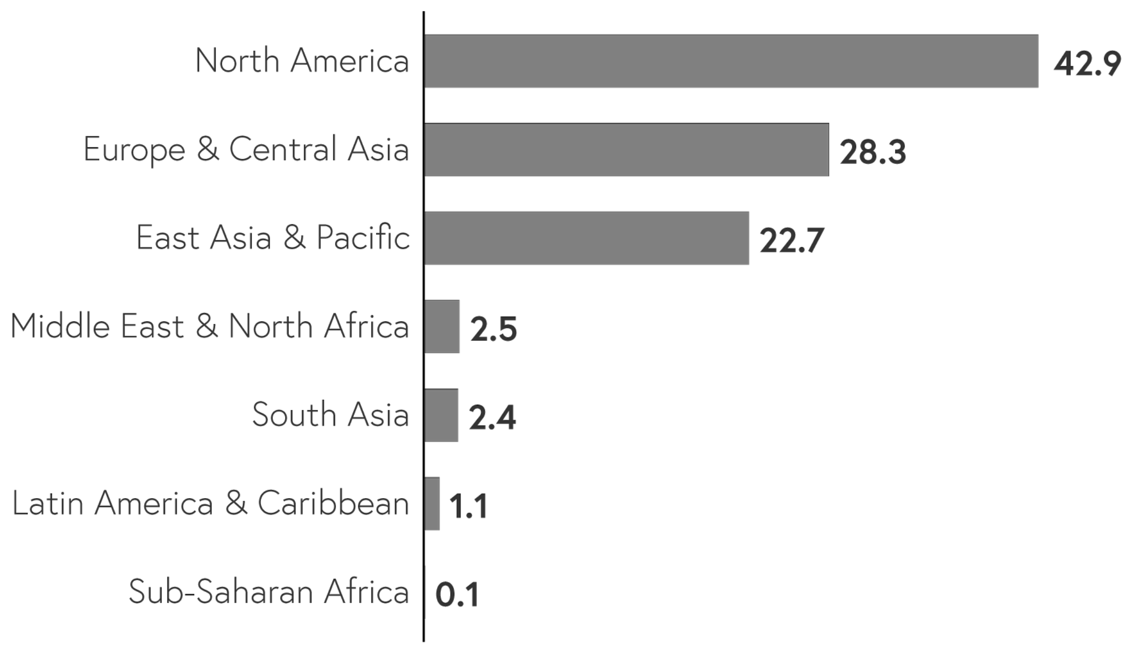

Fig. 1.11.
Note: Percentage of conference citations to unknown country is 12.7%.

Over 40% of world AI conference paper citations are attributed to North America (regional self citation - 17%, from East Asia - 13%, from Europe - 10% of world citation).

AI Patents

Patents on AI technology provide a measurement of AI activity in industry and its potential impact on products.The graph below (Figure 1.12a) shows the share of AI patents on MAG by world regions between 1990-2018. The graph for total number of AI patents published by regions can be found in the Appendix. Over 51% of published AI patents are attributed to the North America, with the share of Europe and Central Asia declining to 23%, close to East Asia & Pacific.

The following graph (Figure 1.12b) shows the total number of AI patents and average per capita AI patent publications between 2015-18. The US published three-folds the number of AI patents of the next country, Japan. Over 94% of AI patents are filed in high income countries, with the share of upper middle-income countries rising to 4% in 2018 (see Appendix Graph). See Technical Appendix for data and methodology.

Share of World AI Published Patents (%), 1990-2018

Source: MAG, 2019.

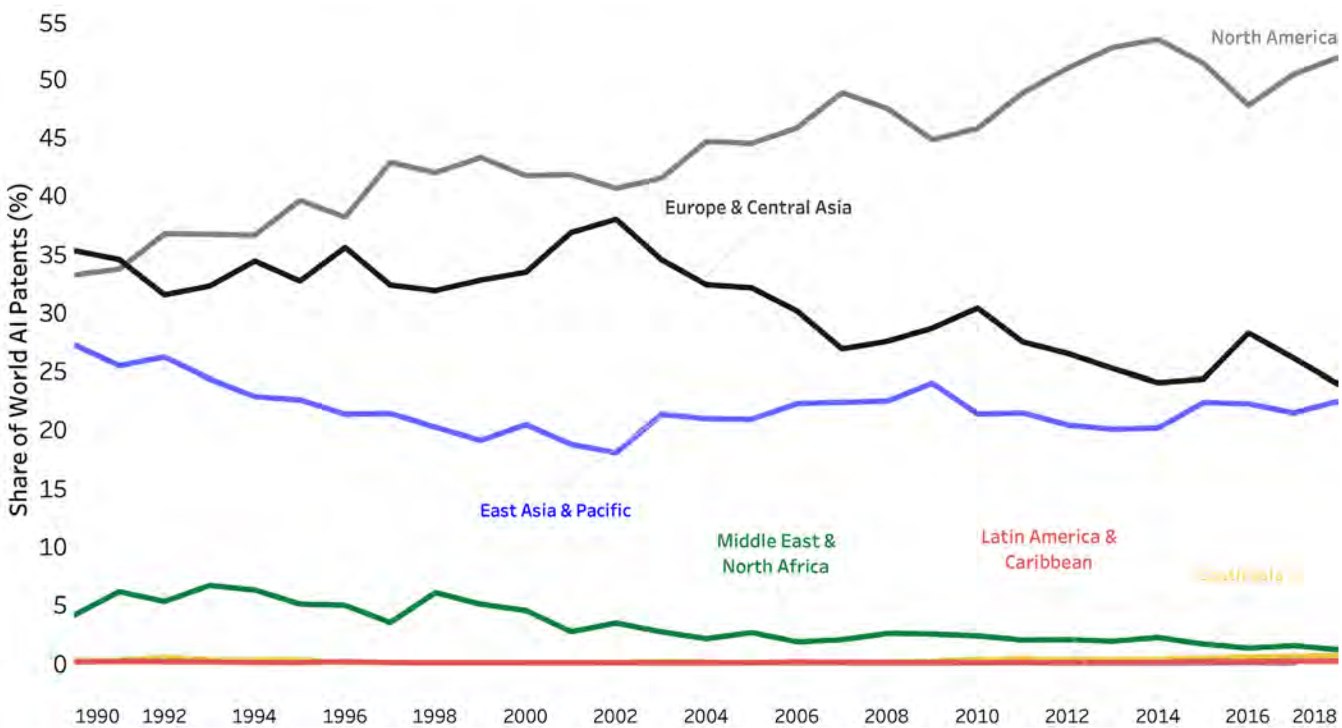

Fig. 1.12a.

AI Patents

Total Volume and average annual per capita AI Published Patents, 2015-2018

Source: MAG, 2019.

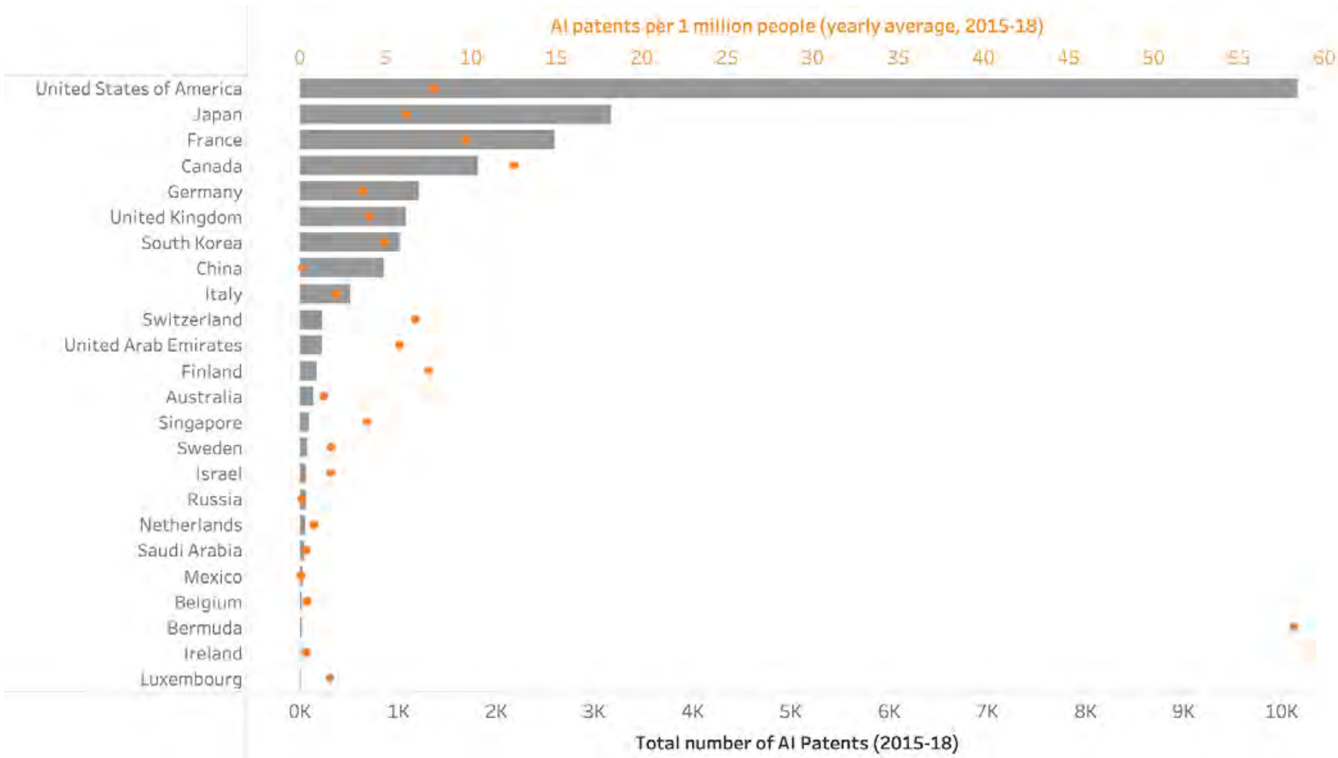

Fig. 1.12b.

Over 94% of AI patents are filed in high income countries, with the share of upper middle-income countries rising to 4% in 2018.

AI Patents Citations

The box below (Box 1.3) presents highlights AI patent citation from all patents. The insights on patent citation is revealing. Majority of world AI patent flow is dependent on North America. The interactive graphs are available on the web.

Citations to North American AI patents accounted for over 60% of world patent citation activity; followed by East Asia with over 22%, and Europe with over 17% of AI patent citation (Figure 1.13).

Box 1.3.
•Over 60% of AI patent citation activity is related to North America, with almost 45% (of world AI patent citation) self-citation, 9% from East Asia patents, and 7% from European patents

•North American patents cited European and East Asian patents around 6,000 times betwen 2015-18, with the individual regions accounting for 6-7% each of world patent citations

note: Percentage of patent citations to unknown country is 37.2%

AI Patent Citation Attributed to Region (% of world journal citations), 2014-18
Source: MAG, 2019.

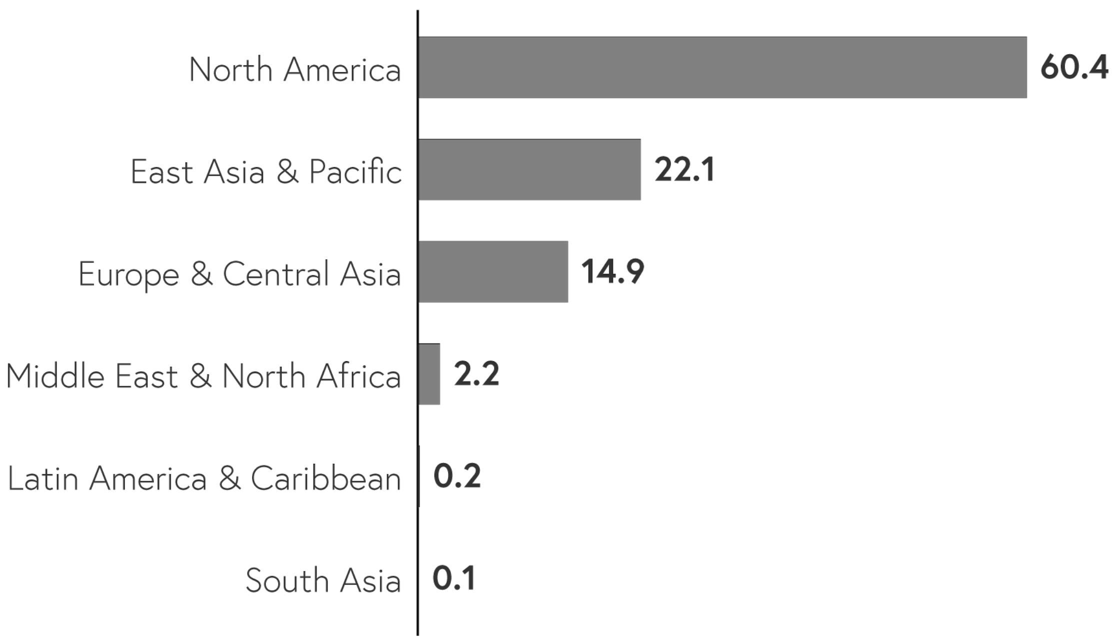

Fig. 1.13.
Note: Percentage of patent citations to unknown country is 37.2%.

North America accounts for over 60% of global AI patent citation activity between 2014 and 2018.

GitHub stars

GitHub is a website where developers upload, comment on, and download software code. Stars indicate a person has expressed interest in a particular piece of code and/or project on GitHub, similar to how 'likes' on social media services like Twitter and Facebook can indicate popularity of a given post. GitHub Stars therefore provide a rough measure of the popularity of various AI-programming frameworks. The graphs below show the number of times various AI and ML software packages have been starred on GitHub (Figure 1.14a and 1.14b).

One noticeable trend is the emergence of corporate-backed research frameworks, like Tensorflow (which was developed predominantly by Google) and PyTorch (which was developed predominantly by Facebook). Note that Keras popularity appears to tail off, but Keras has subsequently been integrated into TensorFlow, so its popularity is partially reflected in that metric. Two non-industry frameworks, sci-kit learn and Caffe, continue to show growing popularity, but their growth trajectories appear lower than those of the corporate frameworks.

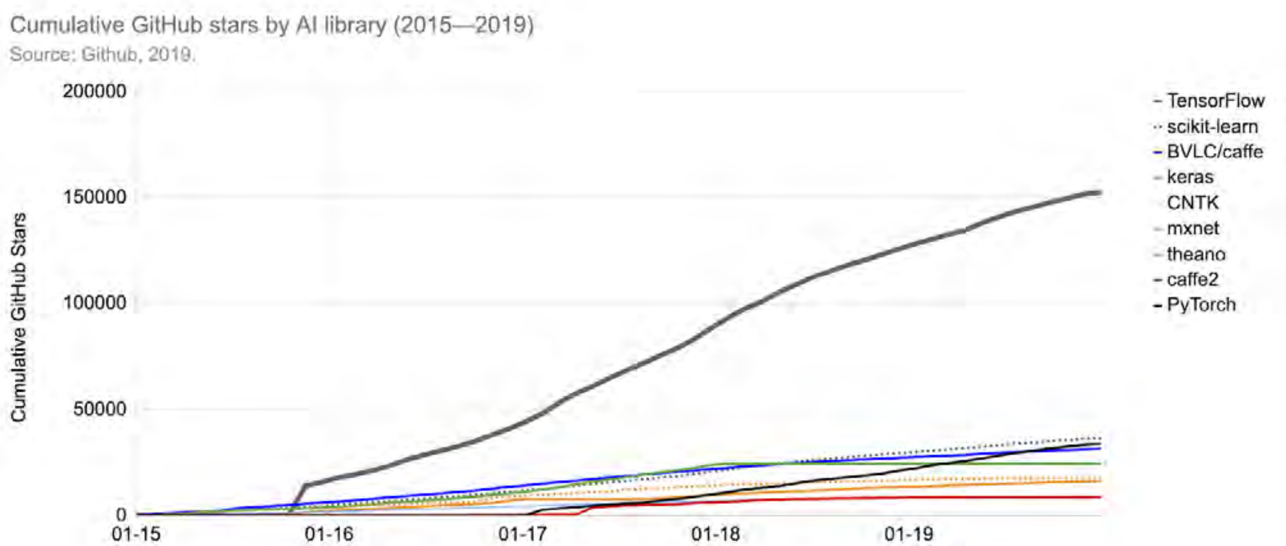

Fig. 1.14a.

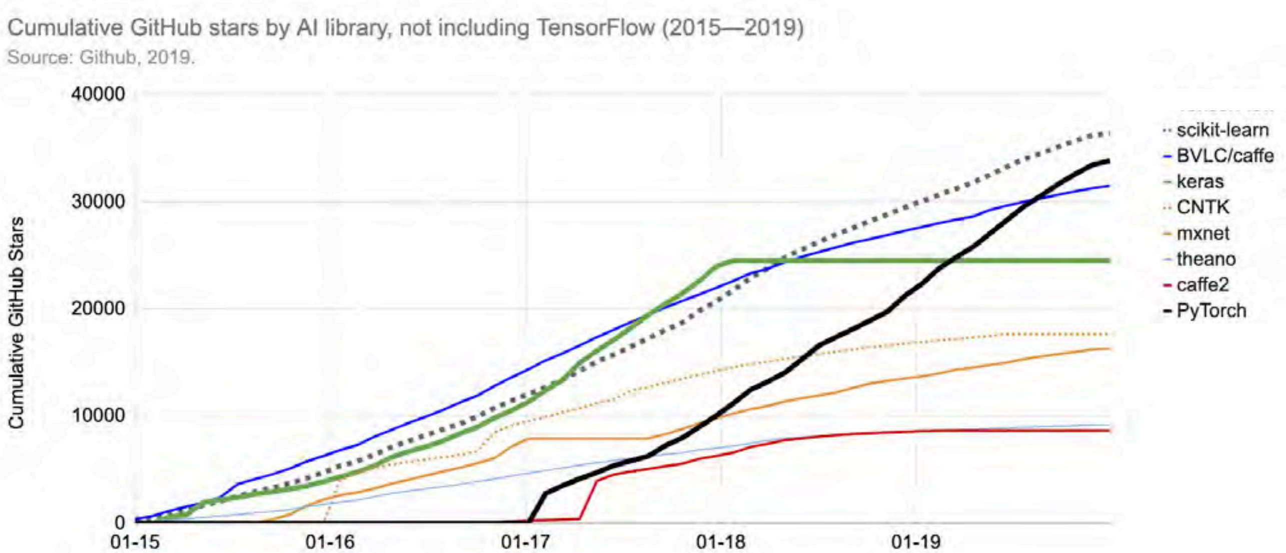

Fig. 1.14b.

Women in AI

There are significant international differences in the gender diversity of AI researchers. Half of the authors could be gender-labelled by first name with a high degree of confidence (China, one of the world leaders in AI research, is excluded from the sample due to a lower confidence in gender-labelling authors by name, and will be included in 2020). Countries with less than 5,000 publications on arXiv are not considered in this analysis. Technical Appendix provides details on data and methodology.

The differences between the share of female authors in AI and non-AI (refers to publications in all fields) papers within countries are presented below (Figure 1.16a). Over 41% of the AI papers in the Netherlands and over 39% of AI papers in Denmark had at least one female co-author. By contrast, only 10 per cent and 16 per cent of those with Japanese and Singaporean affiliations had a female co-author.

Countries such as Malaysia, Denmark, Norway and Israel show a stronger presence of women in AI research relative to non-AI papers.

The Women in AI report from NESTA can be found here. The longitudinal country data showing the share of female authors in AI and non-AI publications from NESTA is available here with the 30 countries with most publications. The change in share of women authors in AI is presented from 2000-2018, showing growth in AI publications with female authors from Europe (Figure 1.16b). Several countries have women as authors of over 30% of AI papers on arXiv including Argentina, Canada, Iran, and many European countries (Portugal, Spain, France, Belgium, Italy, Netherlands, Denmark, Ireland, Hungary). In the United States, the share of women authors in AI has decreased slightly over this period.

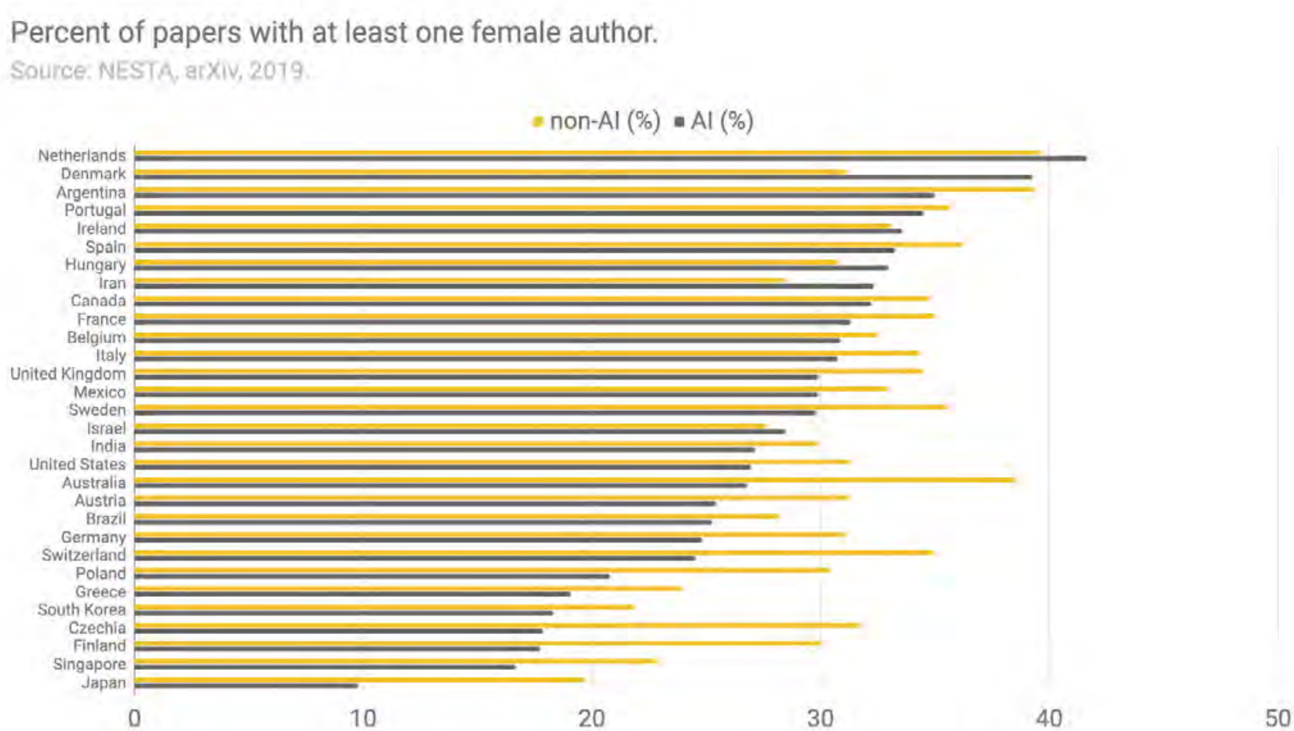

Fig. 1.16a.

"Our findings suggest that both geography and research domains play a role in influencing participation of women in AI publications. This means that national policies and institutions and social norms in research communities will both need to play a role in increasing female participation in AI research."
Kostas Stathoulopoulos and Juan Mateos-Garcia, NESTA

Women in AI

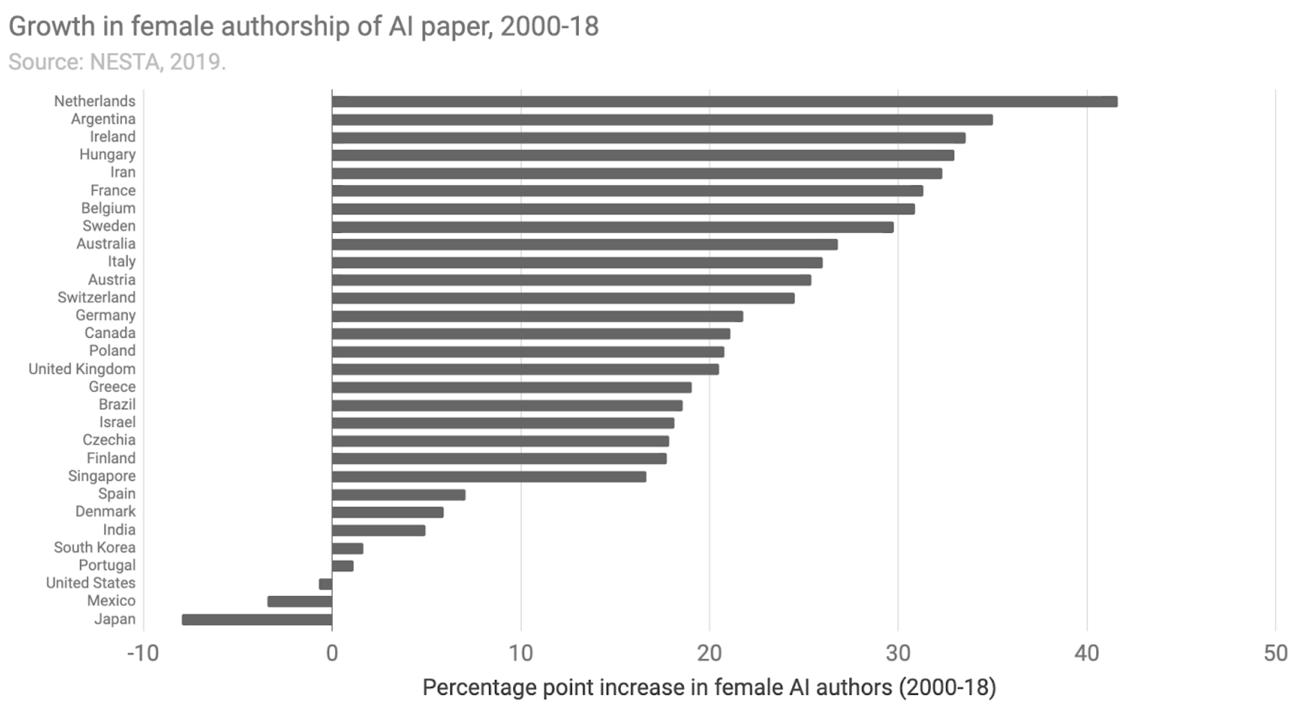

Fig 1.16b.

Many Western European countries as well as Argentina, Canada, and Iran show relatively high presence of women in AI research.

Measurement Questions

Some questions implied by the data in this section include:

- What is the best way to weight the relative importance of paper publications on preprint services like arXiv versus traditional journal publications?
- What tools are available to help us neatly attribute papers to a specific region or originating institution and/or funding source?
- Is it possible to measure and assess the gender of AI researchers without the addition of specific metadata to preprints and published papers?

Chapter Preview

Chapter 2: Conferences

[Table_of_Contents]
[Conferences_Technical_Appendix]

Introduction

This chapter presents data from a variety of sources on AI conference attendance and summary of topics and policy milestones achieved. First, the attendance at large and small AI conferences is presented. Second, acceptance rate by countries and AI subject areas at the AAAI conference is presented. Similar trends can be identified for other key AI conferences in the future. Third, growth in attendance and participation is presented for gender diversity organizations, and the mention of ethics at select conferences, highlighting the growing interest at the intersection of human rights and AI.

It should be noted that this data does not include the full scope of organizations dedicated to increasing participation of underrepresented individuals in AI, of which there are many, and which will be covered in the 2020 edition. The AI Index is still gathering data for organizations that measure racial and ethnic diversity in the field. For instance, Black in AI, is a vibrant effort. Other conferences that have a specific workshop or component dedicated to ethical challenges include ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*), AAAI/ACM Artificial Intelligence, Ethics, and Society conference (AES), FAT/ML at ICML), NeurIPS Joint Workshop on AI for Social Good.

Large AI Conferences

Conferences strongly indicate a level of industry and academic enthusiasm of a subject. AI conferences have grown not only in size but in number and prestige as well. The graphs below show attendance at large AI conferences from 1984 to 2019 (Figure 2.1a), and growth of large conference attendance relative to 2012 (Figure 2.1b). Large AI conferences are defined as those with over three thousand attendees in 2019. In 2019, NeurIPS 2019 will have 13,500 people, CVPR had about 9,227 people, ICML had about 6,400 people, and IJCAI-19 had 3,015 people. NeurIPS (formally NIPS), CVPR, and ICML, remain the most attended AI conferences. NeurIPS and ICML are growing at the fastest rate — with over eight-fold increase relative to their 2012 attendance. Publication venues are provided in the Technical Appendix.

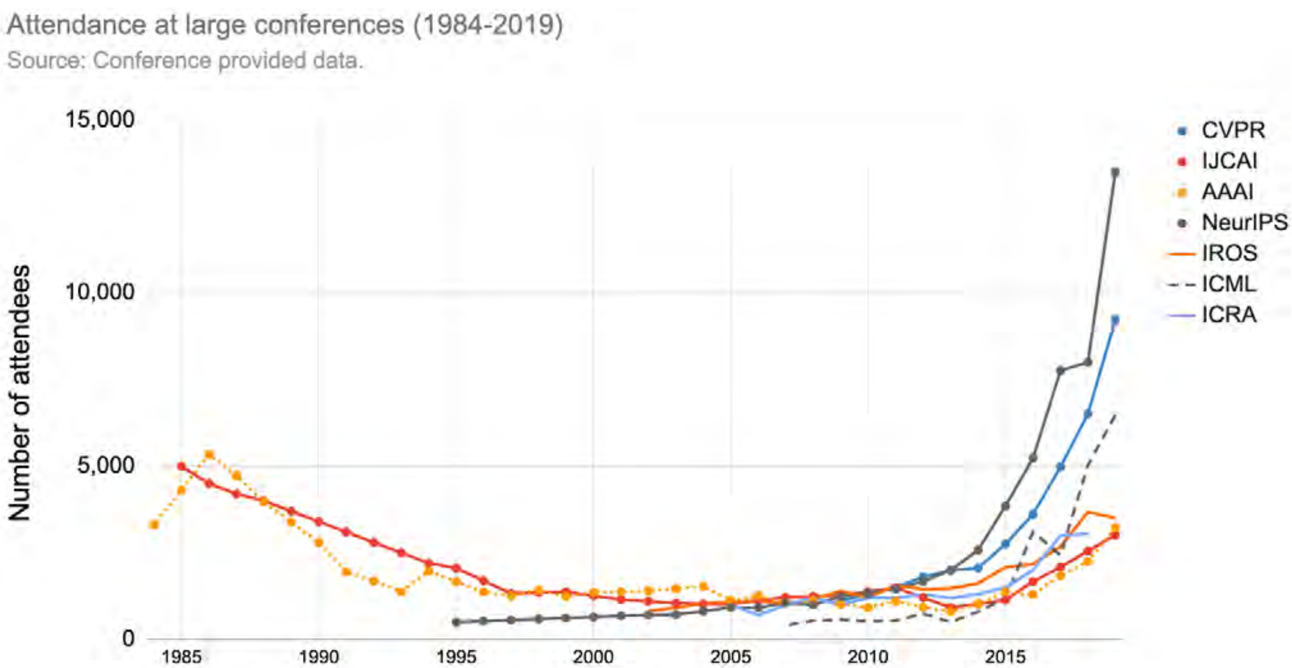

Fig. 2.1a

Note: IJCAI occurred every other year till 2014. The missing year between 1984 and 2014 are interpolated as the mean between the two known conference attendance dates to provide a comparative view across conferences.

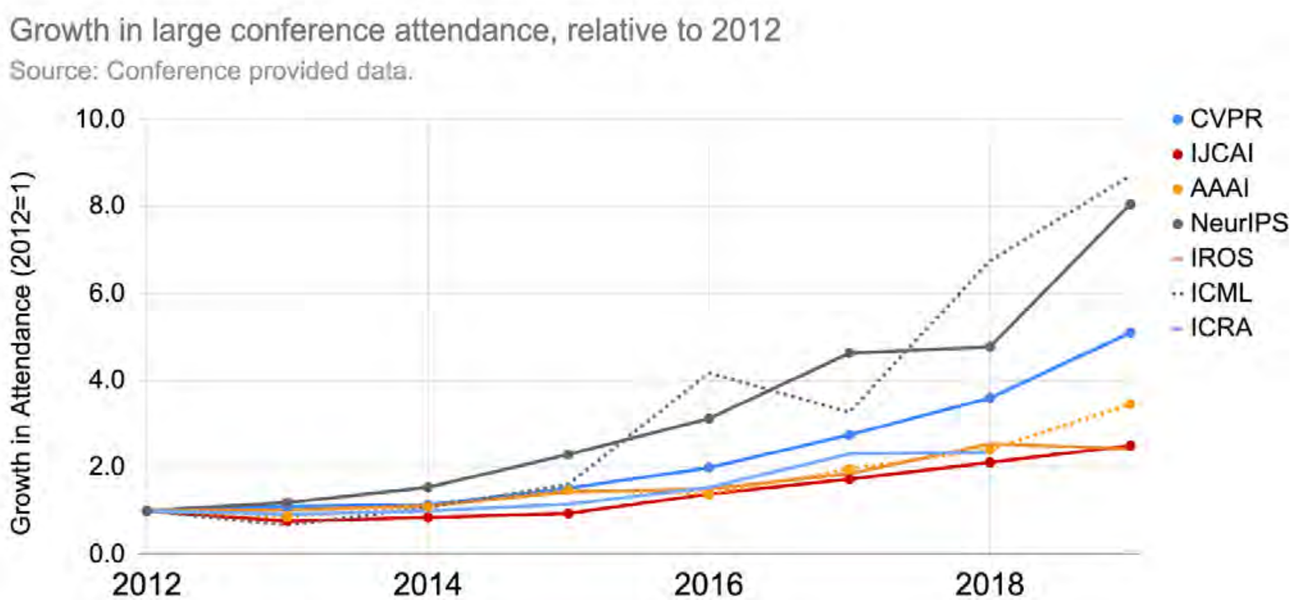

Fig. 2.1b

Small Conferences

The graphs below show attendance at small AI conferences (Figure 2.2a), and growth of small AI conference attendance relative to 2014 (Figure 2.2b). Small AI conferences are defined as those with under three thousand attendees in 2019. ICLR's 2019 attendance was over 15 times that of 2014.

This increase is likely a result of a greater focus on deep and reinforcement learning within AI today. See Appendix for data and methodology. Note that KR takes place every second year, so there was no KR in 2019. From 2020 KR will be an annual event.

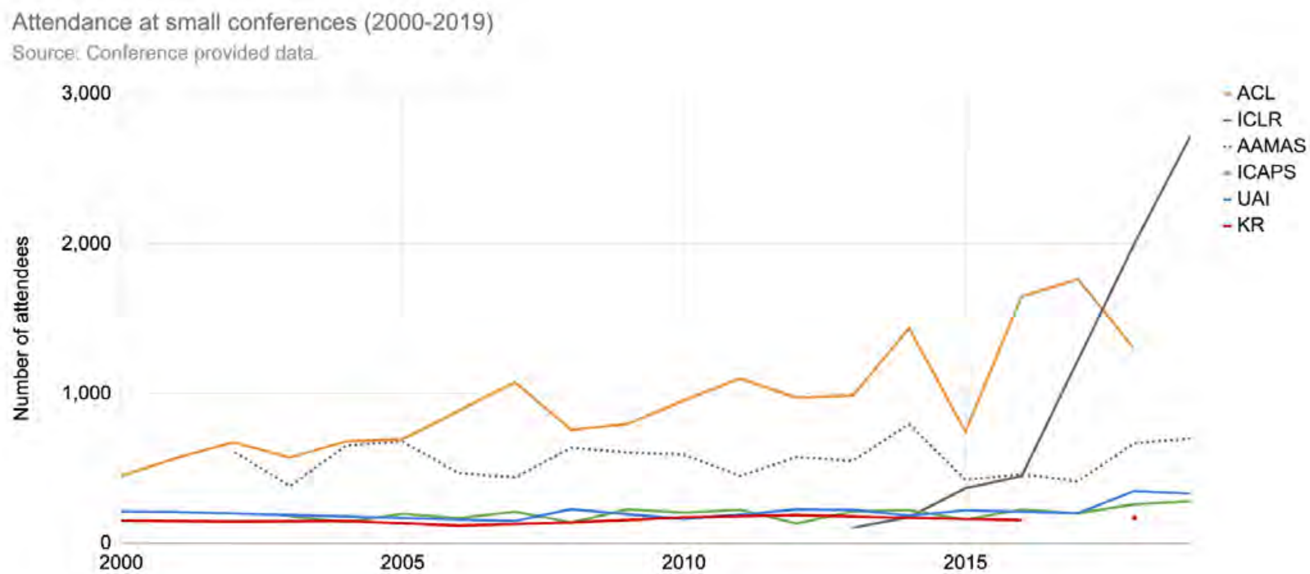

Figure 2.2a

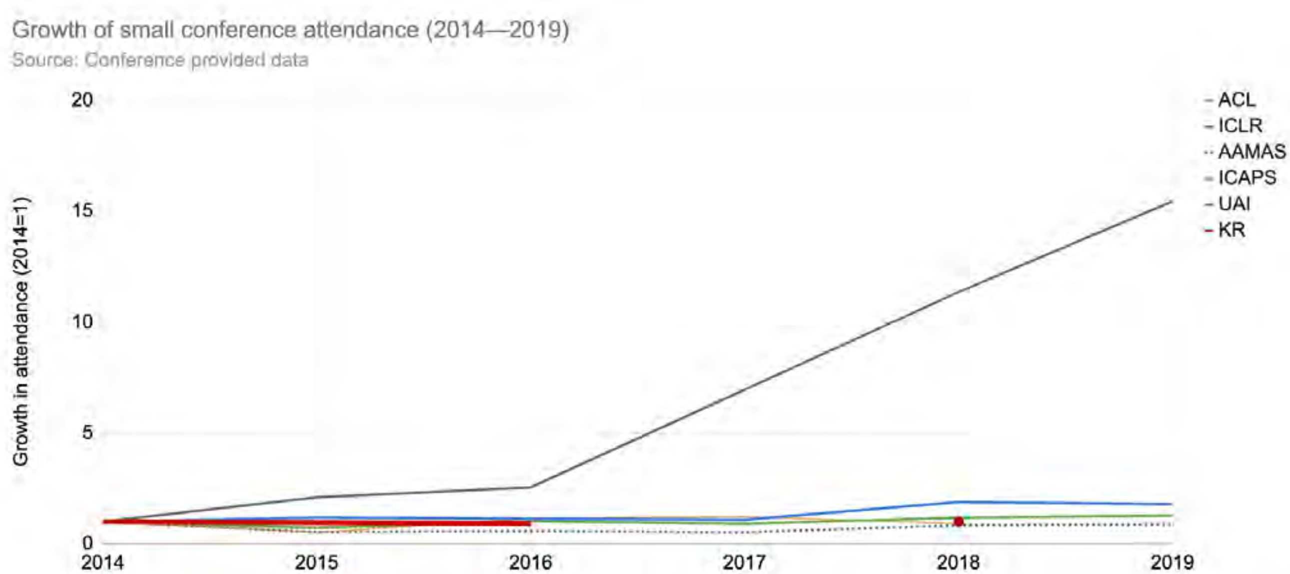

Figure 2.2b

AAAI Paper Statistics

Paper statistics are presented from AAAI - one of the longest running AI conferences that provides a broad coverage of AI topics.[3] The graph below (Figure 2.3a) shows the number of submitted and accepted papers for the 2019 Association for the Advancement of Artificial Intelligence (AAAI) conference, by country. Only countries with more than 10 accepted papers are presented. China had the largest number of submitted and accepted papers. Over 68% of submissions were from student first authors. Israel had the highest acceptance rate (24%), followed by Germany (23%), Canada (22%), the US and Singapore (both 20%).

The next graph (Figure 2.3b) shows the number of submitted and accepted papers by subject areas. Machine Learning, NLP, and Vision remain the top three subject areas. The top three subject areas with submission increase from the previous year were Reasoning Under Uncertainty (194%), Applications (176%), Humans and AI (161%). The top three subject areas with submission decreases were Cognitive Systems (-56%), Computational Sustainability (-34%) and Human Computation and Crowdsourcing (+0.9%). Acceptance rate was highest for Game Theory and Economic Paradigms (32.3%), followed by Heuristic Search (27.5%), Cognitive Systems (27.2%).

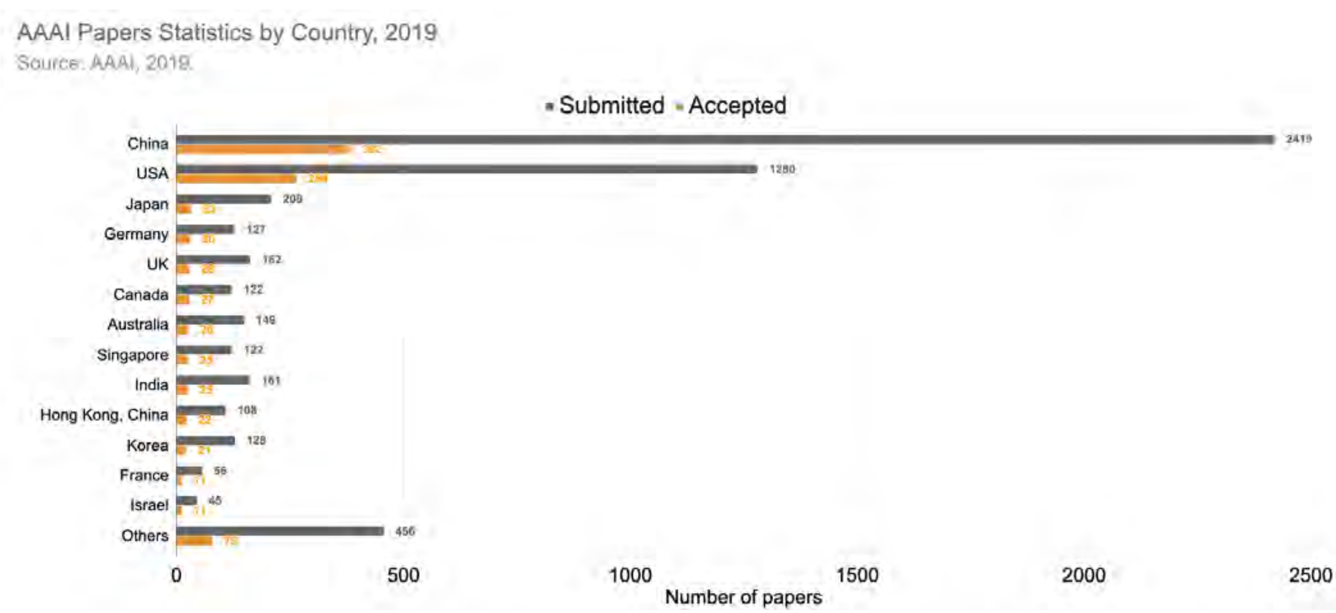

Fig. 2.3a.

Over 68% of submissions were from student first authors. Israel had the highest acceptance (24%), followed by Germany (23%), Canada (22%), the US and Singapore (both 20%).

[3] In the future, AI Index seeks to perform detailed analyses of multiple conferences.

AAAI Paper Statistics

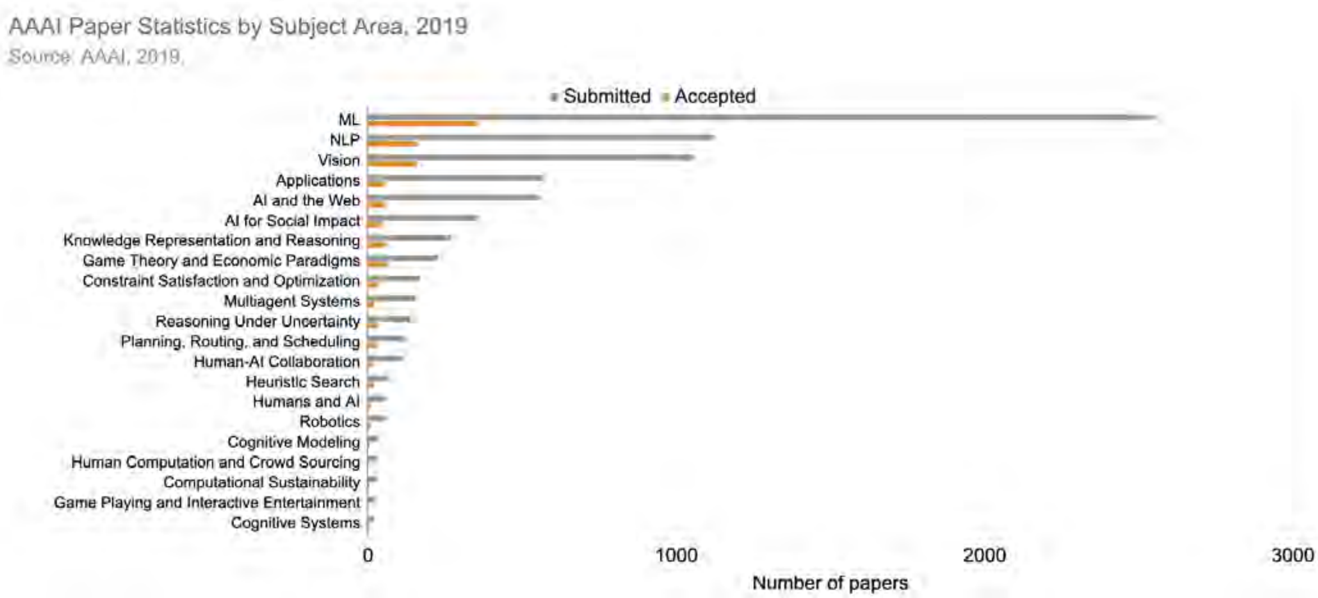

Fig. 2.3b.

Diversity Organizations

The graphs (Figures 2.4a & 2.4b) show the number of registrations for the annual workshop hosted by Women in Machine Learning (WiML), an organization dedicated to supporting women in machine learning, and the number of alumni of AI4All, an AI education initiative designed to increase diversity and inclusion in AI. Both the WiML workshop and AI4All increased program enrollment over the past several years.[4] The WiML workshop has 738% more participants than it had in 2014 and AI4ALL has 2000% more alumni than it had in 2015. These increases reflect a continued effort to include women and underrepresented groups in the AI field.

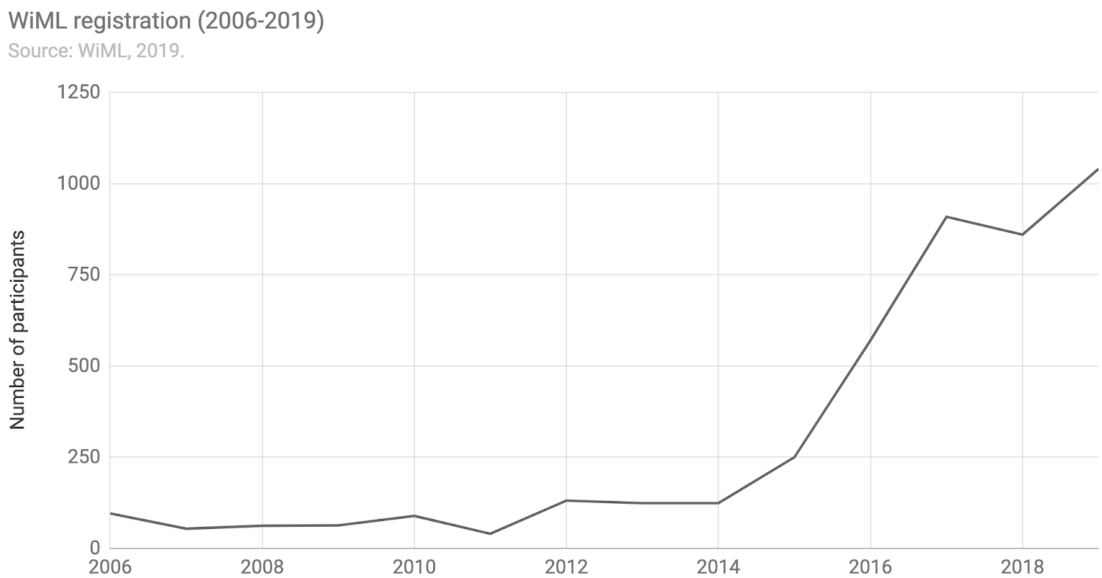

Fig. 2.4a.
Note: WiML workshop registration was slightly inflated in 2017 due to 2-day workshop, rather than 1-day format in other years.

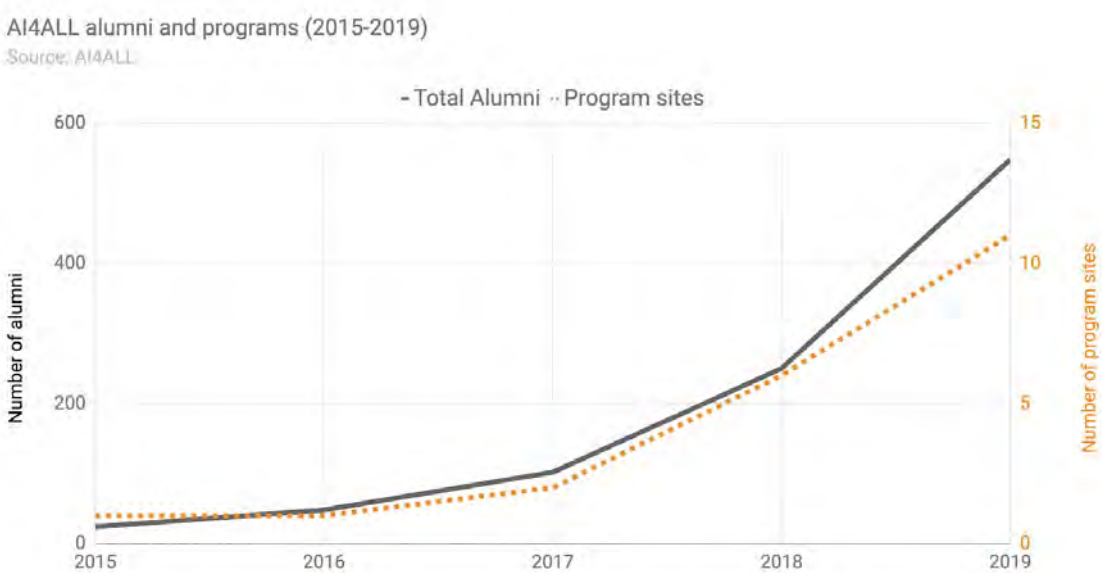

Fig. 2.4b.

[4]In October, 2019 AI4ALL launched a new program called AI4ALL Open Learning. Through the program, teachers and community-based organizations can get access to free, project-based AI curriculum for high school students--no computer science or AI experience is required for students or facilitators. This program is slated to reach over 750 high school students through AI4ALL education partners and other students using the platform by the end of 2019.

Ethics at AI Conferences

To measure Ethics in AI discussions, ethics-related terms are searched for in the titles of papers in flagship AI, machine learning, and robotics conferences and journals. The following statistics were computed on a dataset of a total of 110,108 papers, encompassing 59,352 conference and 50,756 journal entries. The details of conference and publication venue sare provided in the [Technical Appendix](). The total number of papers with ethics related keywords is a small fraction of total papers but is rising fast (Figure 2.5a). The percentage for each category (classical / trending / ethics) is based on the share of papers for which the title (or abstract, in the case of the AAAI and NIPS figures) contains at least one keyword match (Figure 2.5b). The percentages do not necessarily add up to 100% (i.e. classical / trending / ethics are not mutually exclusive). One can have a paper with matches on all three categories. See Appendix for more details.

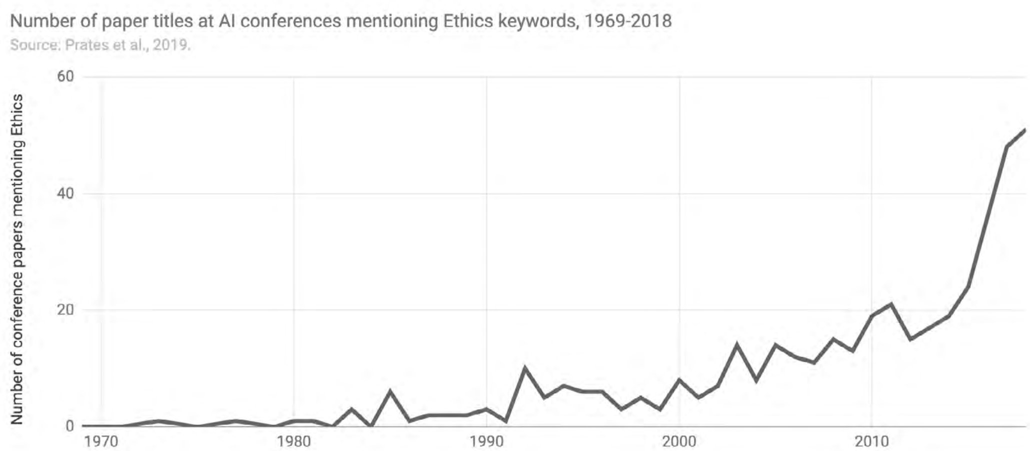

Fig. 2.5a.

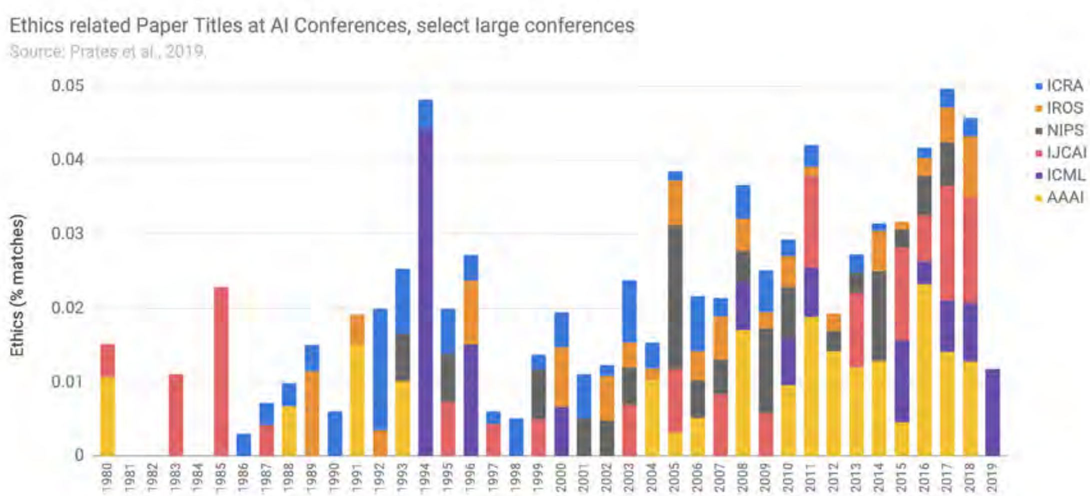

Fig. 2.5b.

"Given the implications of AI and ML in the workforce and society as a whole, the main purpose of this work was to measure the space devoted to ethics in leading AI, ML and Robotics flagship venues by means of a corpus-based approach. The findings suggest that although ethics is a growing trend among AI researchers, it is still substantially overshadowed by other technical topics in the field's flagship venues."
Marcelo Prates, Pedro Avelar, and Luis Lamb
Federal University of Rio Grande do Sul, Porto Alegre, Brazil

Human Rights and AI

RightsCon is one of the world's largest annual summits on human rights in the digital age. The attendance of RightsCon for different years is presented on the left axis (Figure 2.6) and right axis shows the number of AI sessions. 2017 was the first year artificial intelligence appeared as a stand alone track on the program (see chart below for session quantity and percentage of program from 2017 to 2019). Over time, the focus of the artificial intelligence theme has expanded from algorithmic accountability and human rights-based approaches to AI, to include conversations on algorithmic bias and discrimination; privacy and data rights; and the role of AI in the context of governance and elections, censorship and content moderation, and trade and labor. All sessions specifically on and related to artificial intelligence are available here and 2019 here.

Relevant outcomes related to AI in Toronto (2018) and Tunis (2019)

RightsCon Toronto (2018) the Toronto Declaration: Protecting the rights to equality and non-discrimination in machine learning systems was launched by Access Now and Amnesty International.

RightsCon Toronto (2018) Integrate.ai, an artificial intelligence firm launched the first draft of their white paper on Responsible AI in Consumer Enterprise, which provided a framework for organizations to operationalize ethics, privacy and security in the application of machine learning and AI.

RightsCon Tunis (2019) introduced new session format - Solve My Problem - structured to solve specific, defined problems at the intersection of human rights and technology.

RightsCon Tunis (2019) released the shared the RightsCon Community Learnings, providing direction on all tracks covered, including a specific statement on artificial intelligence.

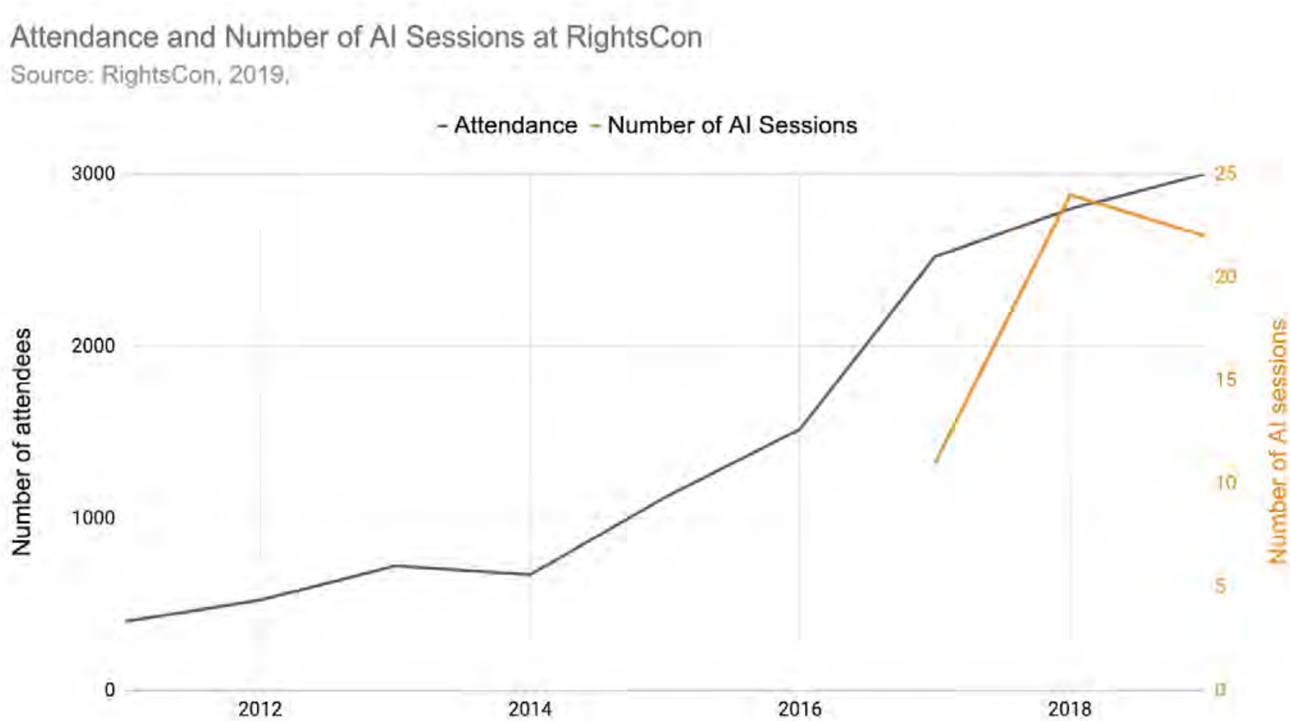

Fig. 2.6.

Measurement Questions

- How can conferences work together to facilitate comparing submissions from one conference to another; for instance, how could we compare a rise in ethics-focused papers at AAAI to a similar rise at CVPR?
- How can conferences enable better tracking of representation within the AI field at large?

Chapter Preview

Chapter 3: Technical Performance

[Table_of_Contents]
[Technical_Performance_Technical_Appendix]

Introduction

The technical performance chapter tracks technical progress in tasks across Computer Vision (Images, Videos, and Image+Language), Natural Language, potential limitations (Omniglot Challenge), and trends in computational capabilities.

Image Classification: ImageNet

ImageNet is a public image dataset of over 14 million images, created by Fei-Fei Li and her collaborators in 2009, to address the issue of scarcity of training data in the field of computer vision. The dataset, and an accompanying yearly competition (ImageNet Large Scale Visual Recognition Challenge, or ILSVRC), have been important catalysts to the developments of computer vision over the last 10 years. It was a 2012 submission to ILSVRC by Krizhevsky et al. that lead to a revival of interest in convolutional neural networks and deep learning.

The database is organized according to the WordNet hierarchy, with images depicting both higher- ("animal") and lower-level concepts ("cat"). A key computer vision task that is studied with this dataset is image classification, where an algorithm must infer whether any of the 1000 object categories of interest is present in the image.

The graph below shows accuracy scores for image classification on the ImageNet dataset over time, which can be viewed as a proxy for broader progress in supervised learning for image recognition.

ImageNet performance is being tracked by looking at scores on the validation set from the ImageNet 2012 dataset reported in published papers. The appendix documents variants of evaluation metrics to assess performance on ImageNet. The graph (Figure 3.1) shows ImageNet performance of the best performing models models trained on the ImageNet Competition training data only (grey points). The first method surpassing human performance[5] was published in 2015, and the ImageNet challenge discontinued in 2017. The dataset continues to be an important benchmark for new computer vision models, and gradual improvements continue to be reported. Three of the most recently published successful methods on this task used additional data for training - they are included as a separate plot on this graph (orange points).

Alternatively, the appendix also shows the performance improvement based on Top-5 accuracy (which evaluates a prediction as successful if the 5 top predictions returned by the model included the correct classification).

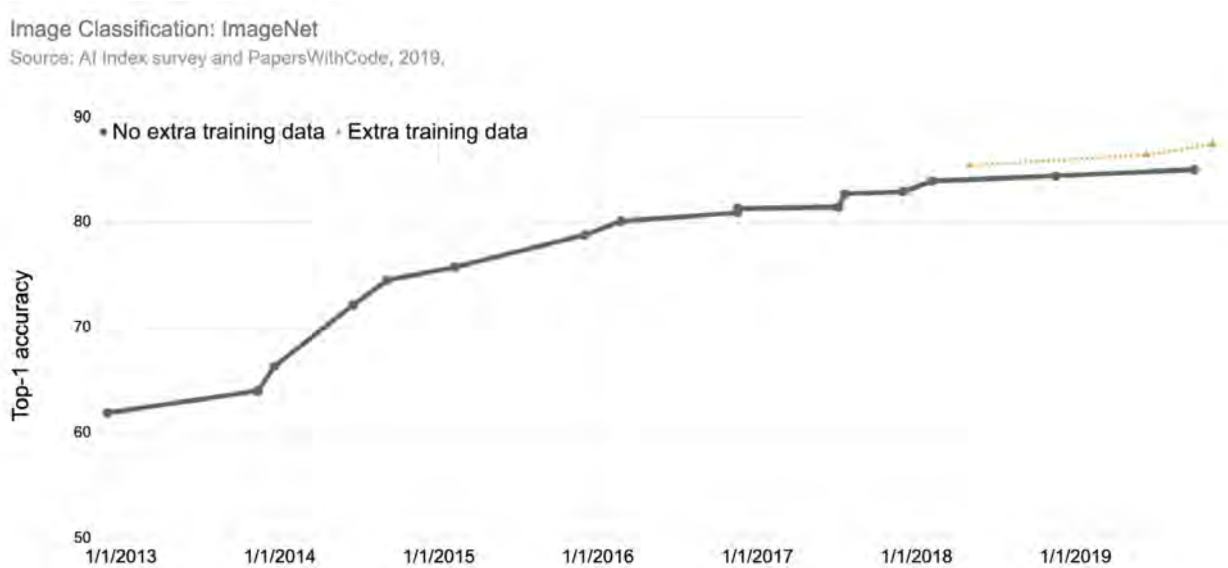

Fig. 3.1.

[5] Note: human performance here is represented by a single person annotating images. It is not representative of "human performance" for a large population.

Image Classification: ImageNet Training Time and Cost

Training Time on Public Clouds

State-of-the-art image classification methods are largely based on supervised machine learning techniques. Measuring how long it takes to train a model and associated costs is important because it is a measurement of the maturity of AI development infrastructure, reflecting advances in software and hardware.

The graph (Figure 3.2a) below shows the time required to train an image classification model to a top-5 validation accuracy of 93% or greater on ImageNet corpora when using public cloud infrastructure. This data is from Stanford's "DAWNBench" project; the data reflects the time it takes well-resourced actors in the AI field to train systems to categorize images. Improvements here give an indication of how rapidly AI developers can re-train networks to account for new data - a critical capability when seeking to develop services, systems, and products that can be updated with new data in response to changes in the world. In a year and a half, the time required to train a network on cloud infrastructure for supervised image recognition has fallen from about three hours in October 2017 to about 88 seconds in July, 2019. Data on ImageNet training time on private cloud instances shows a similar trend (see Appendix).

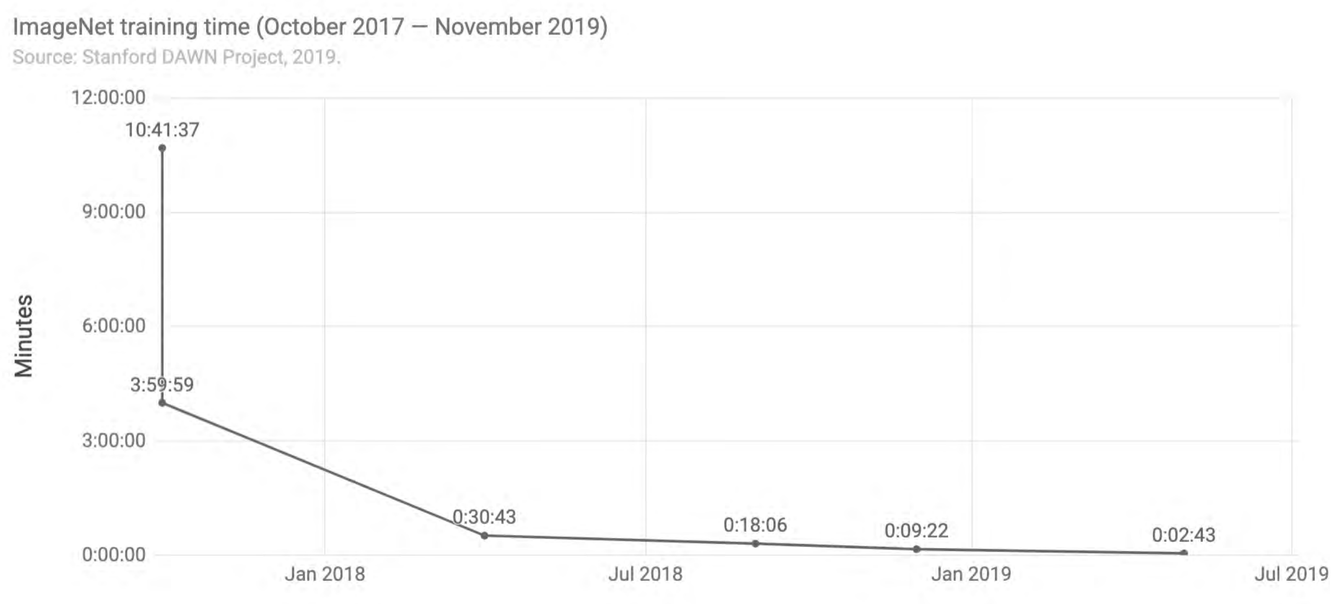

Fig. 3.2a.
Note: DAWNBench will migrate to MLperf. The latest point estimate (not shown) from ML Perf is from July, 2019 at 1 minute and 28 seconds uses Top-1 accuracy versus Top-5 accuracy benchmark shown in the graph above.

In a year and a half, the time required to train a network on cloud infrastructure has fallen from about three hours in October 2017 to about 88 seconds in July, 2019.

Image Classification: ImageNet Training Time and Cost

The next graph shows the training cost as measured by the cost of public cloud instances to train an image classification model to a top-5 validation accuracy of 93% or greater on ImageNet (Figure 3.2b). The first benchmark was a ResNet model that required over 13 days of training time to reach just above 93% accuracy that cost over $2,323 in October, 2017 (see DAWNbench submissions). The latest benchmark available on Stanford DAWNBench with lowest cost was a ResNet model run on GCP cluster with cloud TPU also reaching slightly above 93% accuracy cost slightly over $12 in September, 2018.

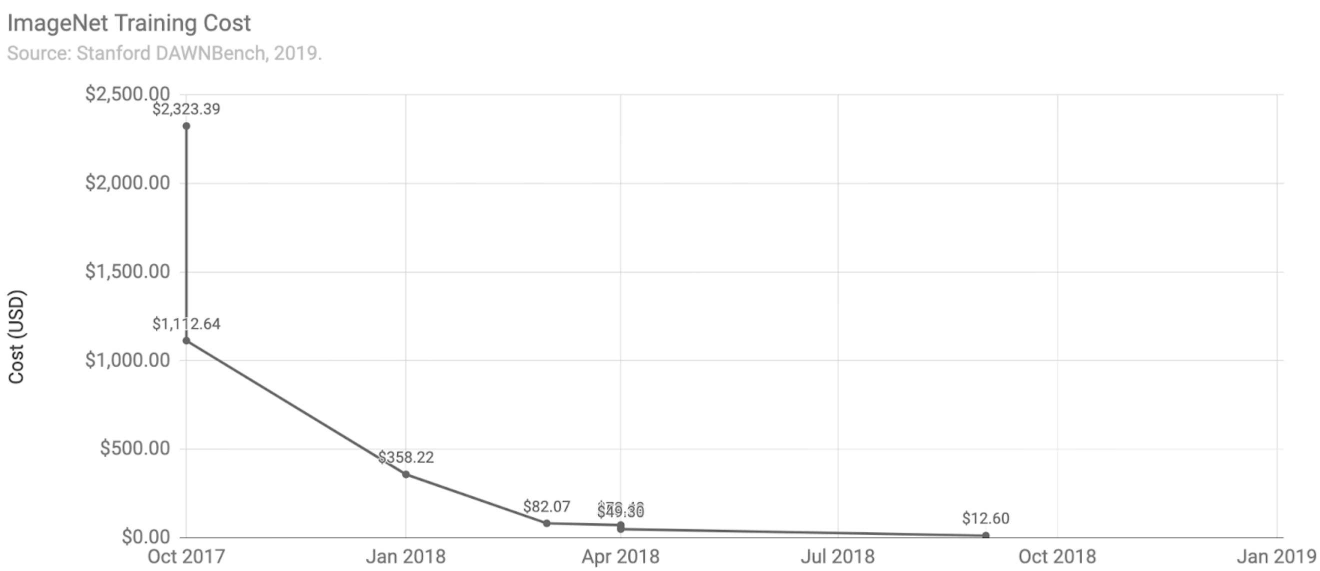

Fig. 3.2b.

Image Generation: CIFAR-10

Image generation has received attention from artists to the general public and policymakers alike. Image generation (synthesis) relies on AI models generating an output image that is meant to approximate (not necessarily replicate) the data distribution the model was trained on. Progress in image generation can be taken as a proxy for the evolution of AI models' ability to generate content in a variety of domains, ranging from images to video to text. However, assessing progress here is difficult, as beyond a certain level of realism, the quality of an image is subjective. In lieu of large-scale qualitative studies, researchers have begun using a metric called FID, which calculates the distance between the feature vectors; using the Inception v3 image model, activations are calculated on real and generated images, then the distance between these activations is calculated, giving a sense of similarity between these two groups of images. When evaluating FID, a lower score tends to correlate with images that better map their underlying data distribution and is therefore a proxy for image quality. (Figure 3.3).[6] Inception score is also reported (see Appendix Graph).

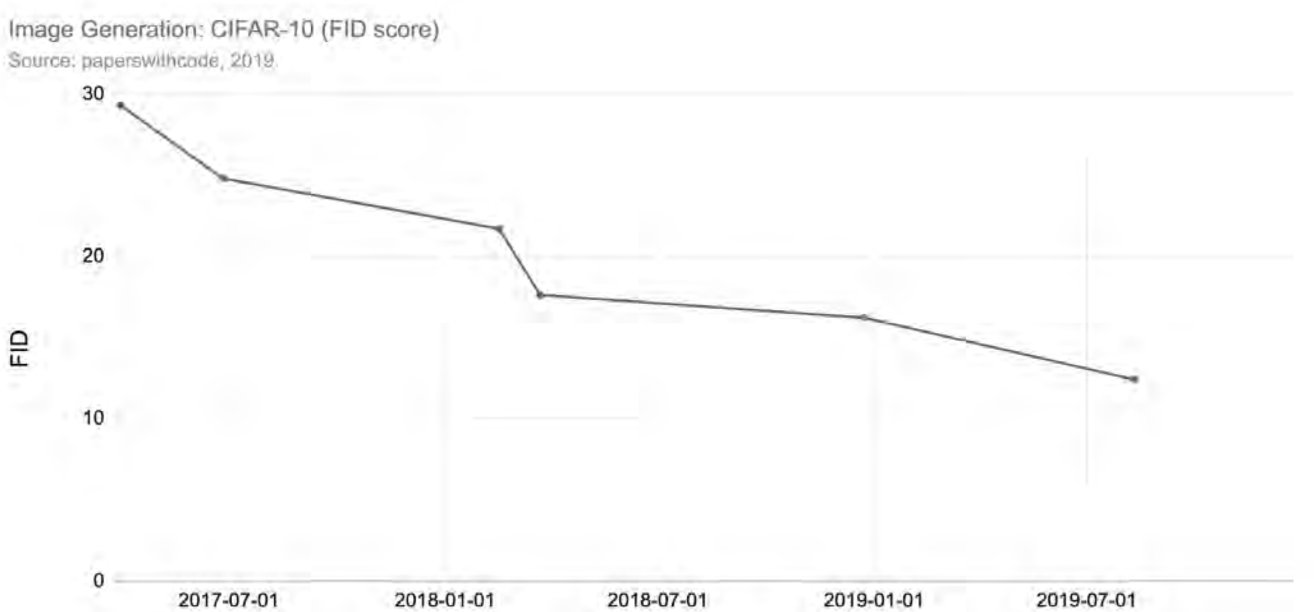

Fig. 3.3.

[1]The inception score is an attempt to remove the subjective human evaluation of images and uses a pre-trained deep learning neural network model for image classification to classify the generated images.

Semantic Segmentation

While image classification can produce a list of objects in the image, many applications require more detailed knowledge of the image contents. For instance, a robot or self-driving car may require to detect the precise boundaries and object categories for all pixels within the image. This corresponds to the task of semantic segmentation, where the algorithm must divide the image into regions and classify each region into one of the categories of interest, producing a pixel-level map of the image contents.

Progress in semantic segmentation is an input to progress in real-world AI vision systems, such as those being developed for self-driving cars. Progress is measured in this domain using the mean intersection over union (IoU) metric on two datasets: Cityscapes (Figure 3.4). Some systems were trained with extra data. See Appendix for details on individual datasets and progress in PASCAL Context

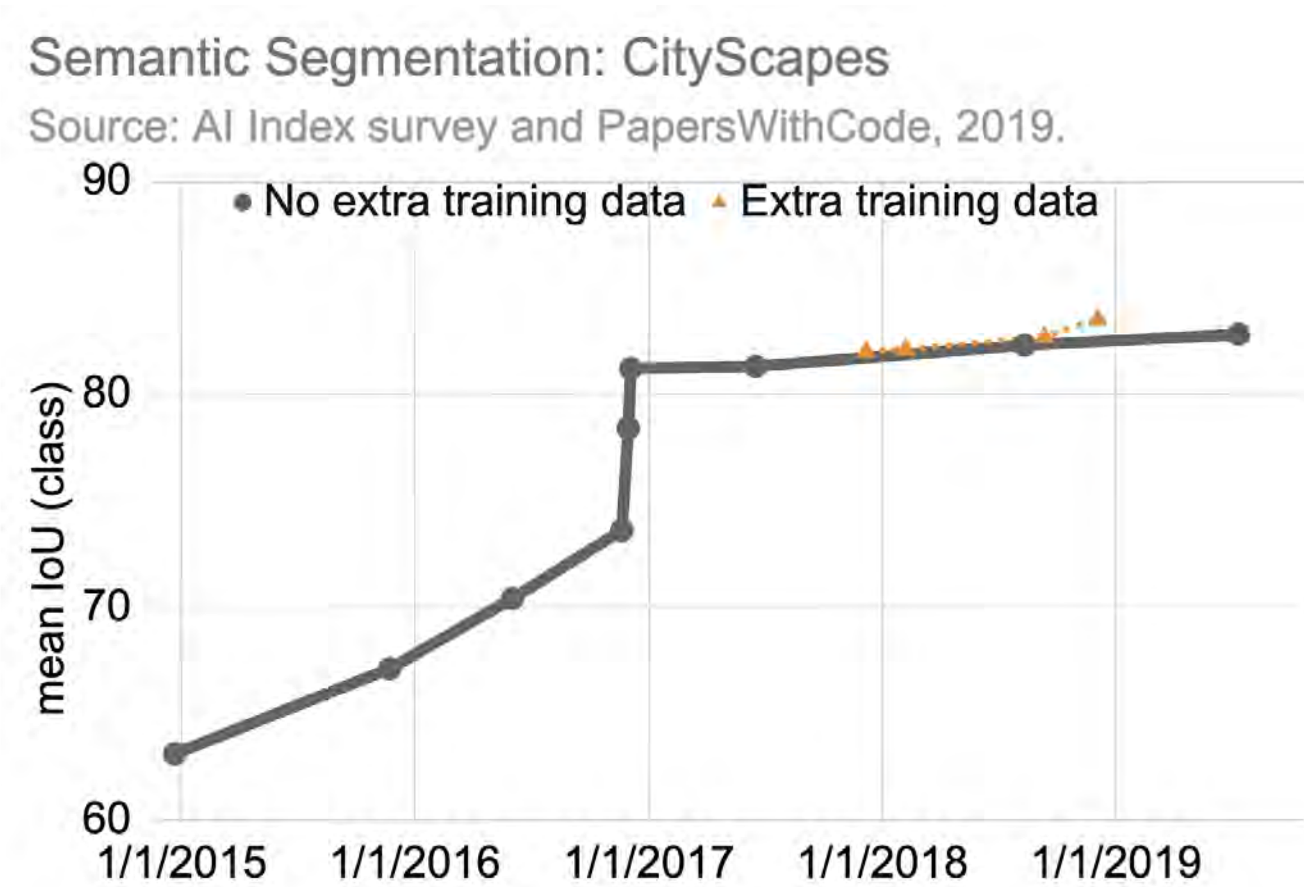

Fig. 3.4.
Note: The orange dots denote tests with additional training data.

ActivityNet

In addition to image analysis, algorithms for understanding and analyzing videos are an important focus in the computer vision research community. Particularly, algorithms that can recognize human actions and activities from videos would enable many important applications. Further discussion of progress in activity recognition in videos appears in the ActivityNet Challenge.

A key task in the ActivityNet Challenge is that of Temporal Activity Localization. In this task, algorithms are given long video sequences that depict more than one activity, and each activity is performed in a sub-interval of the video but not during its entire duration. Algorithms are then evaluated on how precisely they can temporally localize each activity within the video as well as how accurately they can classify the interval into the correct activity category.

ActivityNet has compiled several attributes for the task of temporal localization at the challenge over the last four rounds. Below detailed analysis and trends for this task are presented (e.g. how has the performance for individual activity classes improved over the years (Figure 3.5a)? Which are the hardest and easiest classes now (Figure 3.5b & 3.5c)? Which classes have the leastmost improvement over the years (figure 3.5d)? The ActivityNet statistics are available **here**.

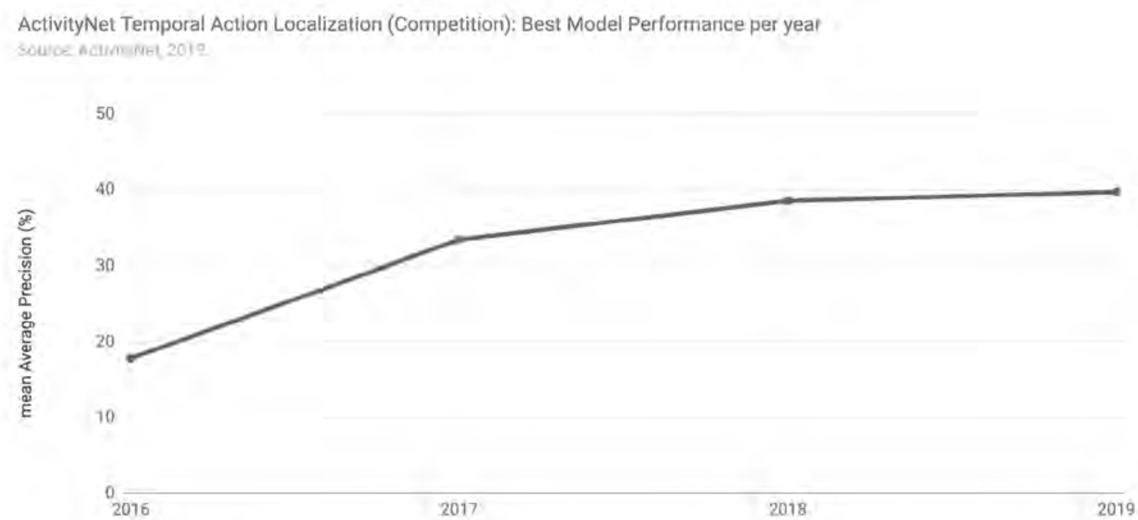

Fig. 3.5a.

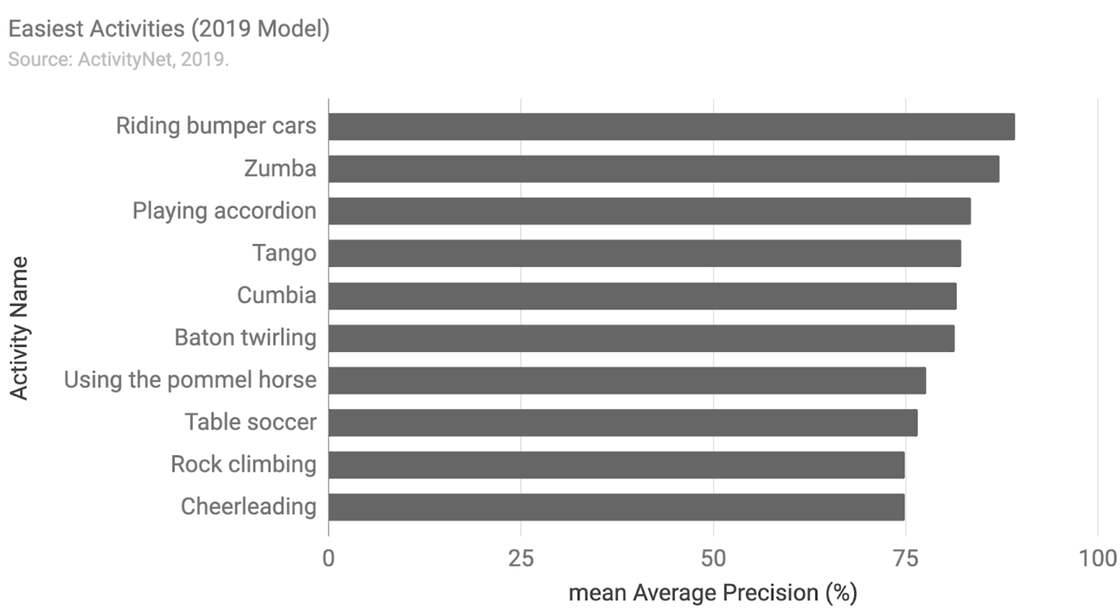

Fig. 3.5b.

Activity Recognition in Videos

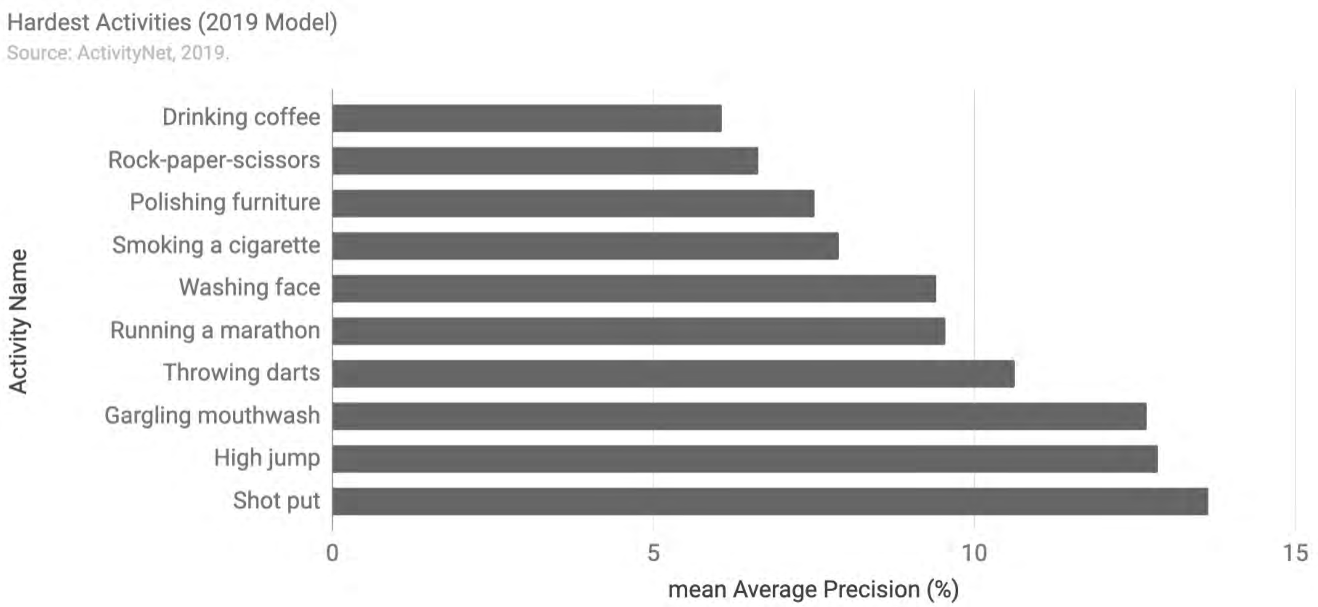

Fig. 3.5c.

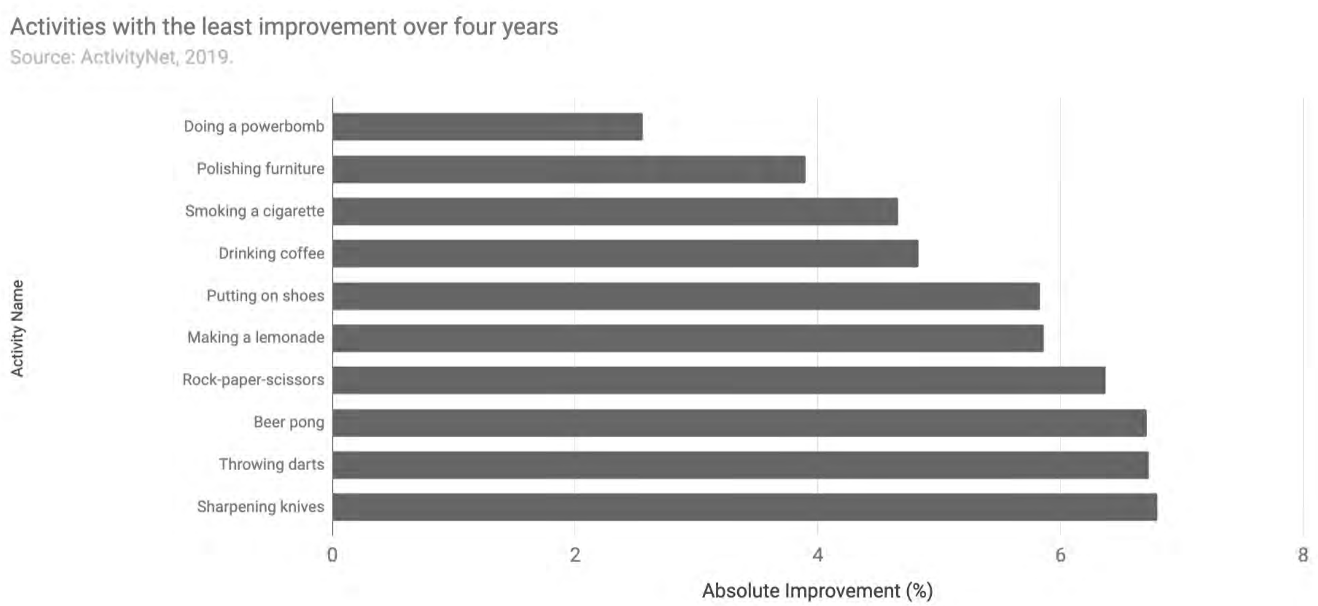

Fig. 3.5d.

"The emergence of large-scale datasets such as ActivityNet and Kinetics has equipped computer vision researchers with valuable data and benchmarks to train and develop innovative algorithms that push the limits of automatic activity understanding. These algorithms can now accurately recognize hundreds of complex human activities such as bowling or sailing, and they do so in real-time. However, after organizing the International Activity Recognition Challenge (ActivityNet) for the last four years, we observe that more research is needed to develop methods that can reliably discriminate activities, which involve fine-grained motions and/or subtle patterns in motion cues, objects, and human-object interactions. Looking forward, we foresee the next generation of algorithms to be one that accentuates learning without the need for excessively large manually curated data. In this scenario, benchmarks and competitions will remain a cornerstone to track progress in this self-learning domain."
Bernard Ghanem, Associate Professor of Electrical Engineering
King Abdullah University of Science and Technology

Visual-Question Answering (VQA)

The VQA challenge incorporates both computer vision and natural language understanding. The VQA challenge tests how well computers can jointly reason over these two distinct data distributions. The VQA challenge uses a dataset containing open-ended questions about the contents of images. Successfully answering these questions requires an understanding of vision, language and commonsense knowledge. In 2019, the overall accuracy grew by +2.85% to 75.28% (Figure 3.6). The 2019 VQA challenge had 41 teams representing more than 34 institutions and 11 countries. Reader refer to the [VQA challenge website] and [Appendix] for more details.

Can you beat the VQA challenge?
To get a sense of the challenge, you can try online VQA demos out at https://vqa.cloudcv.org/. Upload an image, ask the model a question, and see what it does.

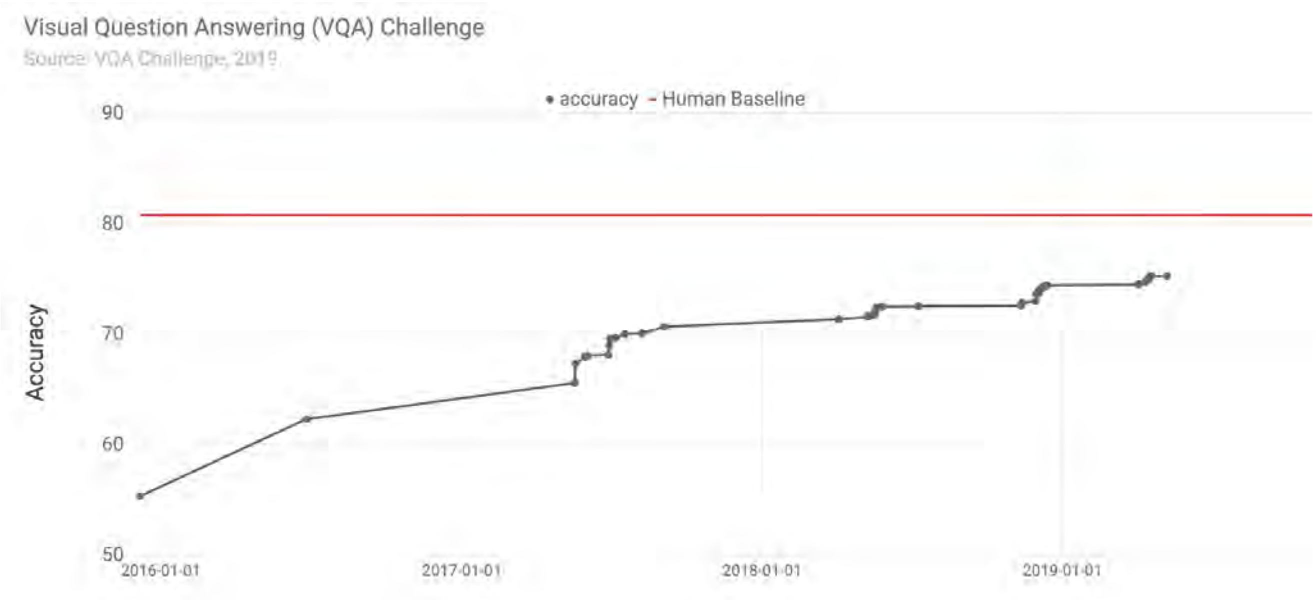

Fig. 3.6.
Note: Human performance is measured by having humans answer questions for images and evaluating their answers using the same metrics as we use to evaluate machines that answer the same questions. Inter-human disagreement, paraphrased answers, spelling errors, etc, contribute to human performance being (quite a bit lower) than 100%.

What explains progress in this domain? *"There's been no silver bullet. Progress has been the consequence of open exploratory research and consistent iterations by researchers in the community -- the vision and language community, the vision community, and the language community. As a community we identified effective multimodal fusion techniques, image representations that are more appropriate for tasks that link to language, convolutional neural network architectures for improved perception, pre-training mechanisms to learn language representations that can be transferred to other tasks."*
Devi Parikh
Georgia Tech | Facebook AI Research (FAIR)

GLUE

Being able to analyze text is a crucial, multi-purpose AI capability. In recent years, progress in natural language processing and natural language understanding has caused the AI community to develop new, harder tests for AI capabilities. In the language domain, a good example is GLUE, the General Language Understanding Evaluation benchmark. GLUE tests single AI systems on nine distinct tasks in an attempt to measure the general text-processing performance of AI systems. GLUE consists of nine sub-tasks — two on single sentences (measuring linguistic acceptability and sentiment), three on similarity and paraphrase, and four on natural language inference, including the Winograd Schema Challenge. As an illustration of the pace of progress in this domain, though the benchmark was only released in May 2018, performance of submitted systems crossed non-expert human performance in June, 2019. Performance has continued to improve in 2019 (Figure 3.7) with models like RoBERTa from Facebook and T5 from Google. More details on GLUE tasks with greater (or shorter) distance to human performance frontier are available (see [Appendix Graph]).

Fig. 3.7.

SuperGLUE

Progress in language-oriented AI systems has been so dramatic that the creators of the GLUE benchmark needed to create a new, more challenging benchmark, so they could test performance after some systems surpassed human performance on GLUE. SuperGLUE contains a new set of more diverse and difficult language understanding tasks, improved resources, and a new public leaderboard.

Within five months of its launch in May, 2019, the T5 model published by Google almost reached human baseline of 89.9 with their at the score of 88.9 (Figure 3.8). This was achieved using a task-agnostic text-to-text framework that utilized an encoder-decoder architecture. The model was pre-trained on a mixture of NLP tasks and fine-tuned on SuperGLUE.

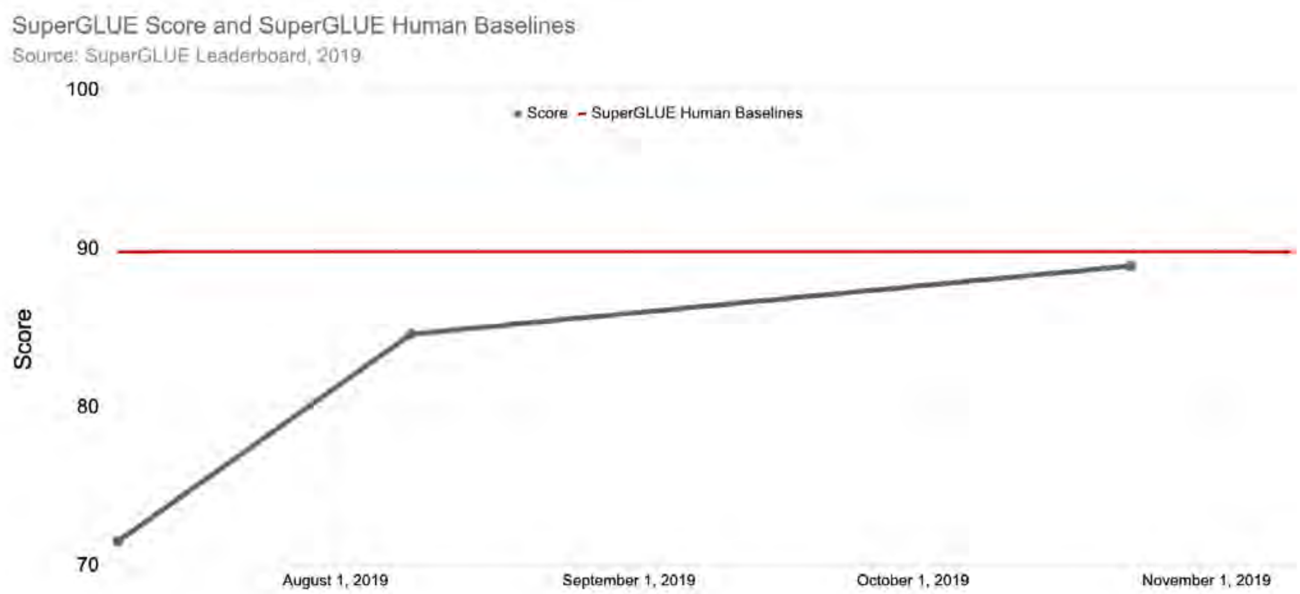

Fig. 3.8.
Notes: Human baseline was estimated by hiring crowdworker annotators through Amazon's Mechanical Turk platform to reannotate a sample of each test set to estimate. More details can be found here.

Since being launched in May, 2019, the T5 Team at Google has almost reached human baseline at the score of 88.9 within five months on SuperGLUE. Human baseline is 89.8.

GLUE and superGLUE

What does progress in natural language understanding mean?

What is the best way to interpret the rapid progress in natural language and what might measures like GLUE and SuperGLUE tell us about progress in this domain? Sam Bowman, an assistant professor at NYU whose group has developed GLUE and SuperGLUE offers:

"We know now how to solve an overwhelming majority of the sentence- or pararaph-level text classification benchmark datasets that we've been able to come up with to date. GLUE and SuperGLUE demonstrate this out nicely, and you can see similar trends across the field of NLP. I don't think we have been in a position even remotely like this before: We're solving hard, AI-oriented challenge tasks just about as fast as we can dream them up," Sam says. "I want to emphasize, though, that we haven't solved language understanding yet in any satisfying way."

While GLUE and SuperGLUE may indicate progress in the field, it is important to remember that successful models could be exploiting statistical patterns in their underlying datasets, are likely to display harmful biases, and when they demonstrate better-than-human performance, they may be doing this unevenly, displaying good performance on some tasks and faulty or inhuman reasoning on others.

"This leaves us in an odd position," Bowman says. "Especially for these classification-style tasks, we see clear weaknesses with current methods, but we don't yet have clear, fair ways to quantify those weaknesses. I'm seeing what looks like a new surge of interest in data collection methods and evaluation metrics, and I think that's a healthy thing for us to be focusing on."

Human Expectations for the SuperGLUE Benchmark

The AI Index has partnered with Metaculus, a crowd forecasting initiative, to source 'crowd predictions' from the general public for the 2019 report. The question went public on August 9, 2019 and will close on closes Dec 30, 2019. Respondents don't predict "yes" or "no," but rather the percent likelihood. At the time of writing this, there were 127 human predictions. Metaculus users were asked the following question:

By May 2020, will a single language model obtain an average score equal to or greater than 90% on the SuperGLUE benchmark?

Results: The median prediction of respondents is a 90% likelihood that a single model will obtain an average score equal to or greater than 90% on the SuperGLUE benchmark.

[Technical_Performance_Technical_Appendix]_[SuperGLUE_Paper][SuperGLUE_Leaderboard]
[Access_Data]

SQuAD

One way to highlight recent progress in natural language processing is to examine performance on the Stanford Question Answering Dataset (SQuAD) challenge. SQuAD is a reading comprehension dataset, consisting of questions posed by crowdworkers on a set of Wikipedia articles. The answer to every question is a segment of text, or span, from the corresponding reading passage, or the question might be unanswerable. SQuAD1.1 is the SQuAD dataset and contains 100,000+ question-answer pairs on 500+ articles. SQuAD2.0 combines the 100,000 questions in SQuAD1.1 with over 50,000 unanswerable questions written adversarially by crowdworkers to look similar to answerable ones. To do well on SQuAD2.0, systems must not only answer questions when possible, but also determine when no answer is supported by the paragraph and abstain from answering. SQuAD2.0 was developed partially because of surprising, rapid performance by entrances on the original SQuAD benchmark. The SQuAD Leaderboard and data are available. The F1 score for SQuAD1.1 went from 67 in August, 2016 to 95 in May, 2019 (Figure 3.9). Progress on SQuAD2.0 has been even faster. F1 score went from 62 in May, 2018 to 90 in June, 2019. CodaLab hosts other active NLP competitions.

The time taken to train QA model to 75 F1 score or greater on SQuAD 1.0 went down from over 7 hours in October, 2017 to less than 19 minutes in March, 2019 (Figure 3.13b). The cost to public cloud instances to train a QA model to has reduced from $8 to 57 cents by December, 2018, and inference time reduced from 638 milliseconds to 7 milliseconds (see Appendix Graph).

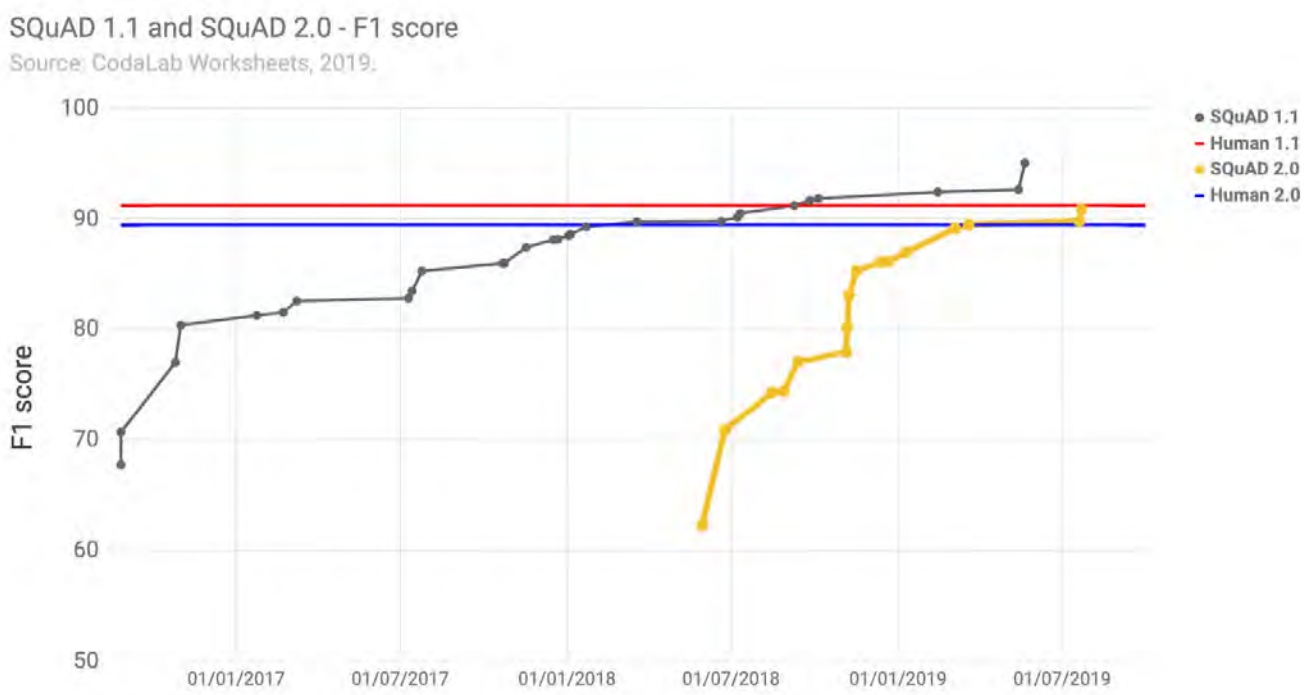

Fig. 3.9.

The F1 score for SQuAD1.1 went from 67 in August, 2016 to 95 in May, 2019. Progress on SQuAD2.0 has been even faster. F1 score went from 62 in May, 2018 to 90 in June, 2019.

Reasoning

The Allen Institute for Artificial Intelligence (AI2) has several initiatives that relate to measuring the advancing capabilities of AI systems and is home to several AI research initiatives including the AllenNLP, Aristo, and Mosaic projects. Several AI2 Leaderboards are publicly available for NLP and commonsense reasoning tasks. Performance improvements in selected tasks are presented below.

AI2 Reasoning Challenge (ARC)

Released in April 2018, the ARC dataset contains 7,787 genuine grade-school level, multiple-choice science questions. The questions are text-only, English language exam questions that span several grade levels. Each question has a multiple-choice structure (with typically four answer options). The questions are accompanied by the ARC Corpus, a collection of 14M unordered, science-related sentences including knowledge relevant to ARC. It is not guaranteed that answers to the questions can be found in the corpus. The ARC dataset is divided into a Challenge Set (2,590 questions) and an Easy Set (5,197 questions). The Challenge Set contains only questions that were answered incorrectly by both a retrieval-based algorithm and a word co-occurrence algorithm.

ARC Easy

The first graph from AI2 shows the progress on the ARC-Easy dataset, 5,197 questions that can be answered by retrieval or co-occurrence algorithms. More details about this task can be found in the Appendix. There have been 20 submissions to the ARC-Easy leaderboard, with the top score yielding 85.4% accuracy on the test set, updated on September 27, 2019 (Figure 3.10).

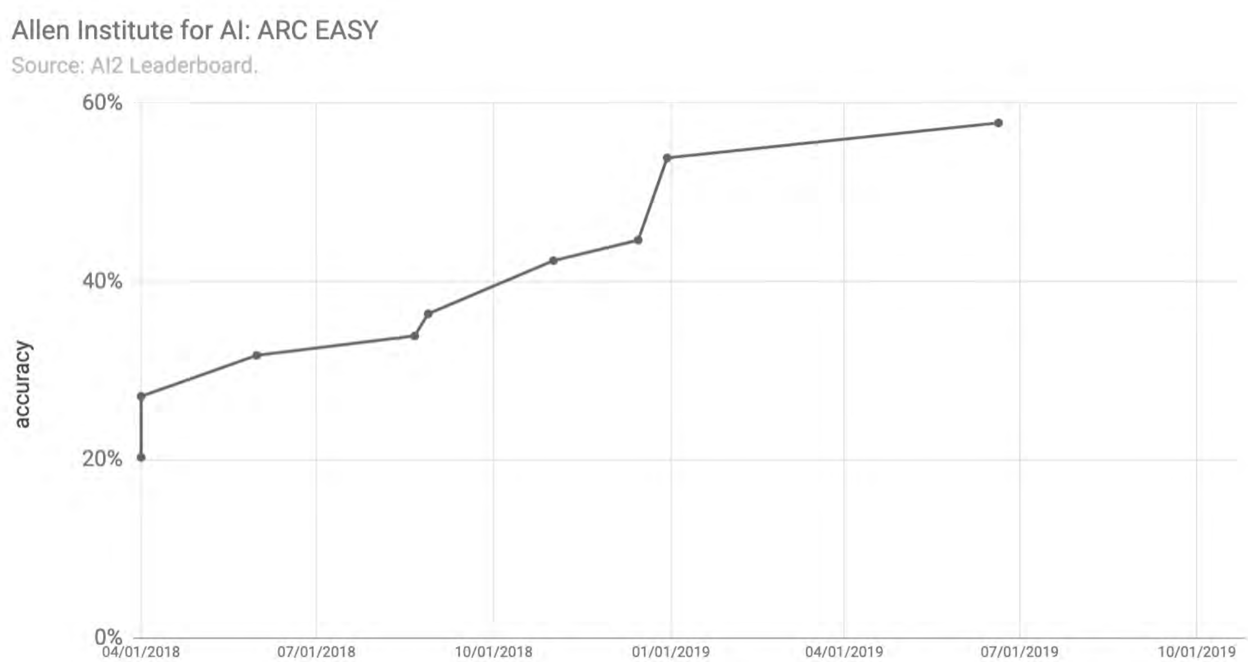

Fig. 3.10.

Reasoning

ARC Challenge Set

The graph below shows performance over time for the ARC Challenge Set. See [Appendix] for data and methodology. There have been 26 submissions to the ARC Challenge Set leaderboard with a top score of 67.7% last updated on September 27, 2019 (Figure 3.11).

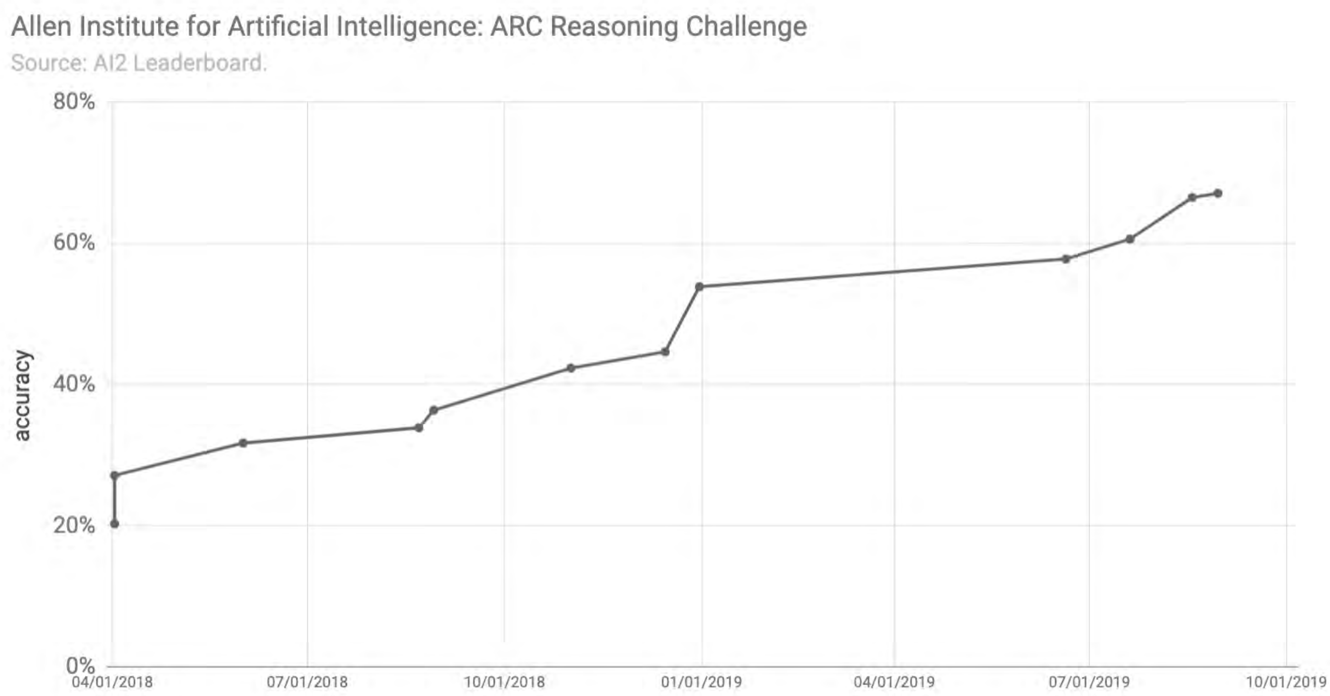

Fig. 3.11.

Commercial Machine Translation (MT)

Translation is one of the more easily applicable capabilities of contemporary language-oriented AI systems. Therefore, examining the number and performance of commercially deployed translation systems gives us a sense of how rapidly technology migrates from research to production, and of what the impact is here.

According to Intento, a startup that provides simple APIs to evaluate third-party AI models in MT from many vendors, the number of commercially available MT systems with pre-trained models and public APIs has grown rapidly, from 8 in 2017 to over 24 in 2019 (Figure 3.12a). Increasingly, MT systems provide a full range of customization options: pre-trained generic models, automatic domain adaptation to build models and better engines with their own data, and custom terminology support.

The growth in commercial MT is driven by engines that excel at their geography and business-related language pairs and domains (Germany, Japan, Korea, China). Since early 2018, the increase in commercial MT system is due to two factors: (1) existing vendors of on-premise and bespoke MT are starting to provide pre-trained models available in the cloud and (2) the technology barrier to fielding translation systems is getting lower as a consequence of more neural machine translation (NMT) frameworks being made available open-source.

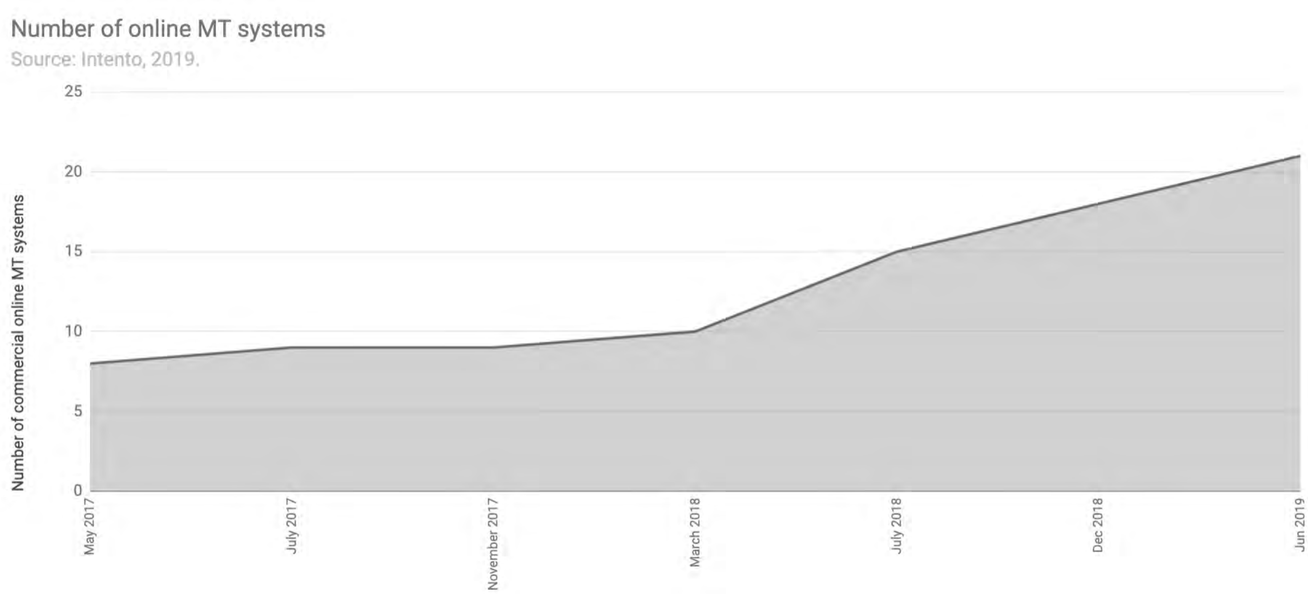

Fig. 3.12a.

Commercial Machine Translation (MT)

Commercial MT quality is evaluated quality using hLEPOR metric, which measure the difference from a human reference translation. hLEPOR scores of 0.7 means almost human-level quality with just a couple of mistakes per sentence. The hLEPOR performance score in language pairs for online systems is presented below (Figure 3.12b). To make the analysis comparable, the presentation is only for pairs including English. It is based on ranking the best online MT system for 48 language-pairs tested. Portugese-English and English-Portugese are pairs with highest hLEPOR score, followed by English to German, and Italian to English. Details on data, methodology, and replicability of results can be found in the Technical Appendix. The next chart shows the ranking of language pairs based on improvement in hLEPOR score between May, 2017 and June, 2019 (figure 3.12c). The fastest improvement was for Chinese-to-English, followed by English-to-German and Russian-to-English. Performance of the baseline models varies widely between different language pairs. The main contributing factor is language pair popularity, which defines how much investment goes into data acquisition and curation. Also, the next-generation translation technology (such as Transformer) is being rolled out to the most popular language pairs first, while rare language pairs may still employ Phrase Based Machine Translation (PBMT) models.

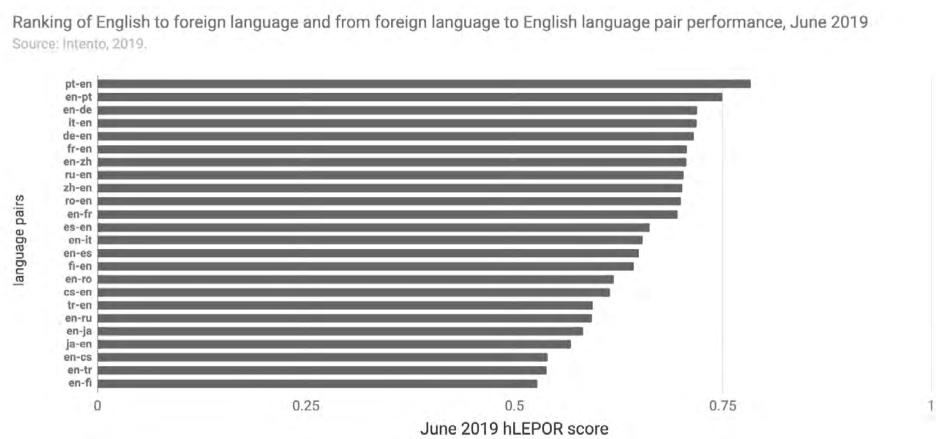

Fig. 3.12b.

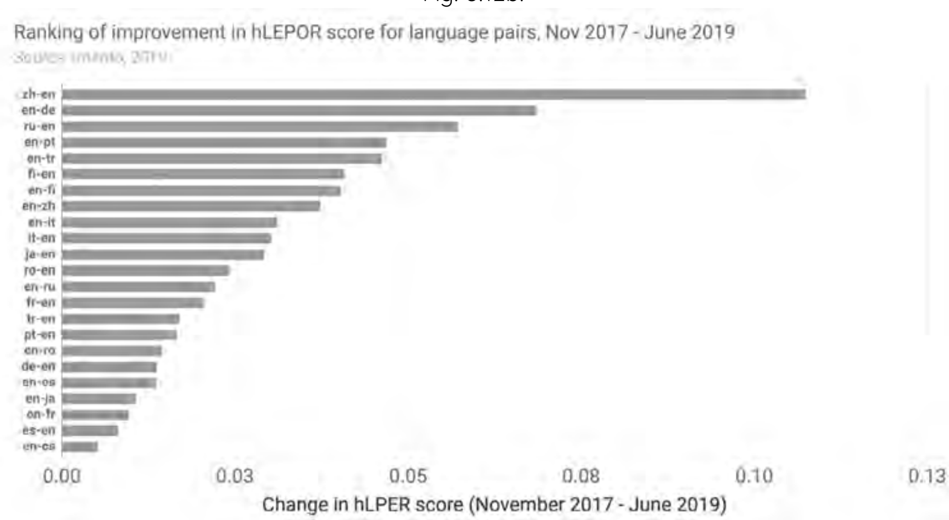

Fig. 3.12c.

"Increased data confidentiality concerns complicate data acquisition for domain-specific models. As a result, we see MT providers putting a lot of effort into building domain adaptation tools for data owners. Those are AutoML-type technology, terminology adaptation, and the ability to improve models based on end-user feedback. We expect these will be the primary technology drivers in the near term."
Konstantin Savenkov, CEO Intento, Inc.

Omniglot Challenge

There has been notable progress on one-shot classification over the last three years; however, there has been less progress on the other four concept learning tasks in the Omniglot Challenge. The Omniglot Challenge requires performing many tasks with a single model, including classification, parsing, generating new exemplars, and generating whole new concepts. Bayesian program learning (BPL) performs better than neural network approaches on the original one-shot classification challenge, despite the improving capabilities of neural network models (Figure 3.13). See tge Appendix for details on the task.

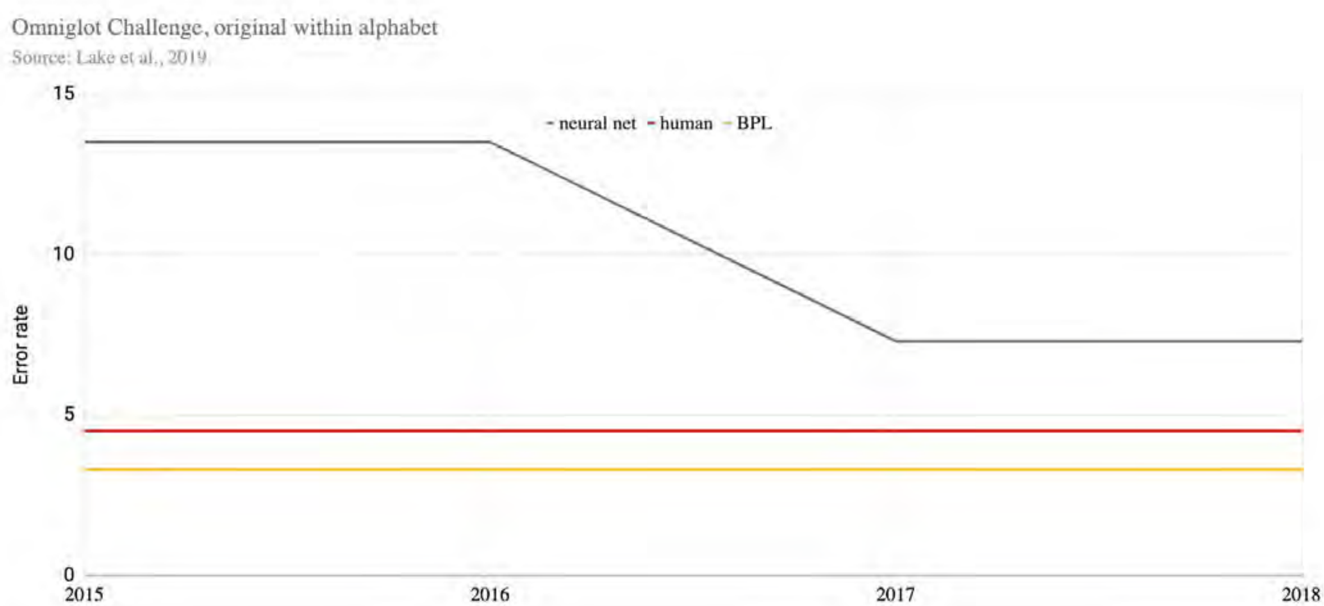

Fig. 3.16.

The Omniglot challenge: a 3-year progress report

Human-level concept learning through probabilistic program induction

"Achieving human-level concept learning will require learning richer representations from less data, and reconfiguring these representations to tackle new tasks" says Brenden Lake, an Assistant Professor at New York University and author of the Omniglot challenge and progress report. Lake further says that *"there is no official leaderboard for Omniglot, and in fact, it's difficult to define an appropriate leaderboard for the entire challenge. Progress on building machines that can learn concepts in a more human-like way cannot be boiled down to just a single number or a single task. Rather, as the progress report states, models need to be developed with a broad competence for performing a variety of different tasks using their conceptual representation."*

Computational Capacity

The amount of computation used in the largest AI training runs has doubled every 3.4 months since 2012 (net increase of 300,000x). The y-axis of the chart shows the total amount of compute, in petaflop/s-days, used to train selected results (Figure 3.14a and 3.14b). A petaflop-day (pf-day) consists of performing 10^{15} neural net operations per second for one day, or a total of about 10^{20} operations. The x-axis is the publication date. Doubling time for the line of best fit shown is 3.4 months. Based on analysis of compute used in major AI results for the past decades, a structural break with two AI eras are identified by OpenAI:

1) **Prior to 2012** - AI results closely tracked Moore's Law, with compute doubling every two years (Figure 3.14a).

2) **Post-2012** - compute has been doubling every 3.4 months (Figure 3.14b). Since 2012, this compute metric has grown by more than 300,000x (a 2-year doubling period would yield only a 7x increase).

Two methodologies were used to generate these data points. When information was available, the number of FLOPs (adds and multiplies) in the described architecture per training example were directly counted and multiplied by the total number of forward and backward passes during training. When enough information to directly count FLOPs was not available, GPU training time and total number of GPUs were used and a utilization efficiency (usually 0.33) was assumed. Technical details on calculations can be found on the OpenAI blog.

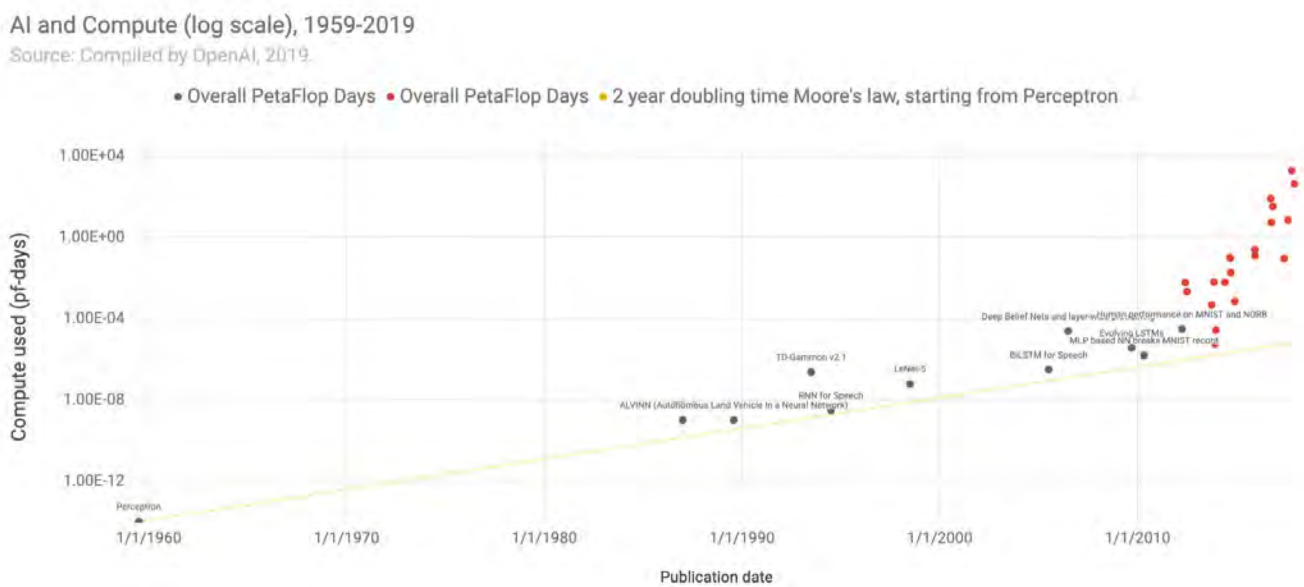

Fig. 3.14a.

Computational Capacity

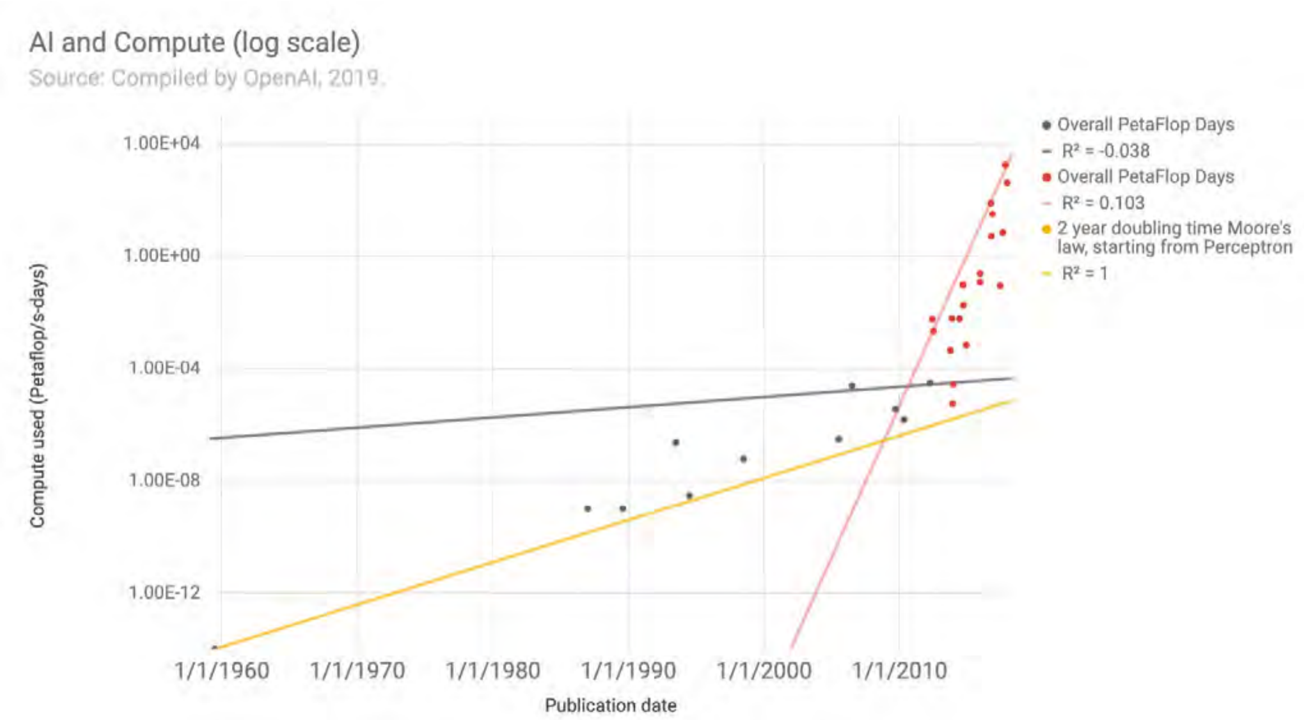

Fig. 3.14b.

Prior to 2012, AI results closely tracked Moore's Law, with compute doubling every two years. Post-2012, compute has been doubling every 3.4 months.

Human-Level Performance Milestones

The inaugural 2017 AI Index report included a timeline of circumstances where AI reached or beat human-level performance. The list outlined game playing achievements, accurate medical diagnoses, and other general, but sophisticated, human tasks that AI performed at a human or superhuman level. This year, two new achievements are added to that list. It is important not to over-interpret these results. The tasks below are highly specific, and the achievements, while impressive, say nothing about the ability of the systems to generalize to other tasks.

1980

Othello

In the 1980s Kai-Fu Lee and Sanjoy Mahajan developed BILL, a Bayesian learning-based system for playing the board game Othello. In 1989, the program won the US national tournament of computer players, and beat the highest ranked US player, Brian Rose, 56—8. In 1997, a program named Logistello won every game in a six game match against the reigning Othello world champion.

1995

Checkers

In 1952, Arthur Samuels built a series of programs that played the game of checkers and improved via self-play. However, it was not until 1995 that a checkers-playing program, Chinook, beat the world champion.

1997

Chess

Some computer scientists in the 1950s predicted that a computer would defeat the human chess champion by 1967, but it was not until 1997 that IBM's DeepBlue system beat chess champion Gary Kasparov. Today, chess programs running on smartphones can play at the grandmaster level.

2011

Jeopardy!

In 2011, the IBM Watson computer system competed on the popular quiz show Jeopardy! against former winners Brad Rutter and Ken Jennings. Watson won the first place prize of $1 million.

2015

Atari Games

In 2015, a team at Google DeepMind used a reinforcement learning system to learn how to play 49 Atari games. The system was able to achieve human-level performance in a majority of the games (e.g., Breakout), though some are still significantly out of reach (e.g., Montezuma's Revenge).

2016

Object Classification in ImageNet

In 2016, the error rate of automatic labeling of ImageNet declined from 28% in 2010 to less than 3%. Human performance is about 5%.

2016

Go

In March of 2016, the AlphaGo system developed by the Google DeepMind team beat Lee Sedol, one of the world's greatest Go players, 4—1. DeepMind then released AlphaGo Master, which defeated the top ranked player, Ke Jie, in March of 2017. In October 2017, a Nature paper detailed yet another new version, AlphaGo Zero, which beat the original AlphaGo system 100—0.

2017

Skin Cancer Classification
In a 2017 Nature article, Esteva et al. describe an AI system trained on a data set of 129,450 clinical images of 2,032 different diseases and compare its diagnostic performance against 21 board-certified dermatologists. They find the AI system capable of classifying skin cancer at a level of competence comparable to the dermatologists.

2017

Speech Recognition on Switchboard
In 2017, Microsoft and IBM both achieved performance within close range of "human-parity" speech recognition in the limited Switchboard domain

2017

Poker
In January 2017, a program from CMU called Libratus defeated four to human players in a tournament of 120,000 games of two-player, heads up, no-limit Texas Hold'em. In February 2017, a program from the University of Alberta called DeepStack played a group of 11 professional players more than 3,000 games each. DeepStack won enough poker games to prove the statistical significance of its skill over the professionals.

2017

Ms. Pac-Man
Maluuba, a deep learning team acquired by Microsoft, created an AI system that learned how to reach the game's maximum point value of 999,900 on Atari 2600.

2018

Chinese - English Translation
A Microsoft machine translation system achieved human-level quality and accuracy when translating news stories from Chinese to English. The test was performed on newstest2017, a data set commonly used in machine translation competitions.

2018

Capture the Flag
A DeepMind agent reached human-level performance in a modified version of Quake III Arena Capture the Flag (a popular 3D multiplayer first-person video game). The agents showed human-like behaviours such as navigating, following, and defending. The trained agents exceeded the win-rate of strong human players both as teammates and opponents, beating several existing state-of-the art systems.

2018

DOTA 2
OpenAI Five, OpenAI's team of five neural networks, defeats amateur human teams at Dota 2 (with restrictions). OpenAI Five was trained by playing 180 years worth of games against itself every day, learning via self-play. (*OpenAI Five is not yet superhuman, as it failed to beat a professional human team*)

2018

Prostate Cancer Grading
Google developed a deep learning system that can achieve an overall accuracy of 70% when grading prostate cancer in prostatectomy specimens. The average accuracy of achieved by US board-certified general pathologists in study was 61%. Additionally, of 10 high-performing individual general pathologists who graded every sample in the validation set, the deep learning system was more accurate than 8.

2018

Alphafold

DeepMind developed Alphafold that uses vast amount of geometric sequence data to predict the 3D structure of protein at an unparalleled level of accuracy than before.

2019

Alphastar

DeepMind developed Alphastar to beat a top professional player in Starcraft II.

2019

Detect diabetic retinopathy (DR) with specialist-level accuracy

Recent study shows one of the largest clinical validation of a deep learning algorithm with significantly higher accuracy than specialists. The tradeoff for reduced false negative rate is slightly higher false positive rates with the deep learning approach.

Measurement Questions

- In recent years, we've seen machine learning based approaches demonstrate increasingly good performance on tasks as diverse as image recognition, image generation, and natural language understanding. Since many of these techniques are data-intensive or compute-intensive, there is a need for metrics that measure the *efficiency* of AI systems, as well as their raw capabilities.

- Moving from single task to multi-task evaluation for AI capabilities, how should the importance of various sub-tasks be weighted for assessing overall progress?

- How can tasks where we're making *no progress* be measured? Many measures of AI progress exist because developers can build systems which can (partially) solve the task - how can areas that are challenging for contemporary systems be assessed?

Chapter Preview

Chapter 4: The Economy

[Table_of_Contents]
[Jobs_Technical_Appendix] [Investment_Activity_Technical_Appendix] [Robotics_Technical_Appendix]

Introduction

This chapter is broken into three sections: Jobs, Investment Activity, and Corporate Activity.

The first section on AI Jobs shows data relating to AI jobs, hiring, and skill levels around the globe as well as in US regions. It includes the AI Hiring Index across countries, sectoral demand for AI jobs, and skill penetration of AI by countries, sector, and gender. The section concludes with trends in skill penetration and labor demand for AI jobs from a sub-regional US perspective. The data on AI hiring, skill penetration by gender and sector are drawn from the LinkedIn Economic Graph. The information about online AI job postings for the US by states and metropolitan areas are based on data from Burning Glass Technologies. According to our sources, there has been a rapid increase in hiring for all categories of AI jobs over the past three years, but they remain a small share of total jobs.
The second section on Investment presents startup investment trends for the world, by countries, and by sectors. The data is sourced to CAPIQ, Crunchbase, and Quid. This is followed by trends in Corporate Investment that includes global AI investment activity by investment types: private startup investment, Mergers & Acquisitions (M&A), Initial Public Offering (IPO), and Minority Stake investments. Finally, public investment trends from the US are presented based on data from BloombergGOV.

The third section on Corporate Activity includes data on adoption of AI capabilities in industry, drawing from McKinsey's Global AI survey. This section also presents global trends in robot installations across countries, drawing from data collected by the International Federation of Robotics (IFR).

Global Hiring

Which countries are experiencing the fastest growth in AI hiring? The hiring rate has been increasing across all the sampled countries, especially for many emerging markets, not just advanced economies.[7] The chart below presents the AI Hiring Index, which is calculated as the percentage of LinkedIn members who had any AI skills (see Appendix for AI Hiring Index definition and Appendix box for the AI skill grouping) on their profile and added a new employer to their profile in the same year the new job began (Figure 4.1.1). The AI hiring rate is normalized for the different countries by dividing over the total number of LinkedIn members in the country. The growth rate is indexed against the average annual hiring in 2015-16; for example, an index of 3 for Singapore in 2019 indicates that the AI hiring rate is 3 times higher in 2019 than the average in 2015-16. The chart shows that the countries with the highest growth in AI hiring on LinkedIn include Singapore, Brazil, Australia, Canada and India.[8] The rapid growth in AI hiring is also confirmed by job postings data from Burning Glass that shows the share of AI jobs (% of total jobs posted online) grew from 0.1% in 2012 to 1.7% in 2019 for Singapore (see Appendix Graph). Similarly, in the US the share of AI jobs grew from 0.3% in 2012 to 0.8% of total jobs posted in 2019. The next section shows the growing role of AI jobs in the US by AI clusters and then economic sectors.

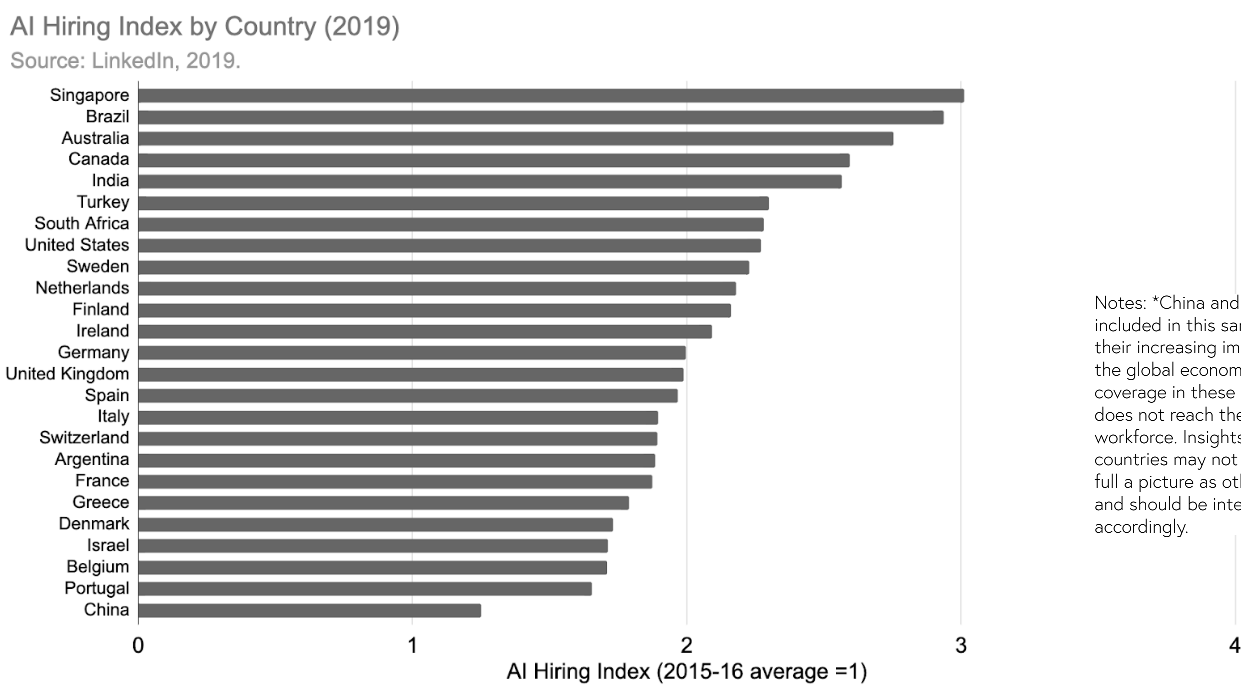

Notes: *China and India were included in this sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly.

Fig. 4.1.1.

"Right now the conversation around AI's impact on individual jobs, and the economy more broadly, is dominated by intensely hyped and alarmist commentary. These discussions need to be grounded in facts and measurement, and this report will hopefully contribute to a more thoughtful, reality-based discussion on trends that could drive big impact in the coming decades."
Guy Berger, Principal Economist at LinkedIn, 2019

[7] Two filters were applied for the countries to be included: 1) countries must have sufficient labor force coverage by our data sources (roughly >40%); and 2) they must have at least 10 AI talents in any given month. Countries and regions with significant representation of their workforce on LinkedIn included in this analysis are United States, Netherlands, Ireland, Denmark, Australia, United Kingdom, Luxembourg, Canada, Singapore, Belgium, New Zealand, Norway, Sweden, United Arab Emirates, France, Portugal, Switzerland, Chile, Spain, Italy, Hong Kong (SAR), Finland, Israel, Costa Rica, Brazil. China and India were included in this sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly. More generally, LinkedIn's Hiring Rate tracks hires or job switches on LinkedIn; this measure has a strong track record in the US tracking government data on job openings (JOLTS) and core capital goods orders (LinkedIn's Economic Graph, 2019).
[8] It should be noted that the analysis depends on the representativeness of LinkedIn users across countries.

US Labor Demand by Job Cluster

Is AI labor demand gaining significance in total jobs posted on the web in the US? Which type of AI jobs witnessed the fastest growth in online job postings in the US? The different clusters of AI job postings from the US are presented by month (Figure 4.1.2). These are mutually exclusive and independent skill clusters for AI jobs. The Appendix provides a graph on total number of jobs by skill clusters and a table, which shows the list of AI skill clusters. Machine Learning jobs increased from 0.07% of total jobs posted in the US in 2010 to over 0.51% in October, 2019, Other important categories of jobs include Artificial Intelligence (0.28%), Neural networks (0.13%), NLP (0.12%), Robotics (0.11%), and Visual Image Recognition (0.10%). The Appendix also provides a breakdown of jobs by AI clusters from Indeed.

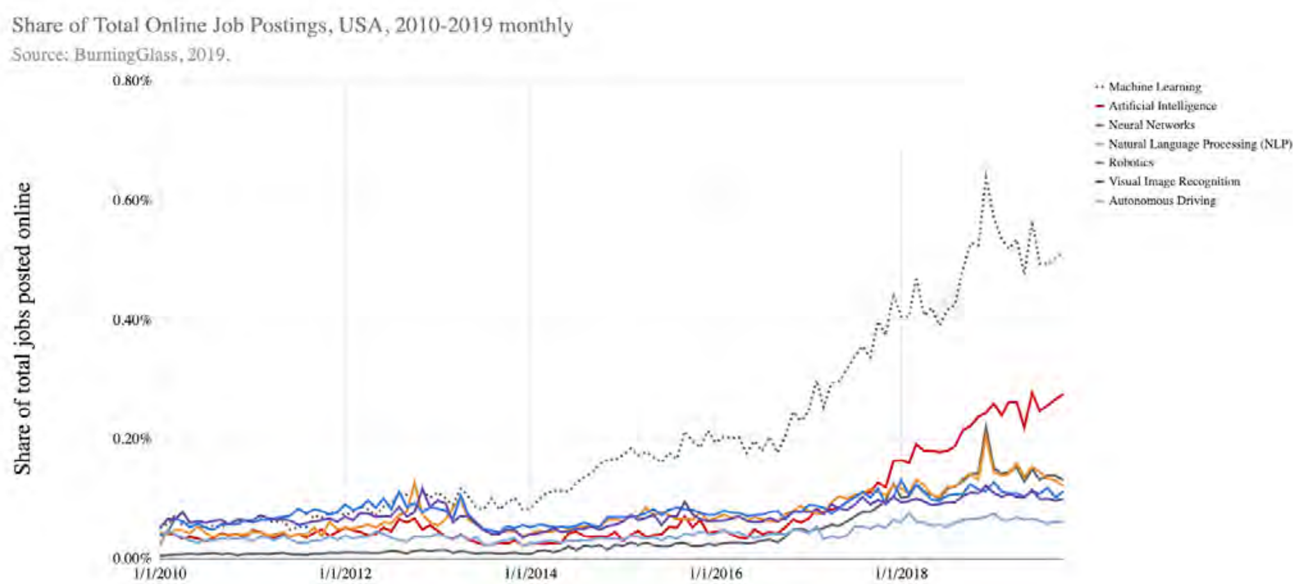

Fig. 4.1.2.

Machine Learning jobs increased from 0.07% in 2010 to over 0.51% in October, 2019 of total jobs posted in the US, followed by Artificial Intelligence jobs (0.28%), Neural networks (0.13%), NLP (0.12%), Robotics (0.11%), and Visual Image Recognition (0.10%).

US Labor Demand By Sector

Which sectors in the US labor market are experiencing stronger AI diffusion via AI job demand? Among sectors, tech, service sectors and manufacturing show the greatest rise in demand for AI skills. The charts below plot the number of AI jobs posted as a percentage of the total jobs posted by sectors in the US. The first provides the ranking of industries with highest demand (percent of total jobs posted) in 2019 (Figure 4.1.3); while the second chart provides a time-series view for the individual sectors (Figure 4.1.4).

Tech service sectors like Information have the highest proportion of AI jobs posted (2.3% of the total jobs posted), followed by Professional, Scientific and Technical Services (over 2%), Finance and Insurance (1.3%), Manufacturing (1.1%), and Management of companies (0.7%). The demand for AI jobs has increased across all economic sectors. The proportion of AI jobs posted across Information, Professional, Scientific and Technical, Finance and Insurance, Administrative and Waste Management has increased by over one percentage point (in terms of share of total jobs posted). On the other hand, the traditional services sector, which includes construction, arts, public administration, healthcare and social assistance, demonstrates a relatively lower demand for AI jobs.

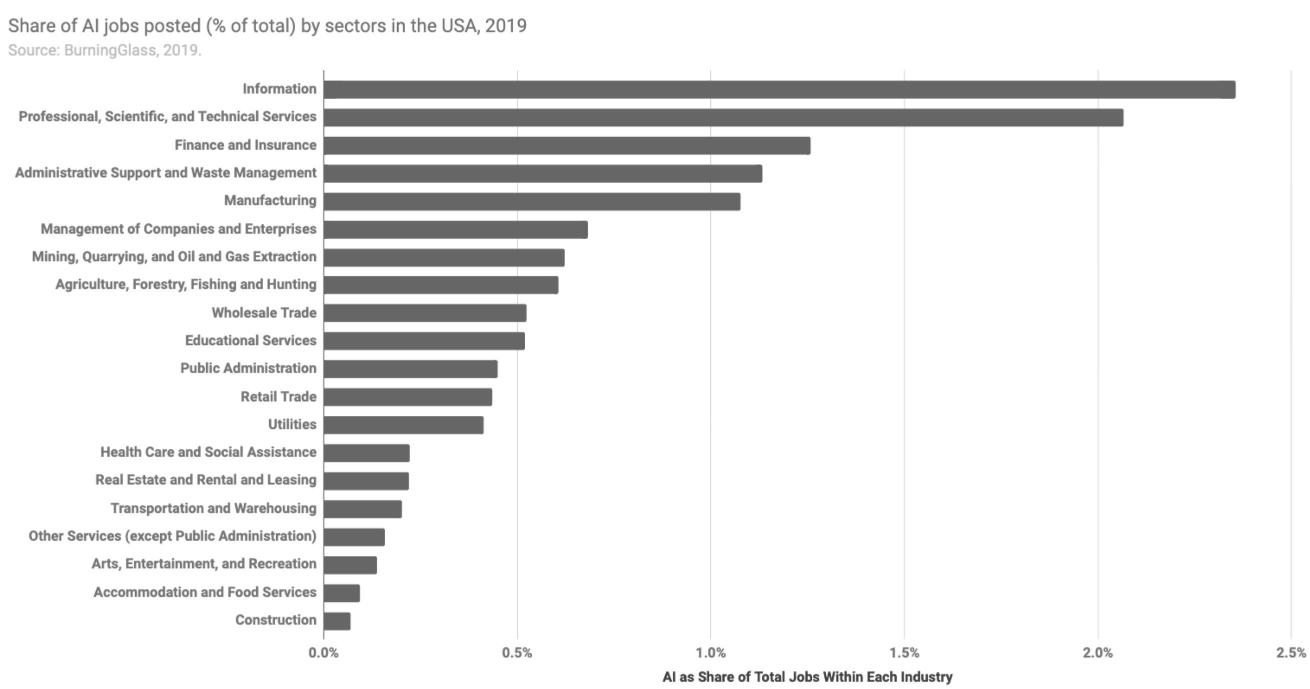

Fig. 4.1.3.

US Labor Demand By Sector

Share of AI jobs posted (% of total jobs posted) by Industry, 2010-19

Source: BurningGlass, 2019

Fig. 4.1.4.

AI labor demand is growing in significance especially in hi-tech services and the manufacturing sector.

[Jobs_Technical_Appendix]
[Access_Data]

Global Skill Penetration

Penetration and Relative Penetration of AI Skills

Using LinkedIn data, the **Penetration of AI Skills** in a given country is defined as the average share of AI skills among all the top 50 skills in each occupation, across all occupations in that country. This metric can also be computed at the sector-country level.

Since different countries have different occupation sets, this penetration rate may not be directly comparable across countries. To allow for cross-country comparisons, the **Relative Penetration of AI skills** is defined as the ratio between the penetration of AI skills in a given country and the average penetration of AI skills across all countries in the sample, considering only the overlapping occupations between the country and the sample.

Skills data are drawn from the member profiles of professionals on the LinkedIn platform. Specifically, the data are sourced from the skills listed on a member's profile, the positions that they hold and the locations where they work.

LinkedIn has categorized and standardized the over 35,000 unique skills on its standard platform into a set of skills clusters using nonlinear embedding spaces. These clusters are seeded by humans and subsequently applied to co-occurrences of skills on profiles across the entire platform. Skills are related by distance in "skill space." Closely-related skills are tagged with a common human-curated cluster name. Skills that co-occur less frequently are classified in separate clusters. Neural skills embeddings are supplied by the LinkedIn engineering team.

In order to compute this metric, LinkedIn first calculates a weight for each skill based on the prevalence of that skill in a particular segment, such as a particular geography, sector, and/or occupation, and compares it to other segments of the labor market. First, all members who hold the occupation during the relevant period are included in the analysis. Then a frequency measure is assigned to each skill by calculating the number of times members list the skill under the "skills" section of their LinkedIn profile. Note that skills are only included in the analysis if they were specifically added during the period for which the individual has held that position. The skills that are added by fewer than or equal to 10 members during the pre-defined period are dropped to reduce 'noise' in the skills data. Skills are only captured if they are relevant to the role and enables a comparison between skills profiles over time. Finally, each occupation-skill pair is weighted following a term frequency–inverse document frequency (TF-IDF) model: skills that are generic and appear in multiple occupations are down-weighted. The result is a list of skills that are most representative of that occupation in that sector and country.

See also: Data Science in the New Economy Report (World Economic Forum, July 2019).

Skill Penetration

Which countries have the highest penetration of AI skills? The relative skill penetration rate metric is based on a method comparing the share of AI skills for each country against a global average/benchmark based on the same set of occupations. For a given country, the relative skill group penetration is the ratio between the penetration rate of a given skill group in each country and the global average penetration rate.

An interesting example is India. The average penetration of AI skills in India in selected sectors is 2.6 times the global average across the same set of occupations. It is interesting to note that India is expected to add over 10 million new young people to the labor force every year over the next decade (Economic Times, 2018). This gain in labor talent raises an interesting question of how India will use its demographic dividend to train, produce, and export sophisticated AI products and services for inclusive growth and development.

The results below are presented for sample countries where there is sufficient coverage (Figure 4.1.5).[9] An occupation on LinkedIn is one of roughly 15,000 job categories added by LinkedIn members; Members have also added 35,000 types of skills to their profiles The horizontal axis of the chart is the number of unique occupations in a country that have any AI skills in their top 50 skills, as reported by LinkedIn members. This is not a per-capita metric. The results represent pooled skill additions between 2015 and 2018. The three step process to calculate relative skill penetration rates are documented in the Appendix. Bar charts in Appendix show the ranking of countries on these measures.

National Comparison of Skill Penetration and Number of Unique AI Occupations

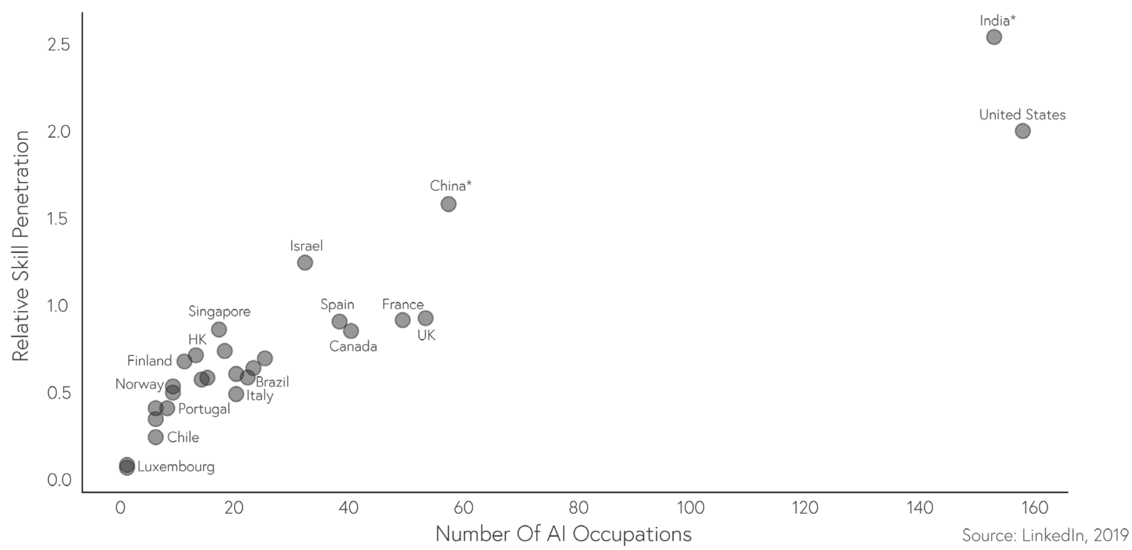

Notes: *China and India were included in this sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly. Number of unique AI occupations refers to the number of unique job titles with high skill intensity.

Fig. 4.1.5.

"While the impact of AI on economies has been primarily concentrated in developed economies on the technological frontier, it's important to note its impact on developing economies. In China and India, the two largest developing economies, we're seeing a similarly large surge in AI skill prevalence."
Guy Berger, Principal Economist at LinkedIn, 2019

[9] Countries and regions with significant representation of their workforce on LinkedIn (roughly >40%) included in this analysis are United States, Netherlands, Ireland, Denmark, Australia, United Kingdom, Luxembourg, Canada, Singapore, Belgium, New Zealand, Norway, Sweden, United Arab Emirates, France, Portugal, Switzerland, Chile, Spain, Italy, Hong Kong (SAR), Finland, Israel, Costa Rica, Brazil. China and India are included in this sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly.

Skill Penetration

In order to provide a deeper sectoral decomposition of AI skill penetration across sectors and countries, the following sample top five sectors with the highest AI skill penetration globally are chosen: Software & IT Services, Hardware and Networking, Education, Finance, and Manufacturing (Figure 4.1.6). India, the US, France, China, and Israel are frequently among the top countries in AI Skill Penetration across all countries. The US ranks in the top 5 countries for AI skill penetration across all sectors. As noted earlier, the large labor pool in India and its IT skills provide hope for cautious optimism as AI could become a driver for occupational diversity, jobs and growth. China only shows up in the top 5 ranking in the education-related skill penetration. Other pockets of specialization worth highlighting include Norway and Israel in AI skills in Software and IT; Norway, France, and Sweden in Hardware and Networking; France, Israel, and Sweden in hardware and networking as well as manufacturing; Spain and Switzerland in education; and the UK and Canada in finance.

Global AI Skill Genomics: Ranking of Sectoral Relative AI Skill Specialization by Countries, 2018

Sectoral Rankings of AI Skill Penetration Scores, by Country

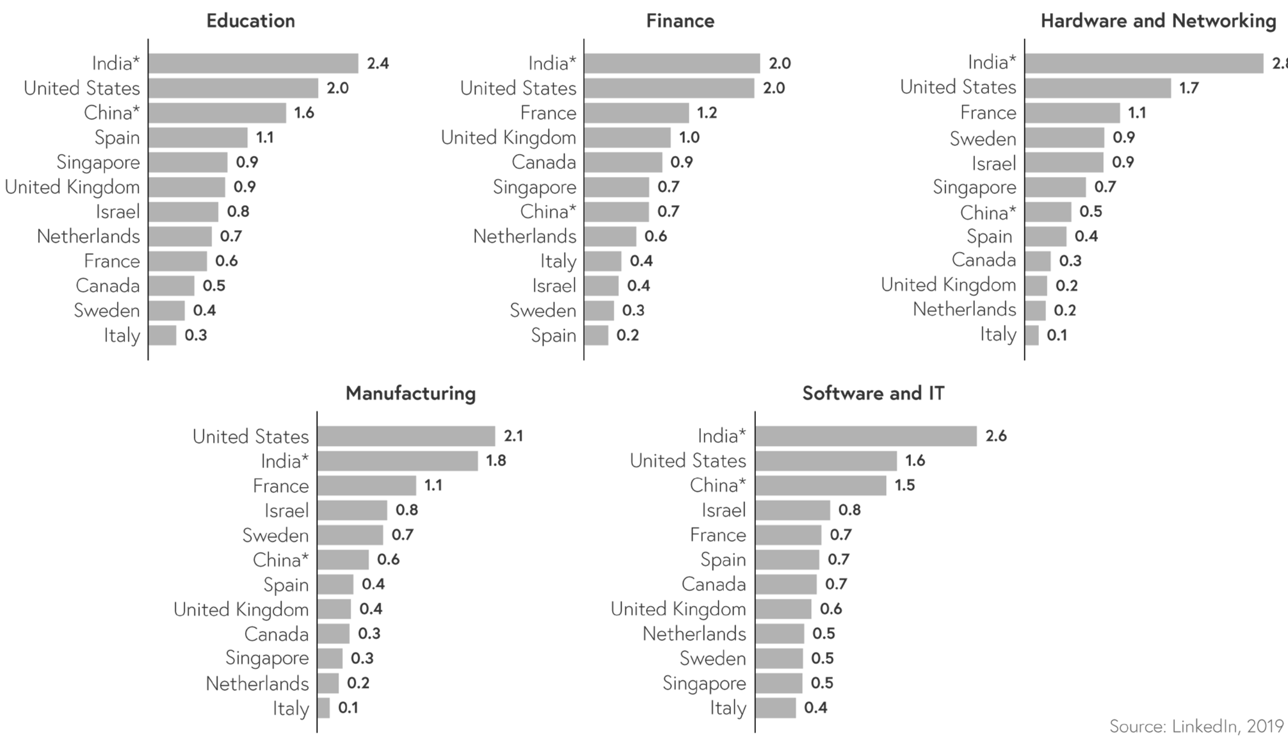

Fig. 4.1.6.

*China and India were included in this sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly.

How will India utilize its demographic dividend to train, produce, and export sophisticated AI products and services for inclusive growth and development?

Inclusion: Global Skill Penetration By Gender

Which countries exhibit relatively higher AI skill intensity by gender? The chart below presents the ranking of countries based on AI skill penetration for female and male labor pools (Figure 4.1.7).[10] Two trends are worth noting. First, men tend to report AI skills across more occupations than women in all countries in the sample. Second, while countries with high AI skill penetration for men are more likely to exhibit high AI skill penetration for women as well, this pattern is not universal. Some European countries --including the Netherlands, Switzerland, and France-- rank significantly higher when considering only women than when considering men. More granularly, the results indicate that the average occupation held by women in India exhibits over 2.6 times the global average AI skill penetration, while the average occupation held by men in India is 2.7 times the global average AI skill penetration. In terms of AI skill reported for women, India is followed by the US (1.5), Netherlands (1), Switzerland (0.94), and France (0.90). For example, India has 55 occupations where women report AI skills whereas men report AI skills in 127 occupations in 2015-2018.

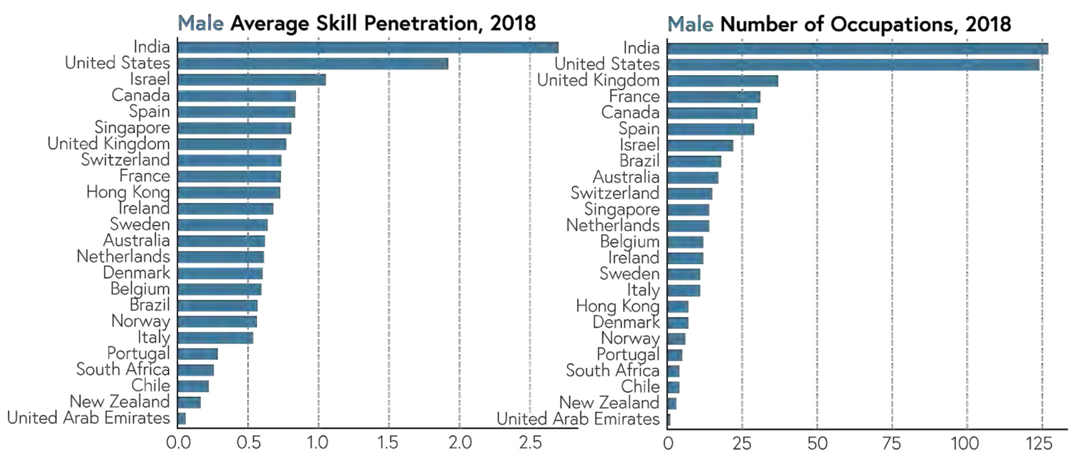

Fig. 4.1.7a.

Source: LinkedIn, 2019

* India was included in this sample due to its increasing importance in the global economy, but LinkedIn coverage does not reach the 40% of the workforce. Insights for this country may not provide as full a picture as other countries, and should be interpreted accordingly. Number of unique AI occupations refers to the number of unique job titles with significant skill intensity.

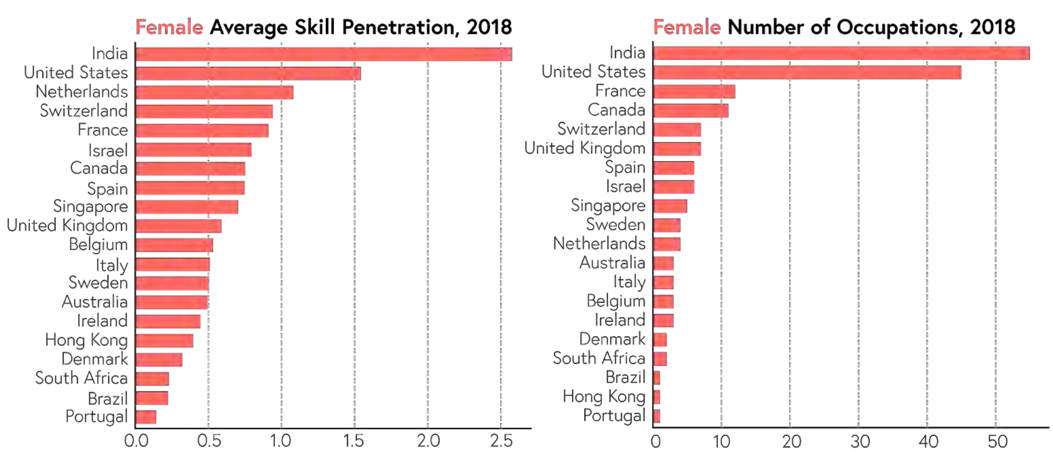

* India was included in this sample due to its increasing importance in the global economy, but LinkedIn coverage does not reach the 40% of the workforce. Insights for this country may not provide as full a picture as other countries, and should be interpreted accordingly. Number of unique AI occupations refers to the number of unique job titles with significant skill intensity.

Source: LinkedIn, 2019

Fig. 4.1.7b.

"Like a lot of other promising -- but not quite mature -- technologies, the AI talent pool is growing at an extremely fast pace. And the pace at which these folks are being hired is growing even faster. More than ever before, this surfaces the need for public and private sector interventions that ensure enough workers are trained and reskilled to meet the rapidly-growing demand for AI skills."
Guy Berger, Principal Economist at LinkedIn, 2019.

[10] "Female" and "male," "women" and "men" are the terms used in the data set. Samples in this analysis consider an additional data filter: having gender data on at least 66% of LinkedIn members. Note that China does not meet this threshold and is thus excluded.

Labor demand and skill penetration by US state

Here the regional AI labor demand and skill penetration by states in the US is examined, followed by metropolitan areas, and cities.

The first chart plots the (relative) importance of AI labor demand as the AI share of total jobs posted on the y-axis and the (absolute) size of labor demand measured as the natural log of total number of AI jobs posted between 2018 and September, 2019 (Figure 4.1.8). Appendix graphs present the ranking of the absolute and relative AI labor demand metrics for US states.

The results show that Washington state has the highest relative AI labor demand with almost 1.4% of total jobs posted are AI jobs. Washington is followed by California with 1.3%, Massachusetts with 1.3%, New York with 1.2%, and the District of Columbia (DC) with1.1%, and Virginia with 1 AI jobs. There are 5 states in addition to Washington, DC where over 1% of total jobs posted are AI jobs. Majority of states lie between 0.2 and 1% of total jobs posted.

In absolute terms California has the largest number of AI jobs posted. Over 93,000 AI jobs were posted in California since 2018. This is three times the volume of the next state, New York, with 30,000 AI jobs posted in AI. Texas was next with over 24,000 jobs posted, followed by Massachusetts with over 19,000, Washington over 18,000, and Virginia over 15,000. The full state level AI labor demand metrics are available here.

Relative importance of AI jobs and absolute size of AI labor demand, 2018-19

Source: Burning Glass, 2019.

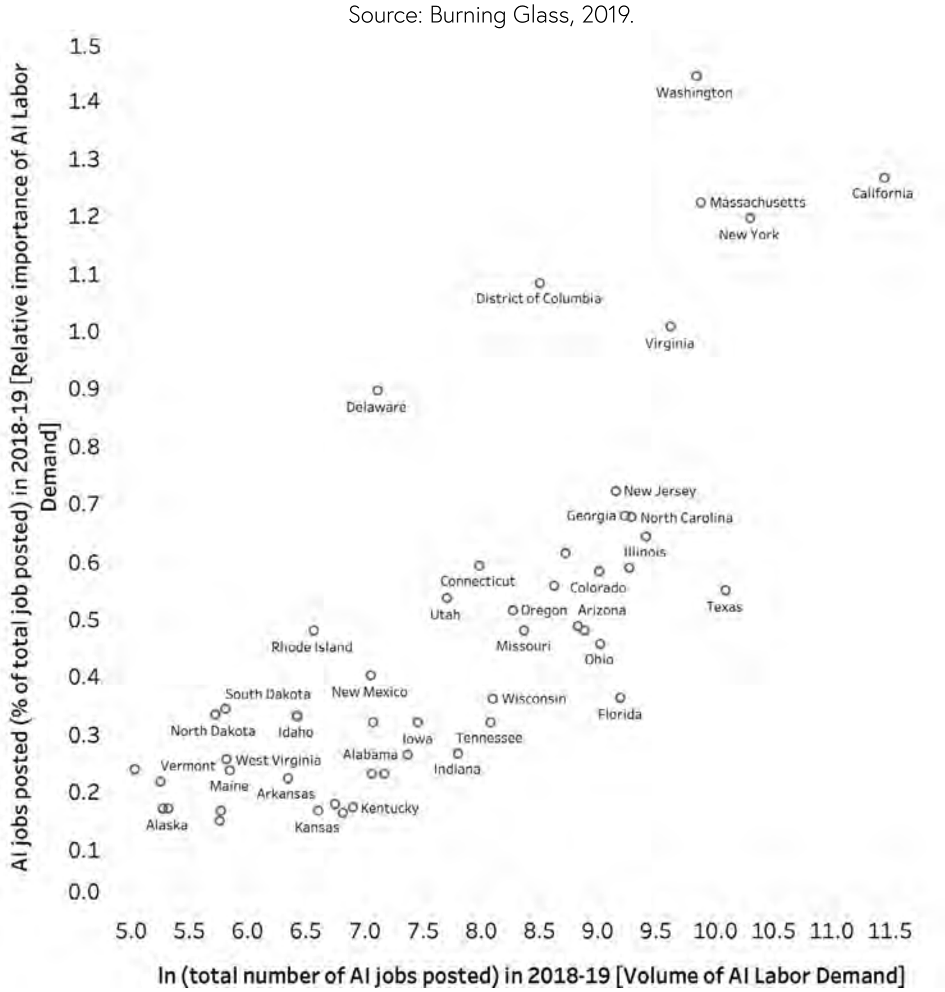

Fig. 4.1.8.
Note: The chart plots the sum of AI job postings in 2018 which includes data up until September of 2019.

Regional Dynamics (US)

Has US AI-related labor demand converged across states over the last decade? The answer is mixed. In terms of absolute labor market demand for AI jobs, the evidence points towards unconditional convergence i.e. the states that had low labor market demand 10 years ago in 2010 witnessed relatively faster growth in AI job postings than big states. Appendix charts show unconditional convergence in absolute labor demand. However, the evidence also points towards unconditional divergence in relative AI labor market demand. Appendix chart on unconditional divergence in relative US state level AI labor demand shows that the relative importance (or the relative size of AI job postings) has grown fastest in initially large AI states. For example, states like Washington, California, Massachusetts, Virginia, New York, Maryland or DC witnessed an increase in AI share of total employment greater than 0.2 percentage points since 2010.

US state maps show the average annual growth in AI jobs between 2010-19 (Figure 4.1.9a) and AI relative skill penetration respectively (Figure 4.1.9b). With convergence in absolute AI job posting growth, initial conditions matter. States like Wyoming starting with a very small base experience faster growth in AI job postings of over 70%, followed by North Dakota with over 65%, Nevada with over 50%, Rhode Island and Montana with over 45% average annual growth between 2010-10. However, in terms of AI skill penetration only states such as California, New York, and Texas appear to have higher relative AI skill penetration.

Average annual growth in AI jobs postings for US States, 2010-19

Source: Burning Glass, 2019.

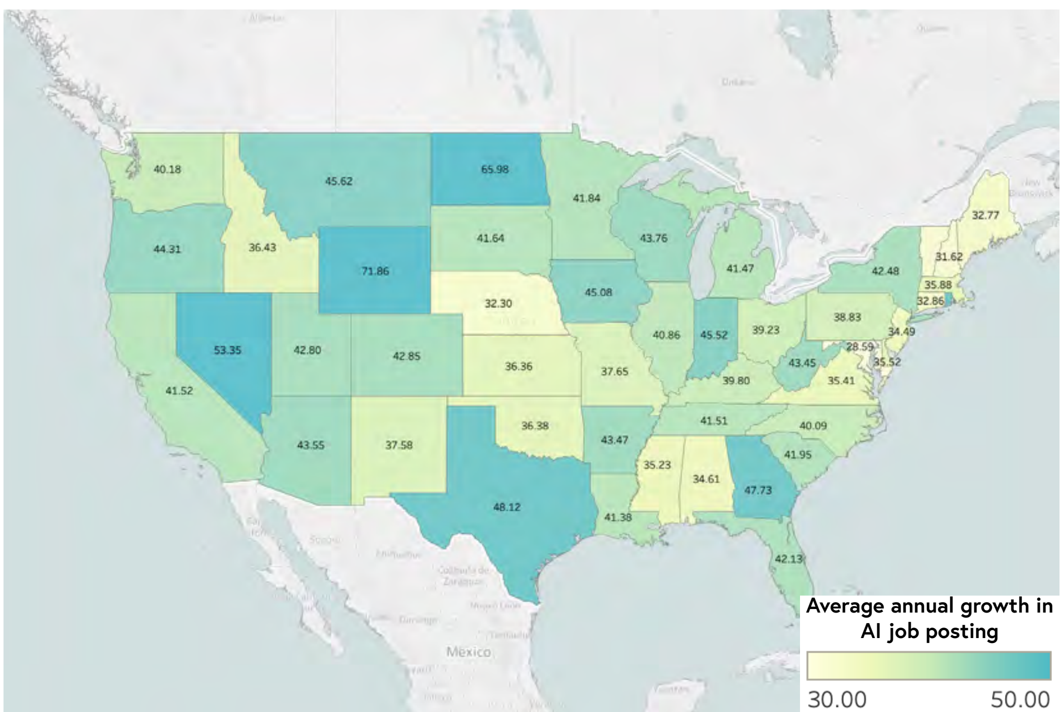

Figure 4.1.9a

Note: The color represents the average annual growth in AI job postings as measured by the natural log difference between the sum of AI jobs posted between 2018 and September, 2019 - the natural log of total AI jobs posted between 2010-13, divided by the time-period difference.

The states that had low labor market demand 10 years ago in 2010 also witnessed fast growth in AI job postings along the big states.

Regional Dynamics (US)

US States AI Skill Penetration, 2018

Source: LinkedIn Economic Graph, 2019.

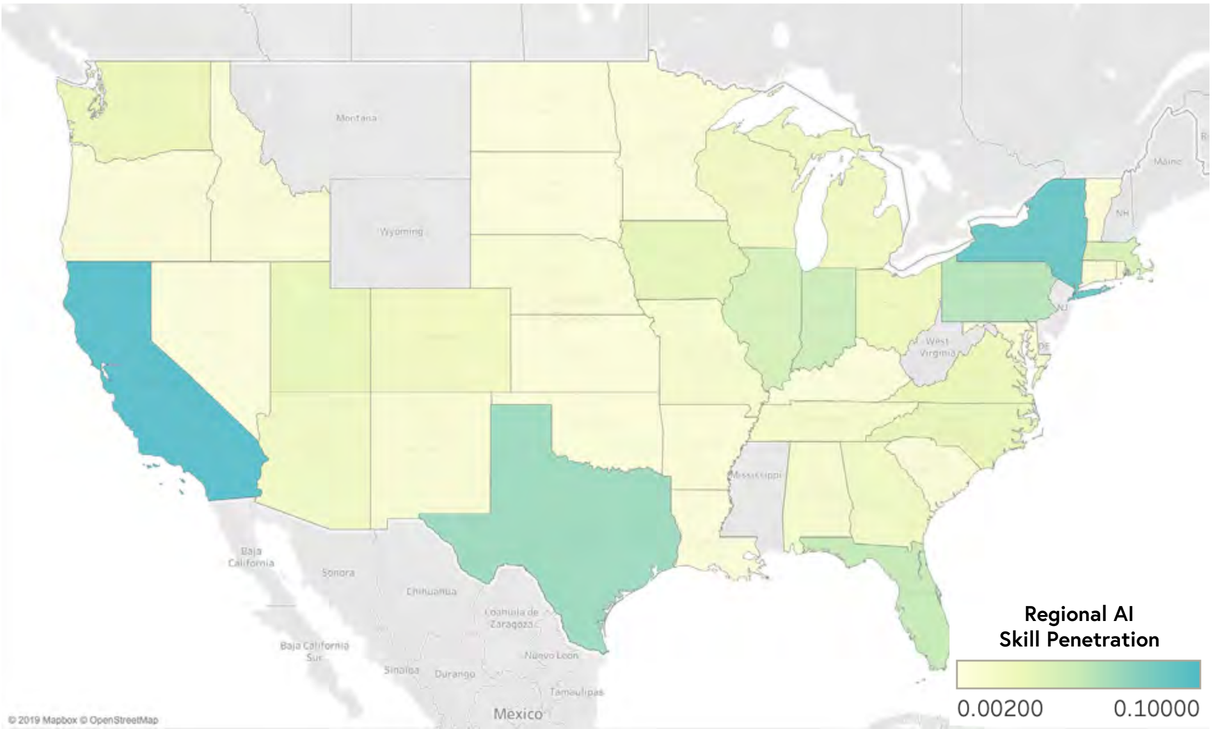

Fig. 4.1.9b.

Labor Demand and Skill Penetration by US Metropolitan Areas and Cities

What are the deeper regional dynamics of AI job demand in the US? Is demand primarily concentrated in tech epicenters, or is it dispersing across the country? The map of the US for Metropolitan Statistical Areas (MSA's) is presented below (Figure 4.1.10). The size of the bubble represents the absolute size of labor demand, i.e., total number of AI jobs posted. The largest bubble size represents the total number of AI jobs posted 20,000 jobs in a given MSA. The color schematic represents the relative importance of AI labor demand, with the shade of blue representing any MSAs with greater than 1 percent share of AI jobs in total AI jobs posted. Readers should note that the sample size of smaller MSAs is not reliable for a small sector like AI; hence the data is missing.

In addition to details on the data and methodology, readers can also observe the evolution of AI jobs and the economic impact across different regions. The methodology is discussed in Appendix.

Regional Dynamics of AI labor demand in the US

Source: Burning Glass, 2019

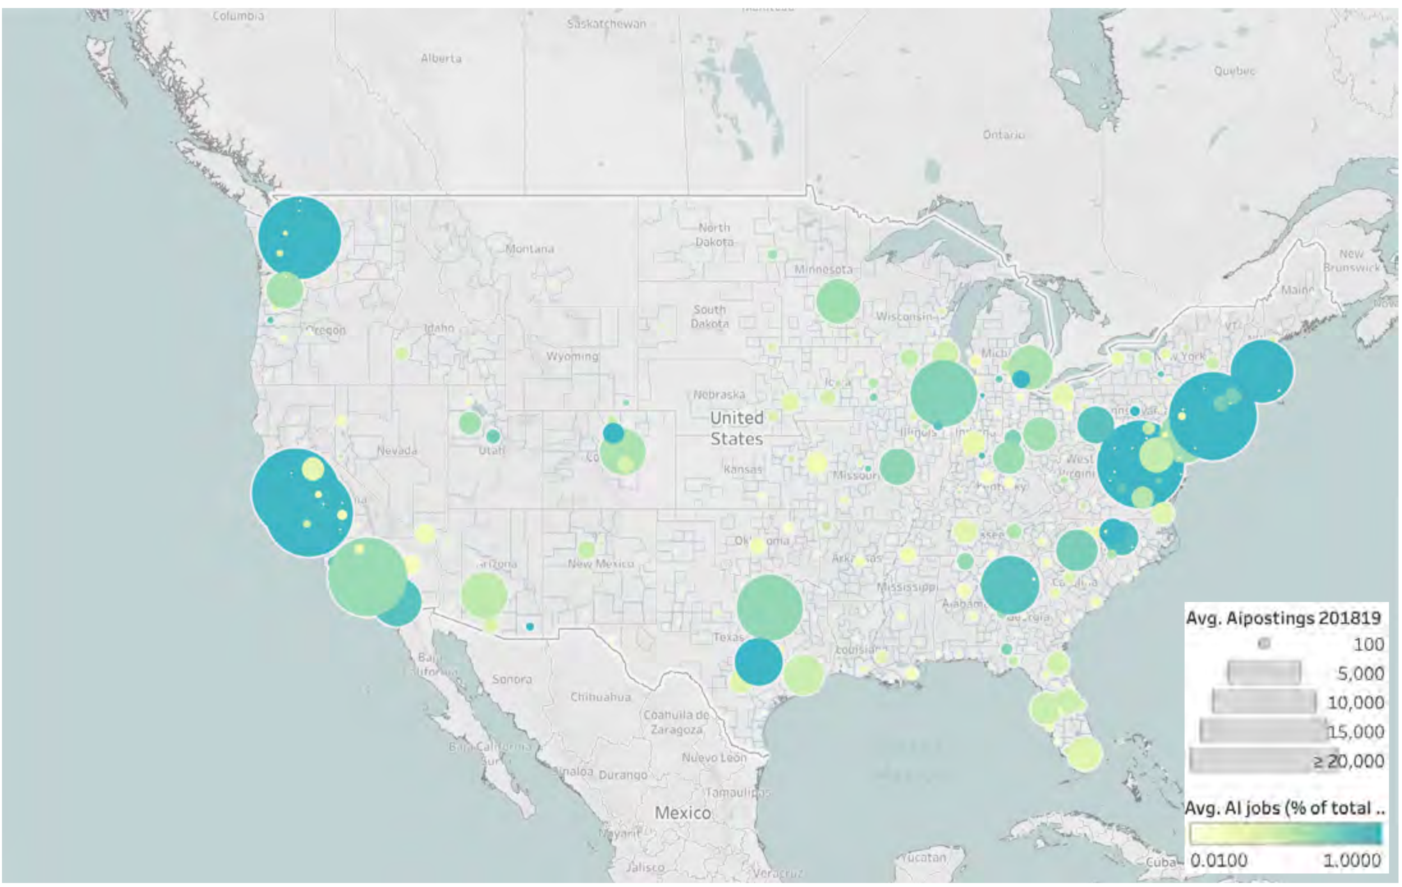

Fig. 4.1.10.
Notes: Alaska and Hawaii have not been presented for presentational brevity.

Labor Demand and Skill Penetration by US Metropolitan Areas and Cities

Is there a convergence in AI job posting across metropolitan areas across the US? The chart below plots the average annual growth in total number of AI jobs as a share of IT jobs between 2010 and 2019 for almost 400 MSAs on the vertical axis and the natural log of total number of AI jobs posted in 2010 on the horizontal axis (Figure 4.1.11). The results are again mixed but with no convergence across MSAs for total number of AI jobs posted and unconditional divergence in relative AI labor demand. [The detailed graphs are presented in Appendix](). In the chart below, the graph is broken into four quadrants. The top right quadrant represents the areas that already had high AI job demand and also witnessed rapid growth over the last decade. The top left quadrant represents the areas that are emerging hubs of AI job demand. The bottom left quadrant had a relatively low stock of AI jobs ten years ago and further shrinking since then, while the bottom right quadrant had a relatively high stock of AI jobs in the past but shrinking AI demand since then.

In absolute terms many emerging areas have high growth in AI labor demand. Columbus, Ohio; Knoxville, Tennessee; Jacksonville and Gainesville, Florida; Beckley, West Virginia witnessed the fastest absolute growth in AI job posting starting from a very small base. Knoxville has not been widely discussed. roximity to Oak Ridge National Lab (ORNL) may have influenced its growth. ORNL and Department of Energy (DOE) are significantly ramping up their AI activities and adding to their workforce in this field. This growth could also contribute to local businesses who might work with ORNL, or work in related areas. Since ORNL is a major employer in a relatively small metropolitan area, their ramp-up in AI would be statistically significant to the workforce opportunities in the area. As a side note, anecdotally, in the past it has been mentioned that Oak Ridge has one of the highest concentrations of PhDs in the country, again because the town is small and ORNL is large. The other emerging areas of AI job demand include Asheville, North Carolina; Pittsburg, Pennsylvania; Ann Arbor, Michigan; Fargo, North Dakota; Virginia Beach-Norfolk, Virginia and North Carolina. [Ranking of top MSA with high absolute and relative growth in AI labor demand and top MSA with shrinking AI labor demand are presented in the Appendix graphs.]()

No clear convergence: Many small metropolitan with low initial stock of AI jobs also experienced fast growth in AI labor demand (2010-19)

Source: Burning Glass, 2019

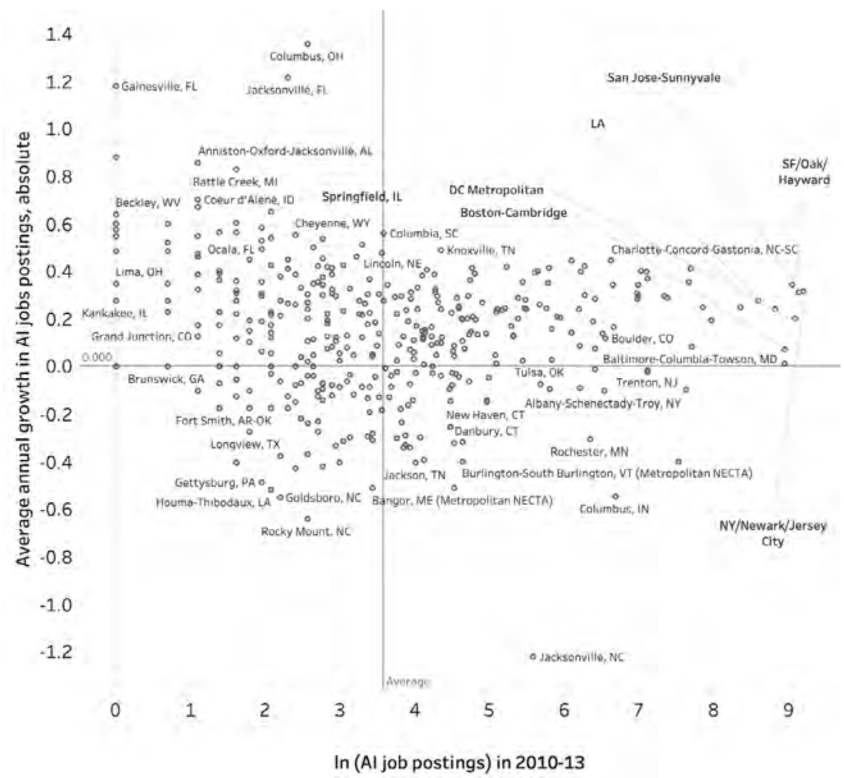

Fig. 4.1.11.

"The growth of AI labor demand in smaller cities and regions of the US illustrates the tremendous potential of AI to generate new types of work across our Nation. Policy strategies for AI education and workforce training – including the President's American AI Initiative and the National Council for the American Worker – will ensure that America's workers are capable of taking full advantage of the opportunities of AI."
Lynne Parker, Deputy US Chief Technology Officer

Labor Demand and Skill Penetration by US Metropolitan Areas and Cities

Table 4.1.1 shows the ranking of AI skill penetration for US regions based on LinkedIn data. Bryan College Station in Texas has the highest relative AI skill penetration in the country, followed by San Francisco Bay Area, Lafayette, Indiana, Binghamton, New York, and Urbana-Champaign, Illinois. This evidence points to greater occupational skill diversity in emerging hubs in addition to Silicon Valley and New York City. Appendix table provide detailed ranking for major US cities on AI skill penetration and provides related results based on LinkedIn data that show unconditional divergence in AI skills across the US regions indicating that the growth in AI skill penetration is faster in areas that initially had high skill penetration. However, the time sample is limited to three years.

Ranking of AI Skill Penetration for US Cities, 2018
Source: LinkedIn, 2019.

City	Rank	City	Rank
Bryan-College Station, TX	**1**	Santa Barbara, CA	**14**
San Francisco Bay Area, CA	**2**	Springfield, MA	**15**
Lafayette, IN	**3**	Madison, WI	**16**
Binghamton, NY	**4**	Raleigh-Durham, NC	**17**
Urbana-Champaign, IL	**5**	State College, PA	**18**
Pittsburgh, PA	**6**	Austin, TX	**19**
Gainesville, FL	**7**	Provo, UT	**20**
Seattle, WA	**8**		
Rochester, NY	**9**		
San Diego, CA	**10**		
Boston, MA	**11**		
Des Moines, IA	**12**		
Bloomington, IN	**13**		

Table 4.1.1.

"Historically, technology can be a vehicle for rising inequality. Policy and social interventions can either mitigate or worsen those trends, so having access to comprehensive data on AI jobs, skills, and tends is critical. These insights help us avoid the bad interventions, and instead invest in those that equitably share the enormous gains that the next wave of technological innovations could generate."
Guy Berger, Principal Economist at LinkedIn, 2019

[Burning Glass: Jobs_Technical_Appendix]_[LinkedIn: Jobs_Technical_Appendix]
[Access_Data]

Measurement Questions

- Traditional statistics and labor force surveys do not yet include AI and related occupations. Thus, online jobs platforms function as proxy indicators to assess the evolution and growth in AI labor market indicators, and largely demonstrate the demand side of labor market outcomes. How can more direct data about the AI workforce be gathered?

- In regard to the data and methodology, one main area for organization is a standard topology of AI skills and keywords to measure AI job metrics. At the moment different online jobs platforms use different processes for data and may have self-selection bias in different country or regional context. Could platforms define standard ways of tagging AI jobs to facilitate further study?

- Data on AI jobs across countries and within countries is not consistently available. More and better collection of data will be required to consistently track developments.

Global

Globally, investment in AI startups continues its steady ascent. From a total of $1.3B raised in 2010 to over $40.4B in 2018 alone (with $37.4B in 2019 as of November 4th), funding has increased with an average annual growth rate of over 48% between 2010 and 2018 (Figure 4.2.1a). We consider only AI companies that received more than $400k in investment. The number of AI companies receiving funding is also increasing, with over 3000 AI companies receiving funding in 2018 (Figure 4.2.1b). Between 2014 and 2019 (through November 4th), a total of 15,798 investments (over $400K) have been made in AI startups globally, with an average investment size of approximately $8.6M.

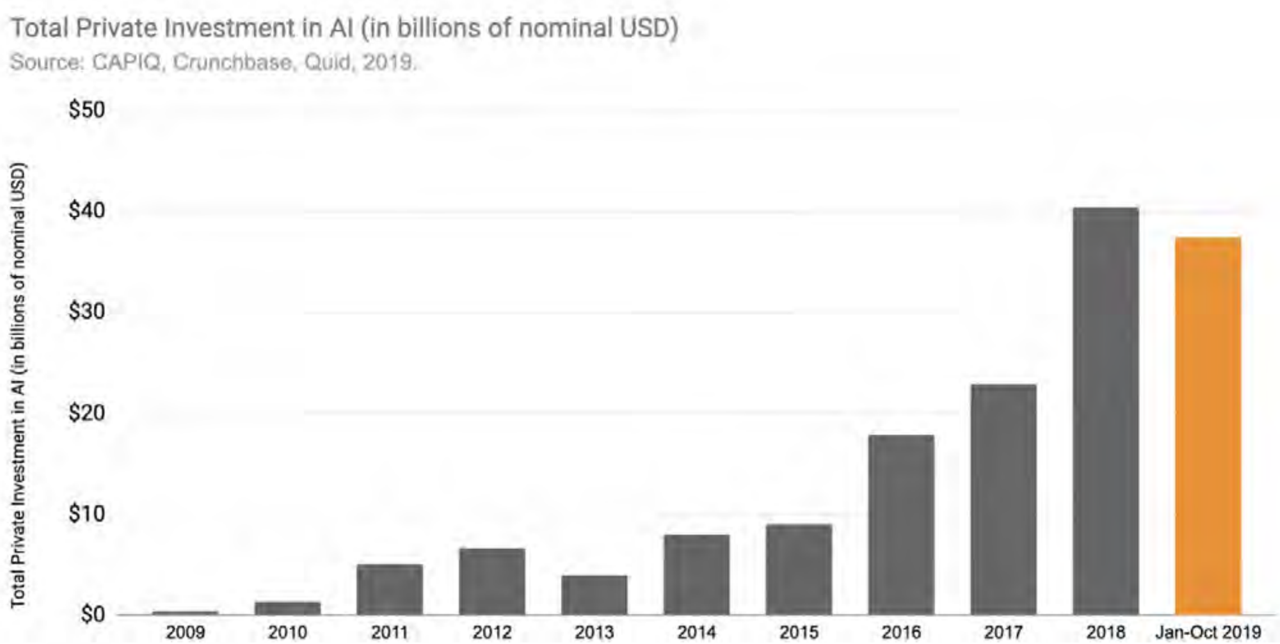

Fig. 4.2.1a.

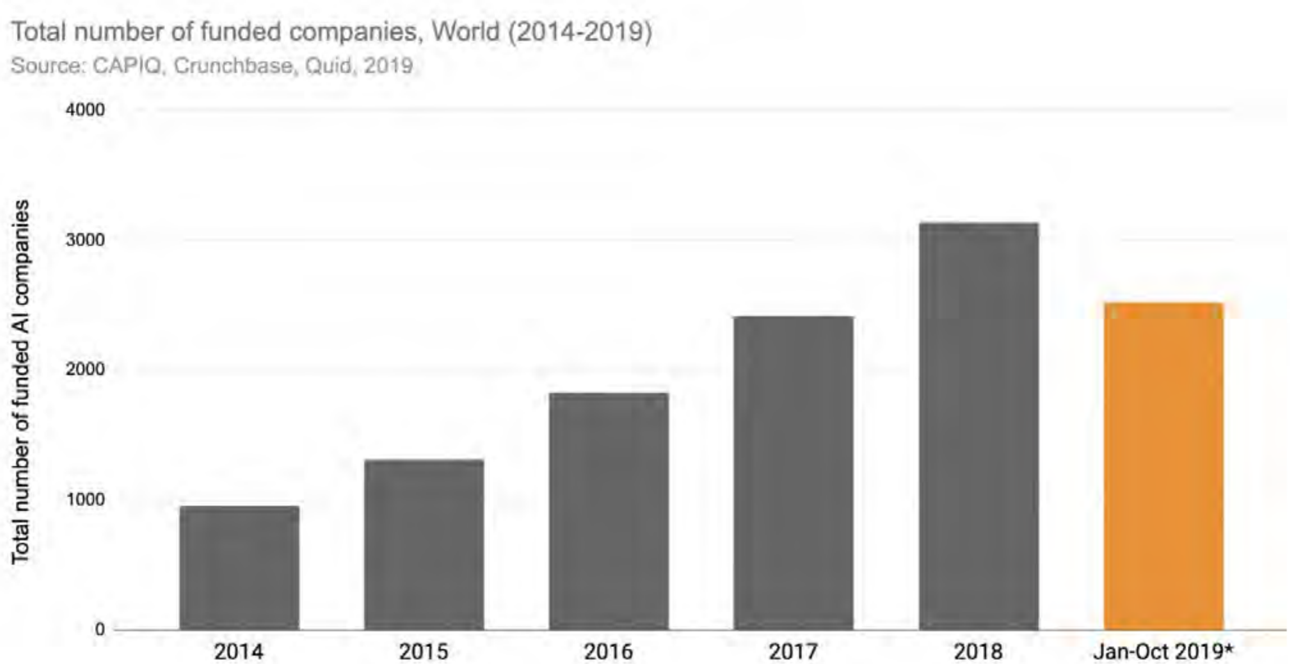

Fig. 4.2.1b.

Country

The United States remains dominant when it comes to the number of funded startups and, in general, has been a consistent leader in AI funding. However, a select few Chinese firms received exceptionally high levels of investment in 2018, which pushed the country closer to parity with the United States (Figure 4.2.2). The underlying detailed time series data can be found here with Appendix graphs providing more detailed country-specific charts.

Which countries appear to be emerging as AI hubs normalized for the size of the country? When adjusted for per capita terms (to reflect the number of startups or investment relative to a country's size), it's actually Israel that has invested the most over the last year, followed by Singapore and Iceland (Figure 4.2.3). During that period, Israel and Singapore also had the largest number of funded startups, trailed a ways back by Iceland, Switzerland, and Canada.

The two graphs above provide data for select economies, however, the full list of countries is available in the appendix. You can also access underlying time series data or appendix graphs that provide more detail with country-specific charts.

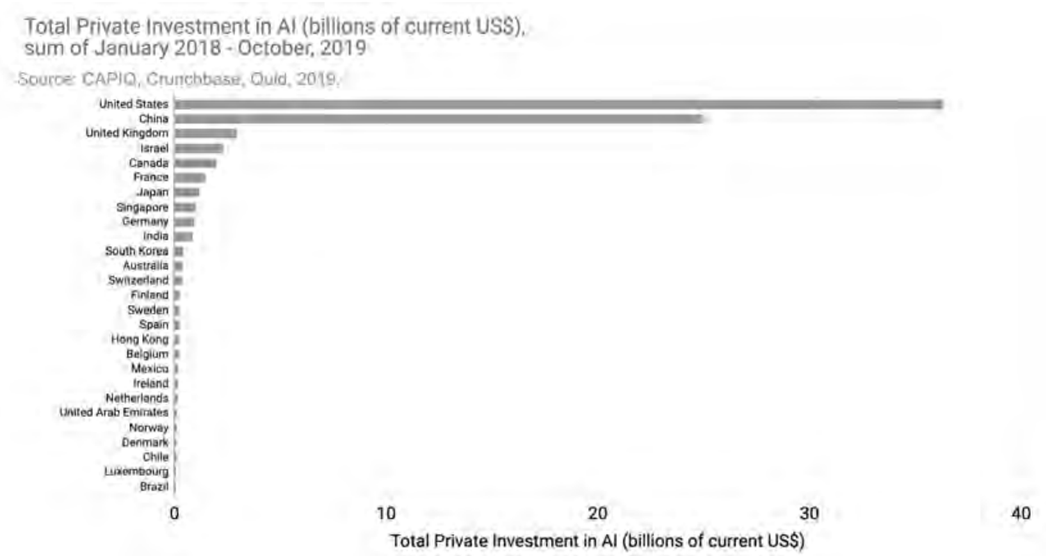

Fig. 4.2.2.

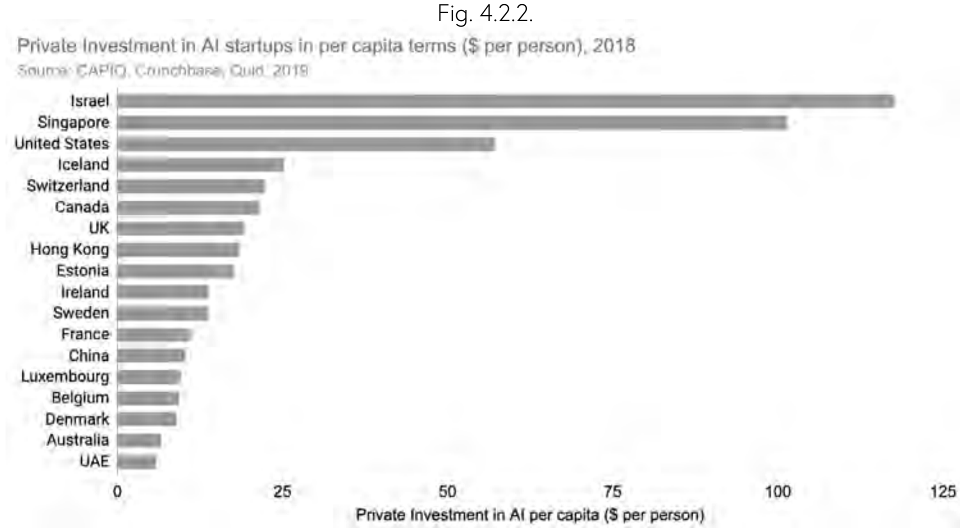

Fig. 4.2.3.

Note: Island economies such as Cayman Islands, British Virgin Islands, Gibralter have been excluded from the sample.

US, Europe, and China take the lion's share of global AI private investment, while Israel, Singapore, and Iceland invest substantially in per capita terms.

Sector

Which are the largest and fastest growing sectors for AI-related investment? Seen in the first graph below (Figure 4.2.4), Autonomous Vehicles (AVs) received the lion's share of global investment over the last year with $7.7B (9.9% of the total), followed by Drug, Cancer and Therapy ($4.7B, more than 6.1%), Facial Recognition ($4.7B, 6.0%), Video Content ($3.6B, 4.5%), and Fraud Detection and Finance ($3.1B, 3.9%).

Which sectors are growing the fastest globally? Seen in the graph below (Figure 4.2.5), robot process automation grew most rapidly (over $1B in 2018), followed by supply chain management (over $500M in 2018), and industrial automation (over $500M in 2018). Other sectors like semiconductor chips, facial recognition, real estate, quantum computing, crypto and trading operations have also experienced substantial growth in terms of global private investment.

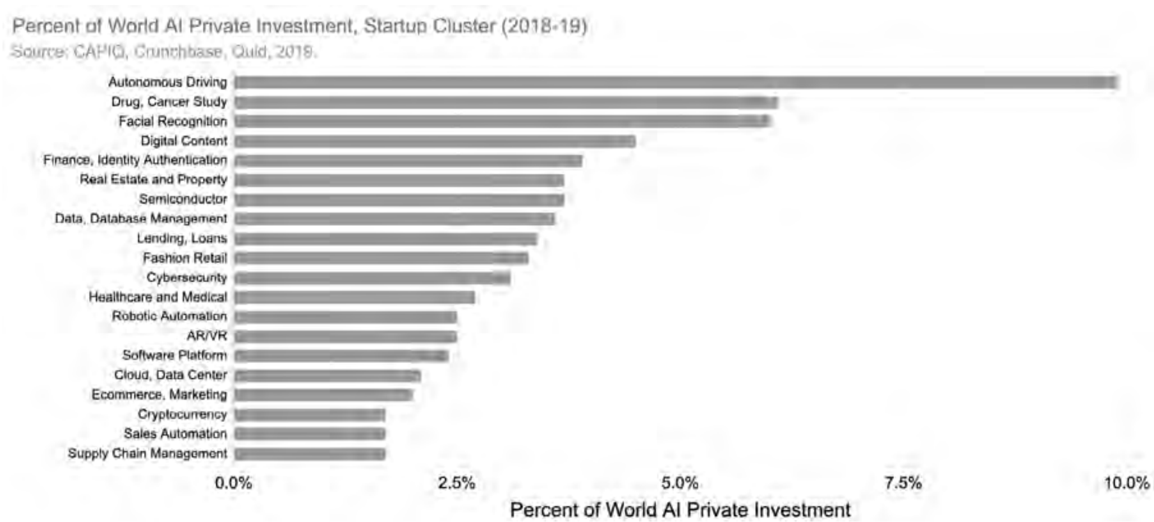

Fig. 4.2.4.
Note: The chart shows the sum of total private AI investments between January, 2018 - October, 2019.

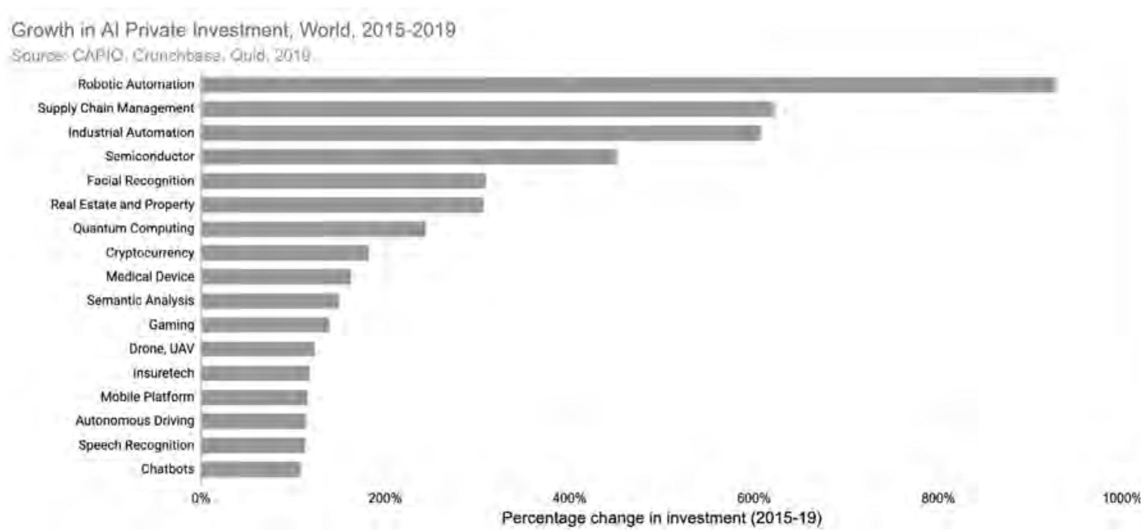

Fig. 4.2.5.
Note: The growth shows growth rate between the 2015-18 (sum) and 2018-19 (sum).

Focus Areas: Global

Given its diverse range of applications—real estate, gaming, finance, healthcare, and security, just to name a few—AI appears to be transforming into a general purpose technology (GPT). Adoption of AI technologies is widely believed to drive innovation across sectors and could generate major social welfare and productivity benefits for countries around the world. One thing is certain: whether directly or indirectly, AI systems play a key role across businesses and shape the global economy for the foreseeable future. New products and processes are developing across a range of industries: supply chains, robotic process automation, speech recognition, sales automation, accounting, natural security, and many more. Using Quid, 36 different global sectors were identified that are currently utilizing AI technologies.

Globally, 4,403 AI-related companies were identified that received investment during the last year. From 36 distinct sectors, top focus areas included **Data Tools** (5.5% of all companies); **Fashion and Retail Tech** (4.7%); **Industrial Automation, Oil & Gas** (4.3%); **Financial Tech** (4.2%); and **Text Analytics** (4.2%). During that time period, these funded startups received a total of $55.7B in private investment, or roughly $12.6M per startup.

Global AI startups that have received funding within the last year (July 2018-July 2019)

Source: CAPIQ, Crunchbase, Quid, 2019.

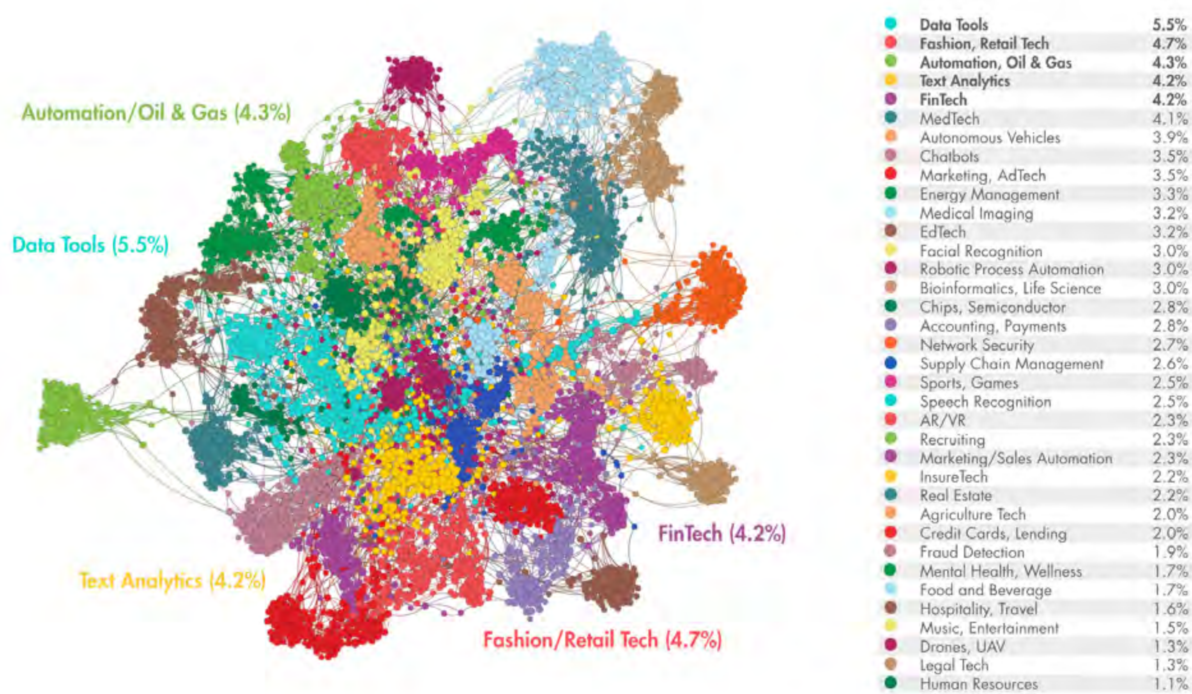

Fig. 4.2.6a.
Network showing 4,403 global AI startups that received investment between July 2018 and July 2019. Colored by sector with top five highlighted.

Appendix: How to Red a Quid Network

AI appears to be transforming into a general purpose technology (GPT). Adoption of AI technologies is widely believed to drive innovation across sectors and could generate major social welfare and productivity benefits for countries around the world.

Focus Areas: Regional

How do key focus areas differ across countries and regions? The following graphs overlap specific country or regional data on the global network map to highlight key differences in the volume and variation of startups for the United States, European Union, China, and India. Seen below, the United States and Europe have the most diverse range of startups—each with some representation across all 36 sectors—even though the US has roughly 70% more companies by volume. In the United States, 1,749 startups were identified that received funding across 36 sectors, with top focus areas including: **Data Tools** (8.1% of all companies); **Medical Tech** (5.3%); **Fashion and Retail Tech** (4.7%); **Text Analytics** (4.7%), and **Chatbots** (3.9%). Most of these categories tracked with global trends; even MedTech and Chatbots ranked highly with the #6 and #8 spots worldwide.

Over the last year, these funded startups received $19.8B of investment, or an average of $11.3M per startup—slightly lower than the global average.
As with the US, each of the 36 global AI sectors has representation in Europe—just on a smaller scale. 993 startups that received funding in the 29 European states were identified during the last year. **Fashion and Retail Tech** (5.7% of all companies) held the top spot, followed by **Medical Tech** (4.4%), **Text Analytics** (4.4%), and a few newcomers to the list: **Marketing and Advertising Tech** (4.3%) and **Autonomous Vehicles** (4%).

During this one year period, funded startups in Europe received a smaller share of the investment pie: a total of $4.6B with an average of $4.7M per startup.

AI startups in the United States: Companies that have received any funding within the last year, by focus area (July 2018-July 2019)

Source: CAPIQ, Crunchbase, Quid, 2019.

Fig. 4.2.6a.
Notes: Network highlighting 1,749 AI startups in the United States that received investment between July 2018 and July 2019. Colored by focus area with top five labeled.

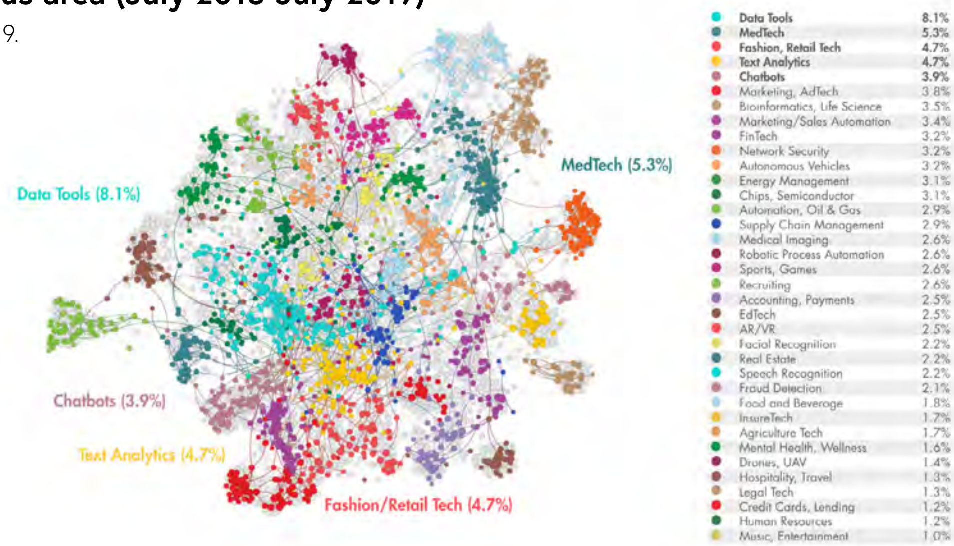

AI startups in the European Union: Companies that have received any funding within the last year, by focus area (July 2018-July 2019)

Source: CAPIQ, Crunchbase, Quid, 2019.

Fig. 4.2.6b.
Notes: Network highlighting 993 AI startups in Europe that received investment between July 2018 and July 2019. Colored by focus area with top five labeled.

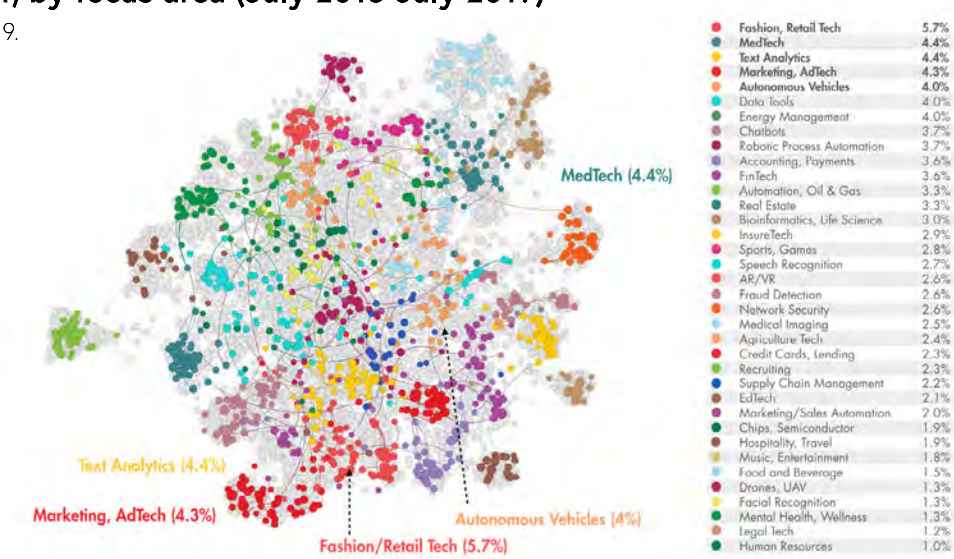

Focus Areas: Regional

AI startups in China received much higher rates of investment during this time period than their Western counterparts. The country's 486 funded startups received a whopping $16.6B in investment, or $34.1M per startup (201% more than startups in the US, and 296% more than the global average).

Though fewer in number, Chinese startups had representation across 35 of the 36 identified global AI sectors. Unlike other countries, **Automation/Oil & Gas** (12%) captured the focus of AI activity, followed by **Facial Recognition** (8.8%); **Education Tech** (8%); **Autonomous Vehicles** (6.4%); and **Mental Health/Wellness** (5%).

India lagged far behind the US, EU, and China when it comes to startup founding and investment. Only 139 startups received funding over the last year, with key focus areas including: **Robotic Process Automation** (6.3%); **Credit Cards/Lending** (5.6%); **Chatbots** (4.9%); **Education Tech** (4.9%); and **Hospitality/Travel** (4.9%). Though sparse, Indian startups were quite diverse in number, matching China and just short of the US and EU with 35 out of 36 focus areas represented.

These startups received $360.1M in private investment, or an average of $2.6M per startup—much lower than the US, Europe, or China.

.

AI startups in China: Companies that have received any funding within the last year, by focus area (July 2018-July 2019)

Source: CAPIQ, Crunchbase, Quid, 2019.

Fig. 4.2.6c.
Notes: Network highlighting 486 AI startups in China that received investment between July 2018 and July 2019. Colored by focus area with top five labeled.

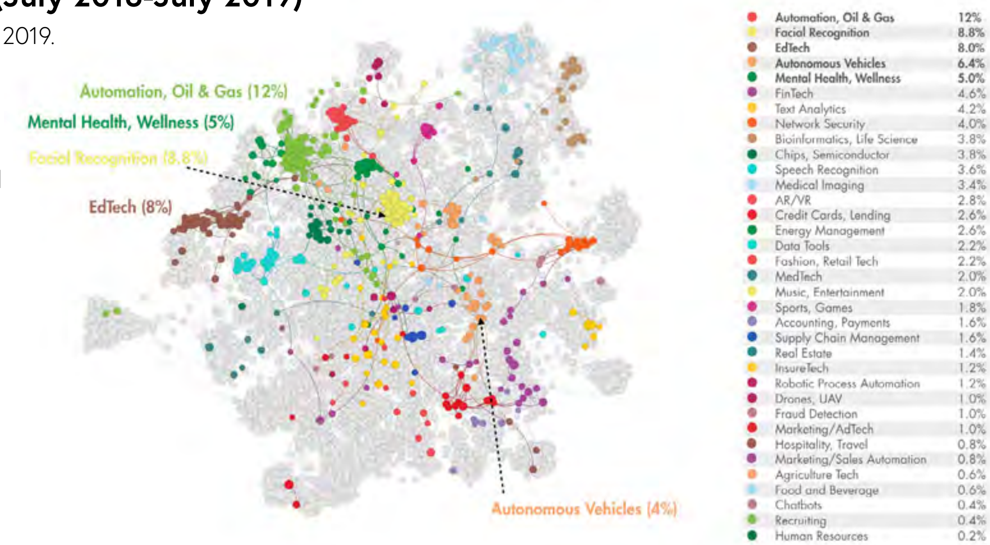

AI startups in India: Companies that have received any funding within the last year, by focus area (July 2018-July 2019)

Source: CAPIQ, Crunchbase, Quid, 2019.

Fig. 4.2.6d.
Notes: Network highlighting 143 AI startups in India that received investment between July 2018 and July 2019. Colored by focus area with top five labeled.

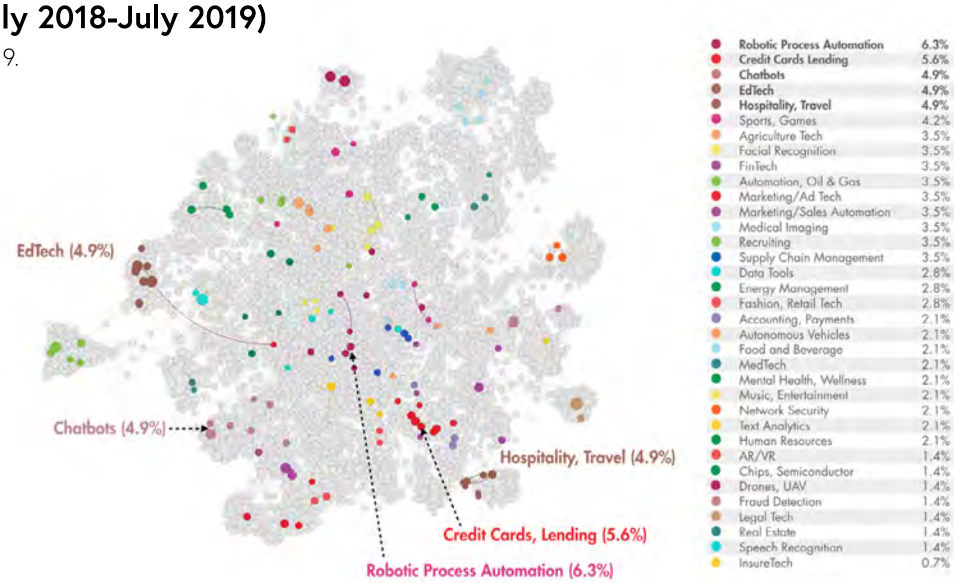

M&As and IPOs

There is growing interest to understand deeper trends in AI Investments. Are M&A, Minority Stake, and Public Offerings equally as big as private investment? The chart below (Figure 4.2.7) plots the volume of different types of investment activity over time. It shows that VC-driven private investment accounted for about half of total investments in AI in 2019, with M&A and Public Offerings taking taking the major share of the remaining half. However, private investment accounted for 92% of the number of deals, with M&A making up just over 4% of deals, and Minority stakes and Public offerings (IPOs) together accounting for 3%. We note that Alibaba's IPO in 2014 accounts for the significant volume of IPO investment in 2014.

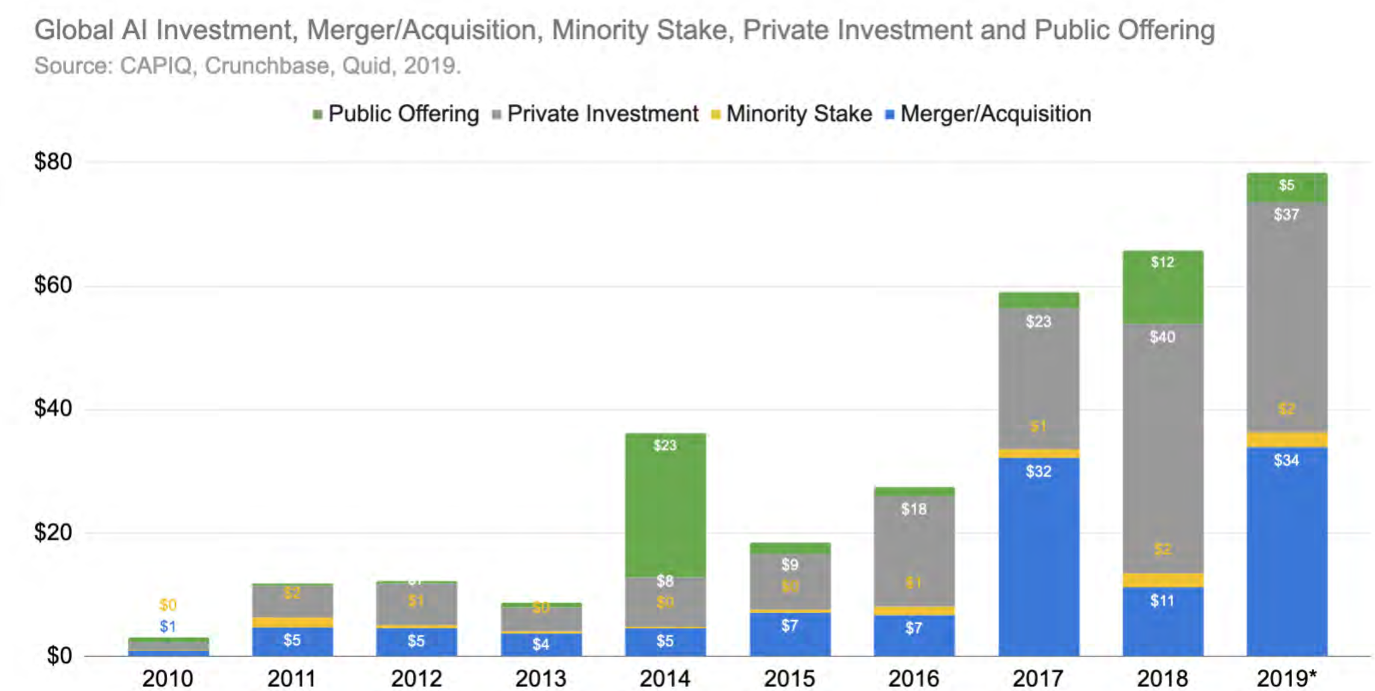

Fig. 4.2.7.
Note: y-axis in billions of US$.* 2019 data is until October, 2019. The jump in 2014 Public Offering reflects Alibaba's IPO.

Mergers & Acquisitions in AI and corporate investment in AI are equally important vehicles for financing AI products and services.

[Investment_Activity_Technical_Appendix]
[Access_Data]

Public Investment

This section considers AI-related public investment for the US only. Reliable cross-country measures on public investment are difficult to obtain since they are no standards in measuring AI investment. Data from Bloomberg Government shows proxy estimates for the Department of Defense (DoD) budget estimates and Contract Spending across US government agencies. Considering federal civilian agencies and DoD budget estimates, the US federal government is projected to invest $4.98 billion in AI R&D in fiscal 2020.

Federal Civilian Agencies' Budgets

In February 2019, the White House issued an executive order that directed US government agencies to, for the first time, quantity their total AI investment and benchmark AI spending year-to-year. In September 2019, the National Science & Technology Council announced that federal civilian (non-Defense Department) agencies expected to invest $973 million on AI, according to a report supplementing the President's Fiscal 2020 Budget Request. The National Science Foundation is the largest civilian funder of AI, with $488 million budget for AI R&D in fiscal 2020, followed by the National Institutes of Health ($203 million), the Department of Energy ($163 million), and the Food and Drug Administration ($39 million). Figures on Defense Department AI R&D were withheld from the report for national security reasons.

Department of Defense (DoD) Budget

The Defense Department is projected to invest another $4.0 billion on AI R&D in fiscal 2020, according to an independent analysis by Bloomberg Government (Figure 4.2.8a). An analysis of the Pentagon's Fiscal 2020 Research, Development, Test & Evaluation (RDT&E) budget request yielded 346 unique budget line items that referenced AI-related keywords in their titles or descriptions. The Defense Advanced Research Projects Agency (DARPA) alone will invest $506 million in fiscal 2020, while the department will allocate $221 million to the Algorithmic Warfare Cross Functional Team, better known as "Project Maven." The cornerstone of the Pentagon's AI program, the Joint AI Center (JAIC), will receive $209 million.

Looking more closely at the DOD's RDT&E budget, the following graphs show the department's AI R&D budgets broken out by programmatic spending area and agency. Applied Research will receive the largest volume of funding ($908 million), followed by $821 million for Rapid Growth Advanced Component Development and Prototyping (ACD&P), and $398 Operational System Development (OSD) (Figure 4.2.8b). Rapid growth in these areas indicates that the Pentagon's focus is scaling and fielding AI prototypes in addition to basic and applied research.

The top AI funding entities within the DOD are the Office of the Secretary of Defense ($1.3 billion), which presides over the department's sprawling Research & Engineering (R&E) enterprise, DARPA ($506 million), and the military services, which collectively will invest $1.57 billion (Figure 4.2.8c).

Department of Defense (DoD) Fiscal 2020 Research, Development, Test & Evaluation (RDT&E) Budget, Artificial Intelligence-specific
Source: Bloomberg GOV, 2019.

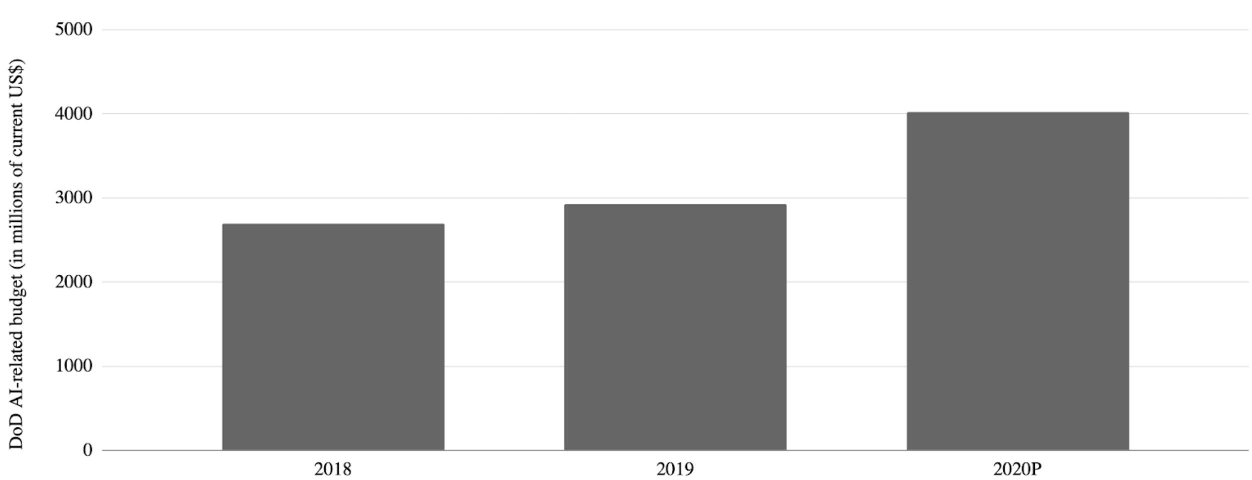

Fig. 4.2.8a.

Public Investment

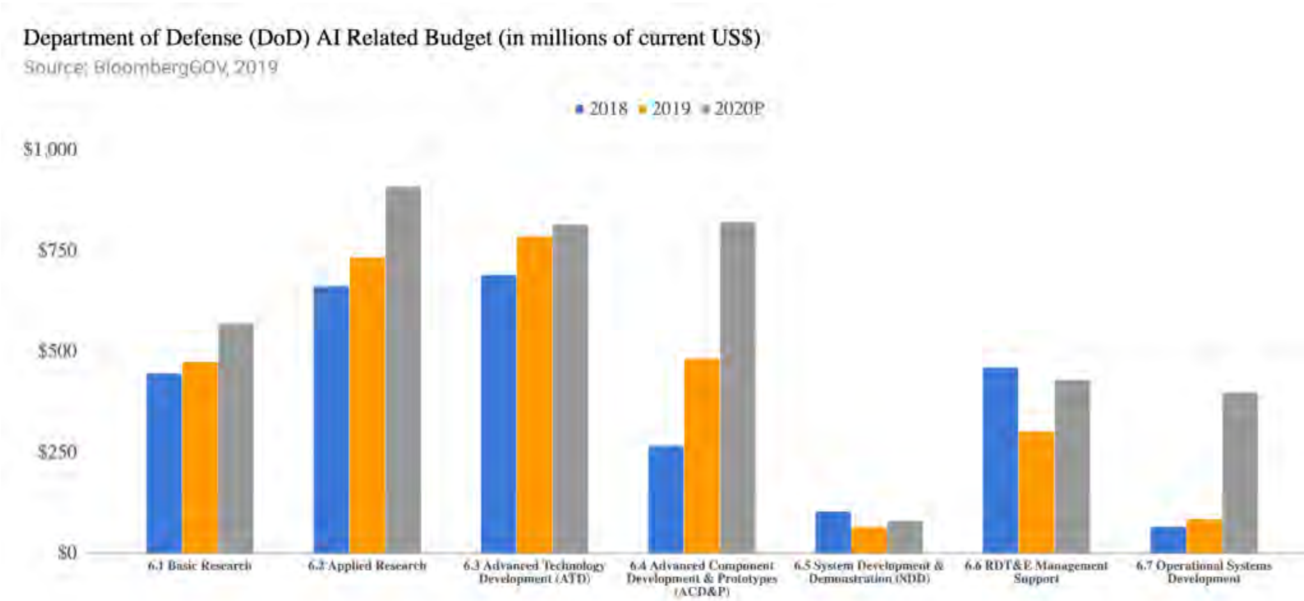

Fig. 4.2.8b.

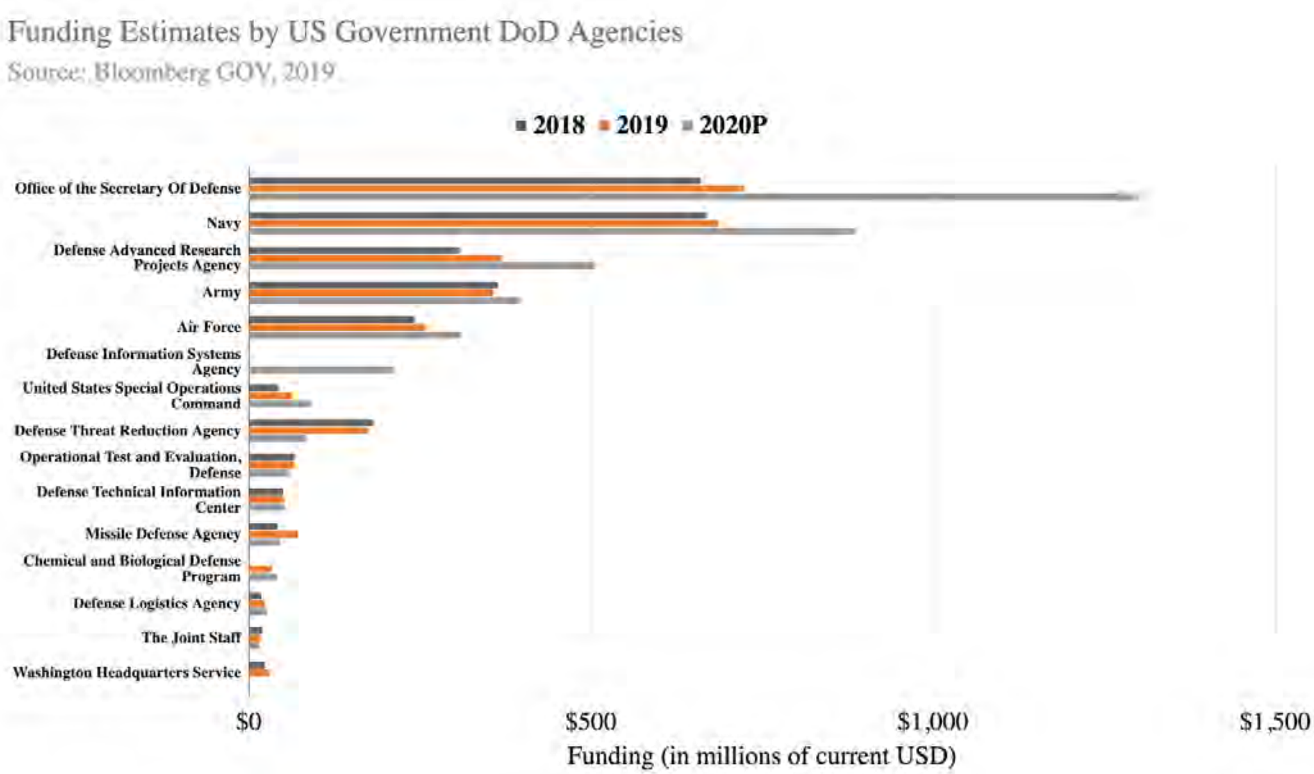

Fig. 4.2.8c.

US Government Contract Spending

Another method of assessing public investment is studying the data on government contracts. The data (Figure 4.2.9a & 4.2.9b) below represents government spending transactions on AI projects between fiscal years 2000 to the present, as defined by Bloomberg Government. Bloomberg built its model using spending data reported by agencies to the [Federal Procurement Data System-Next Generation (FPDS-NG)](). To capture AI spending, Bloomberg first identified all spending transactions associated with R&D and IT projects (GSA Category Management Levels 1 and 17), then identified those that matched with a set of over 100 AI-related keywords (e.g., artificial intelligence, machine learning, neural network).

In fiscal 2018, the latest year in which complete contracting data is available, federal agencies spent a combined $728 million on AI-related contracts, an almost 70% increase above the $429 million that agencies spent in fiscal 2017. Since fiscal 2000, the Pentagon has accounted for the largest share of AI spending of any federal agency ($1.85 billion), followed by NASA ($1.05 billion), and the departments of the Treasury ($267 million) and Health and Human Services ($245 million).

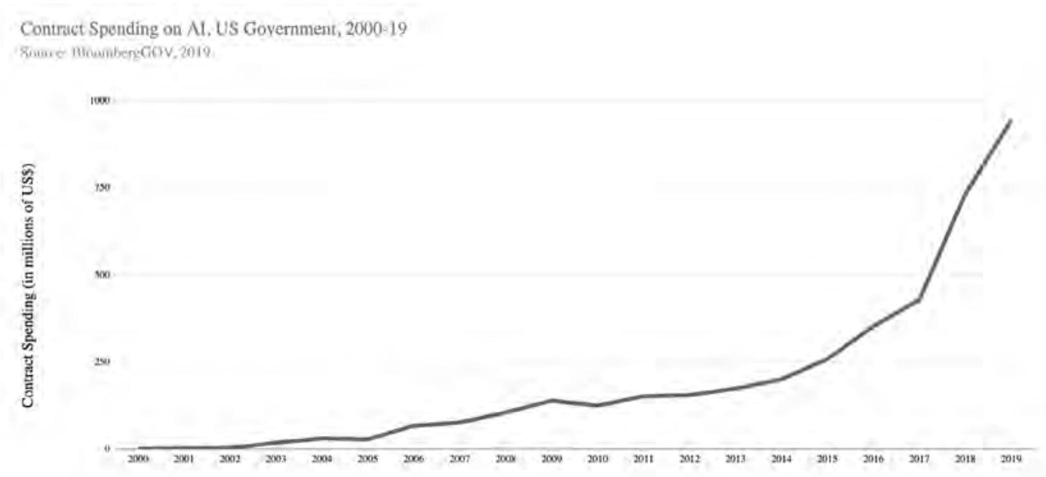

Figure 4.2.9a.

Accounting for Contract Spending across all US Government Agencies

Source: Bloomberg Government based on contract analysis of over 200 government agencies

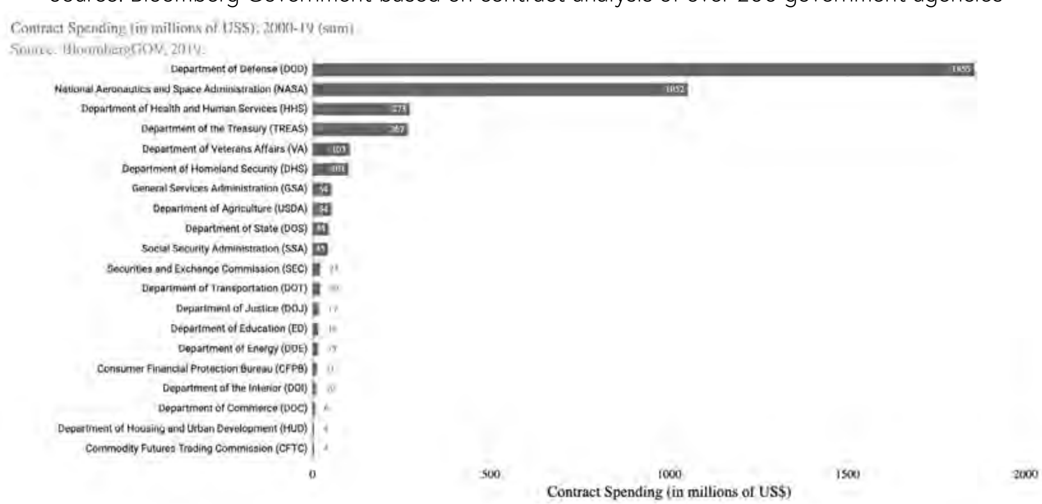

Figure 4.2.9b.

Measurement Questions

- There is no standard consensus on labeling AI related investment activities. For example, startups that could be producers of new AI technologies, or consumers of AI, or others who are not actually involved in AI. It could be interesting to have a more standard labeling mechanism for AI VC investment, as well as corporate investment activities.

- Standard economic measurements can be applied to new data; however, accounting for AI in national accounting or balance of payments is an important discussion for national statistical agencies. There are no existing measurement and accounting standards for public investment or expenditure in artificial intelligence.

- Since AI is a technology that can be produced, transmitted, and consumed across borders, deeper data to uncover growing trading of AI across borders will be an important measurement question for policy decisions.

- Data on public investment is not consistently available across countries. The data here reflect public investments in the US While some data is available regarding announcements that some governments have made, how much of this has actually been invested is less clear. It will be important to continue to track such public investments.

Industry Adoption

The graphs on the following pages show the result of a McKinsey & Company survey of 2,360 company respondents, each answering about their organizations. The full results of this survey, which include insights about how high-performing companies have adopted AI, the capabilities required to scale AI across the business, and the financial outcomes that companies have experienced by adopting AI, are published in McKinsey & Company's "[Global AI Survey: AI proves its worth, but few scale impact.]"

AI adoption by organizations is increasing globally

The results suggest a growing number of organizations are adopting AI globally. Fifty-eight percent of respondents report that their companies are using AI in at least one function or business unit#, up from forty-seven percent in 2018 (Figure 4.3.1a). Adoption appears to be more equally distributed across regions than in 2018, with about six out of ten respondents in most regions reporting their organizations have embedded AI. Across regions, respondents in developed Asia–Pacific report the largest growth since 2018, with a 19-percentage-point increase in companies embedding AI in at least one business function or business unit.

AI adoption *within* businesses has also increased. Thirty percent of respondents report that AI is embedded across multiple areas of their business, compared with 21 percent who said so in 2018 (Fig 4.3.1b).

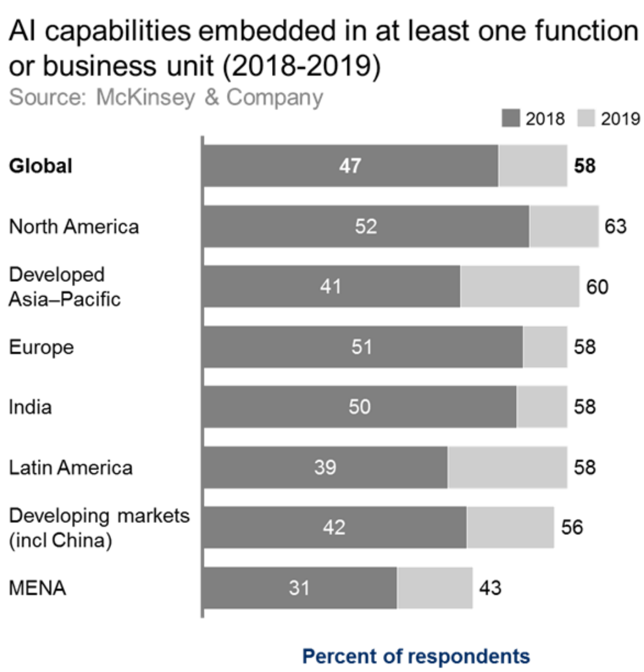

Fig. 4.3.1a.

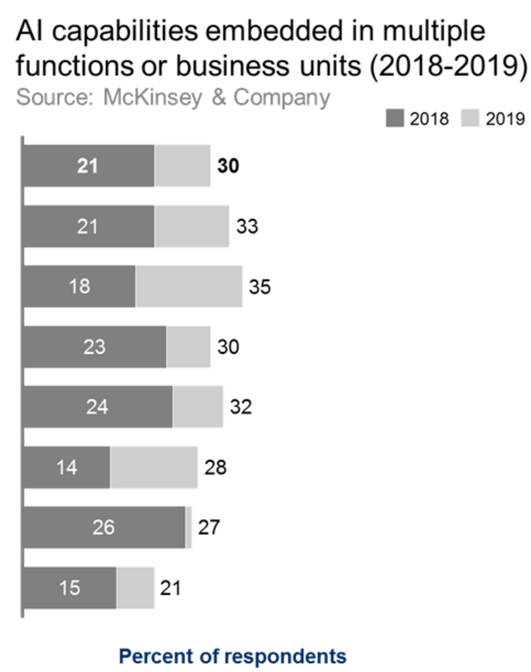

Fig. 4.3.1b..

Industry Adoption

Organizations adopt AI in business functions that provide most value in their industry

Continuing the trend of 2018, companies are most likely to adopt AI in functions that provide core value in their industry (Figure 4.3.2).

For example, respondents in the automotive industry are the most likely to report adoption of AI in manufacturing, and those working in financial services are more likely than others to say their companies have adopted AI in risk functions. Telecom companies are most often adopting AI in service operations, while companies in the pharmaceutical industry tend to apply AI in product development and manufacturing. Respondents in consumer-packaged goods, travel and logistics, and retail are the most likely to report adoption of AI in supply-chain management.

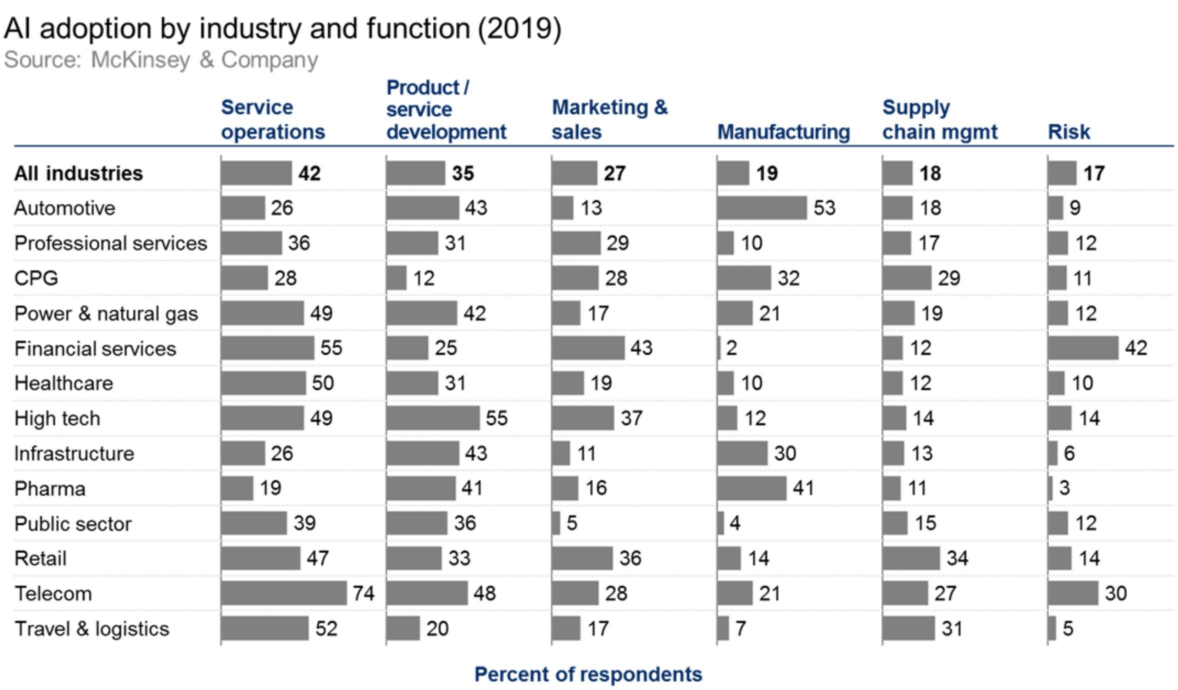

Fig. 4.3.2.

Industry Adoption

The AI capabilities that organizations adopt differ significantly by industry

Across industries, respondents are most likely to identify robotic process autmation, computer vision, and machine learning as capabilities embedded in standard business processes within their company (Figure 4.3.3). However, the capabilities adopted vary substantially by industry.

For example, natural language capabilities—including both understanding and generation of natural language text and speech—are adopted most often in industries with large volumes of customer or operational data in text form, including high tech, telecom, retail, financial services, and healthcare. By contrast, physical robotics is most frequently adopted in industries where manufacturing or transport of physical goods plays an important role in the supply chain, including automotive, consumer packaged goods, and pharma.

AI capabilities embedded in standard business processes (2019)

Source: McKinsey & Company

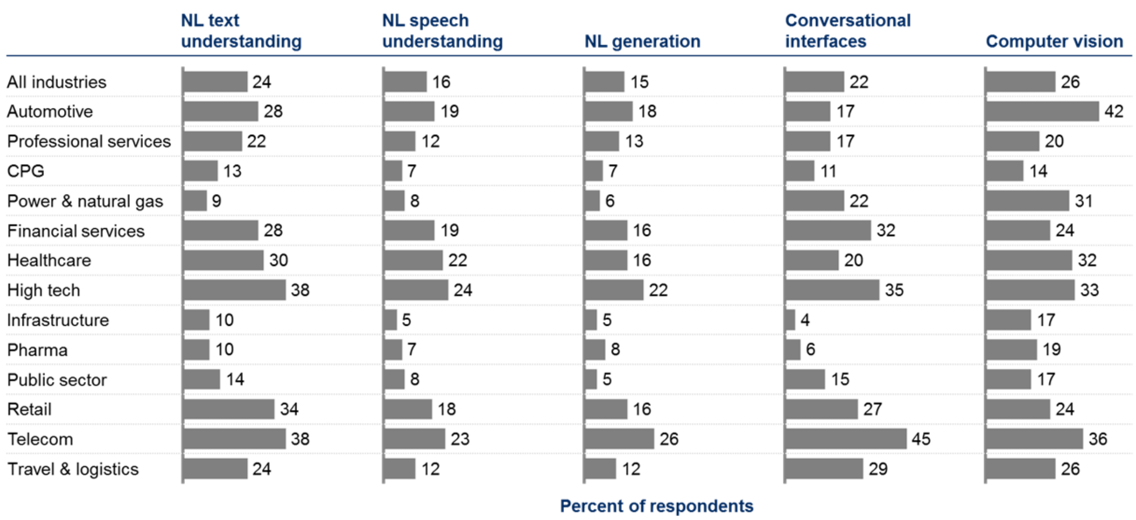

Fig. 4.3.3a.

AI capabilities embedded in standard business processes (2019)

Source: McKinsey & Company

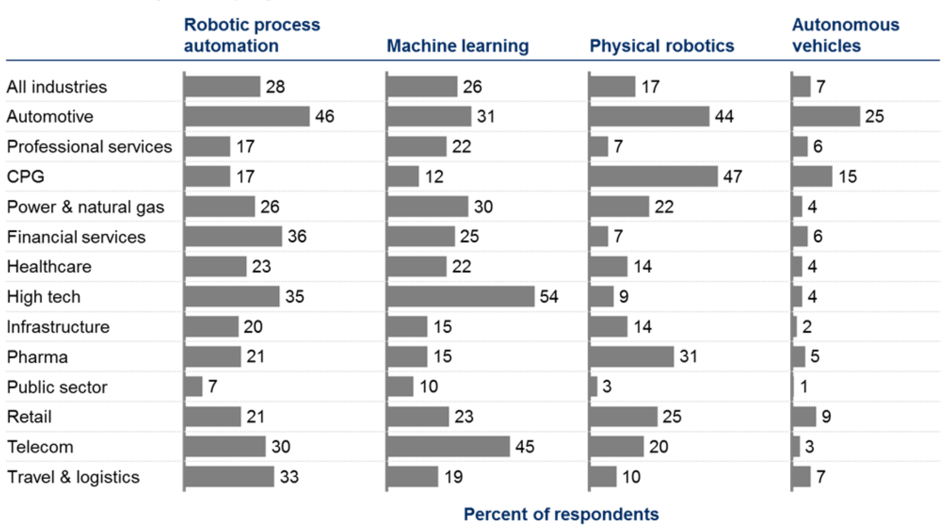

Fig. 4.3.3b.

Industry Adoption

Many companies applying AI do not report taking steps to mitigate the risks

McKinsey's study surveyed respondents on ten of the most widely recognized risks related to AI, including regulatory compliance, equity and fairness, cybersecurity, and personal and individual privacy.

Cybersecurity is the risk respondents most often say their companies are mitigating, cited by 48 percent of respondents from companies that have adopted AI. Thirty-five percent say their organizations are taking steps to mitigate risks associated with regulatory compliance, and three in ten say the same about personal and individual privacy.

Despite growing recognition of the importance of addressing ethical concerns associated with usage of AI, only 19 percent of respondents say their organizations are taking steps to mitigate risks associated with explainability of their algorithms, and 13 percent are mitigating risks to equity and fairness, such as algorithmic bias and discrimination (Figure 4.3.4).

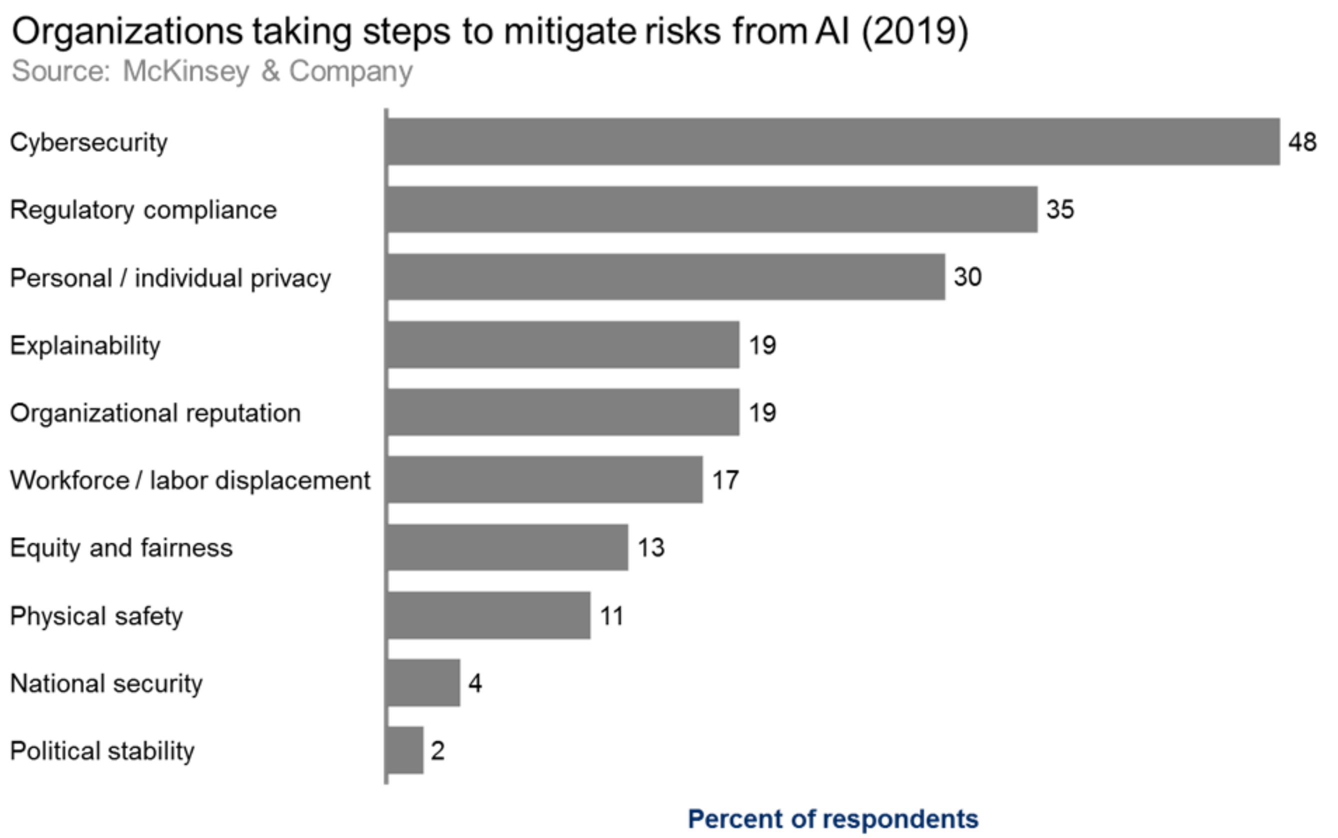

Fig. 4.3.4.
Note: Respondents who said "don't know / not applicable" are not shown.

Robot Installations

The graphs below show annual installations of industrial robot units for the world (Figure 4.3.5). In 2018, global robot installations increased by 6% to 422,271 units, worth USD 16.5 billion (without software and peripherals). The International Federation of Robotics (IFR) computed the operational stock of robots at 2,439,543 units (+15%). The automotive industry remains the largest customer industry with 30% of total installations, ahead of electrical/electronics (25%), metal and machinery (10%), plastics and chemical products (5%) and food and beverages (3%).[11] As mentioned in earlier AI Index Report, the numbers do not provide any indicator on how many of the systems actually use any means of AI, however they provide a measurement of installed infrastructure susceptible of adopting new AI technologies.

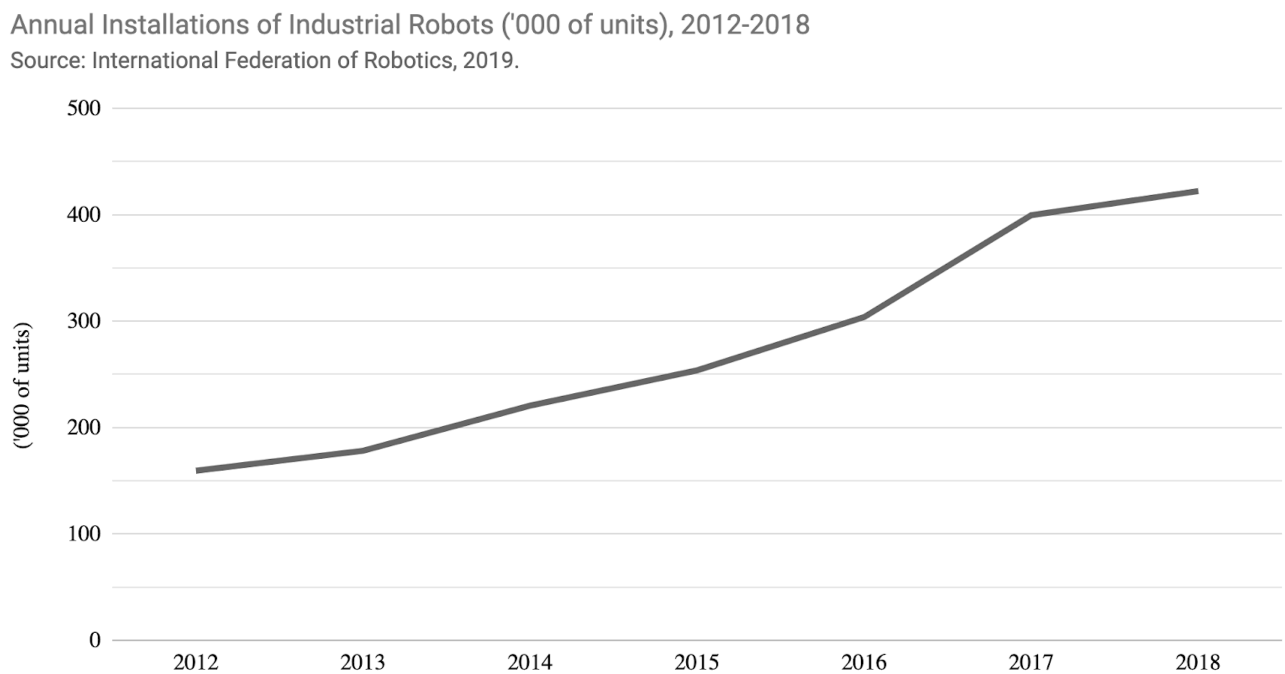

Fig. 4.3.5.

Global Robot Installations in 2018 more than 400,000 units

[11] Note that for almost 20% of the robots there is no information on the customer industry.

Robot Installations

The five major markets for industrial robots are China, Japan, the United States, the Republic of Korea, and Germany (Figure 4.3.6). These countries account for 74% of global robot installations. Since 2013, China has been the world's largest industrial robot market with a share of 36% of total installations in 2018. In 2018, 154,032 units were installed. This is 1% less than in 2017 (156,176 units) but still more than twice the number of robots installed in Europe and the Americas together (130,772 units). The main industries using robots in China are Electronics, Automotive & Metals, and the main application areas for industrial robots are handling and welding. Collaborative robots remain a small share compared to traditional industrial robots (Figure 4.3.7).

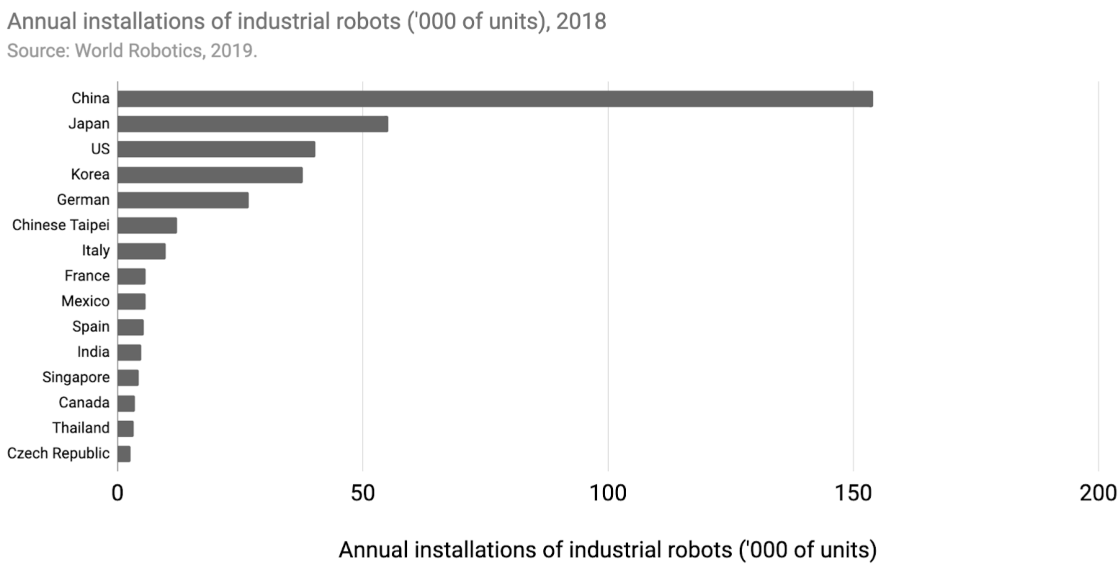

Fig. 4.3.6.

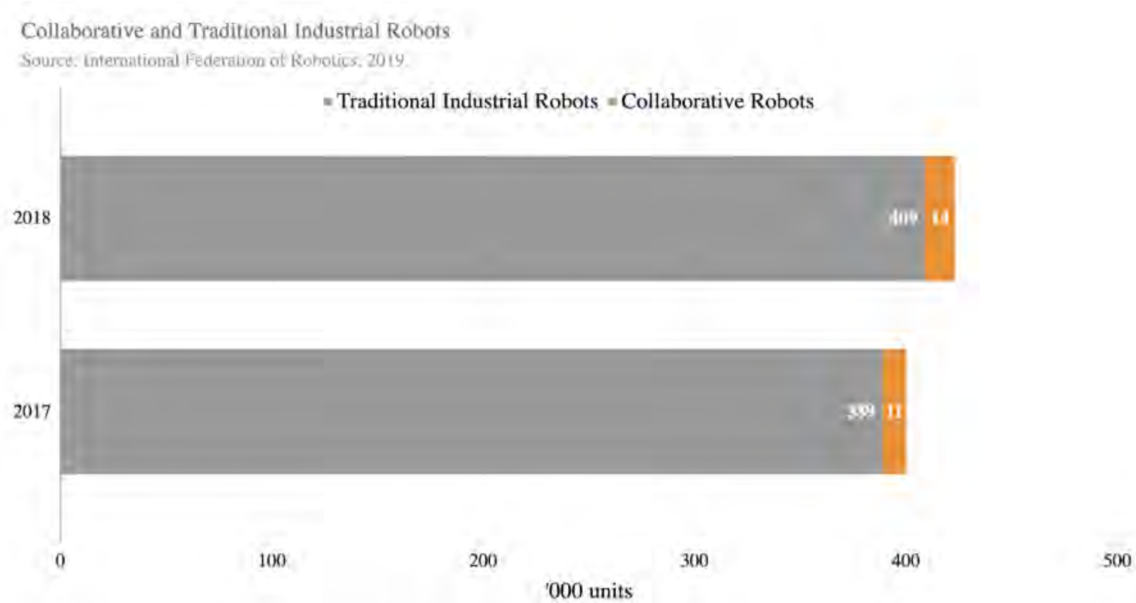

Fig. 4.3.7.

74% of global robot installations concentrated in five countries

Measurement Questions

- Additional firm-level data would be helpful to understand the impacts of AI adoption on firm performance. It would also be valuable to measure the availability and concentration of inputs for AI applications, including data available to countries or to firms, compute power, and talent, to improve understanding of the impact on competition and market power.

- From an economic lens, it would be invaluable to understand the AI components of robotics. Equally important are national and international statistical data on trade flows (imports and exports) of industrial versus service robotics, as a sector in labor force and enterprise surveys. There is also a need to understand the income inequality consequences of robotic automation.

- From a technical performance perspective, it would be essential to measure progress in specific robot tasks (from elementary to complex tasks) in a standardized manner. As observed by Rodney Brooks in the 2018 AI Index Report many sources quote industrial robot shipments that have very little (or no) AI in them, which makes it a poor metric for progress in AI. It could be interesting to look at robots which have an AI component, such as drones (which use SLAM, and other AI algorithms) distinct from home robots such as Roomba, that also have an AI components. Could we identify AI components in distinct robotic systems, and associated failure rates, in addition to their global adoption?

Chapter Preview

Chapter 5: Education

[Table_of_Contents]
[Education_Technical_Appendix]

Introduction

This chapter presents trends in AI education from a variety of data sources, starting first with global data from Coursera and Udacity ML and AI training courses. Second, trends in undergraduate enrollment in introductory ML and AI courses are presented for the US and international universities. Programs from European countries are also identified based on data from Joint Research Center, European Commission and the trends in AI PhD specialization for North America based on the CRA Taulbee Survey. Third, trends in PhD hires on industry hiring, faculty hiring and faculty departures are presented based on the Taulbee Survey and Goffman and Jin (2019). Fourth, trends in gender and international diversity for AI PhDs are presented, along with faculty diversity across select university departments. Included here is a short discussion on ethics courses in computational programs.

It is important to note that there are many other kinds of diversity. The Index continues to gather more numbers on underrepresented minorities, gender minorities, and other groups for 2020.

Coursera

Online Learning

Increasingly, AI education extends beyond the brick and mortar university. Online learning plays a key role in educating and developing AI skills in the workforce around the globe. Many questions arise about what skillsets students gain, where, and how they are meeting demands.

Coursera

Coursera, the world's largest online platform for higher education, serves over 45 million learners around the world by providing access to high quality content from leading universities and companies. The scale of the platform, which includes 3,700+ courses, 400+ specializations, and 16 degrees, creates one of the largest skills databases as millions of learners take graded assessments ranging from multiple choice exams to programming assignments to peer reviewed projects that measure their skill proficiency.

The Coursera Global Skills Index (GSI) ddraws upon this rich data to benchmark 60 countries and 10 industries across Business, Technology, and Data Science skills to reveal skills development trends around the world.

Cousera measures the skill proficiency of countries in AI overall and in the related skills of math, machine learning, statistics, statistical programming, and software engineering. These related skills cover the breadth of knowledge needed to build and deploy AI powered technologies within organizations and society:

•**Math:** the theoretical background necessary to conduct and apply AI research
•**Statistics:** empirical skills needed to fit and measure the impact of AI models
•**Machine Learning:** skills needed to build self learning models like deep learning and other supervised models that power most AI applications today
•**Statistical Programming:** programming skills needed to implement AI models such as in python and related packages like sci-kit learn and pandas
•**Software Engineering:** programming skills needed to design and scale AI powered applications

Below is a world heat map that shows the AI proficiency rankings of the 60 countries covered in the GSI (Figure 5.1). The map shows the quartile ranking category of each country denoted by cutting edge (76%-100%), competitive (51%-75%), emerging (26%-50%), and lagging (0%-25%). Details on the construction of these AI rankings is provided in the Technical Appendix along with a sample skills taxonomy that shows the breakdown of AI skills.

For each major geographic region, you can also see the average country's share of enrollments in AI and the five related competencies (Figure 5.2). The enrollment trends show that South Asia followed by East Asian countries tend to have a higher share of enrollments in AI and related skills.

Note that in terms of country size, there is not a strong correlation between number of users on Coursera and the skill rank of a country in AI. Rather the skill rank of a country correlates much more strongly with metrics like a country's GDP per capita and the level of investment in tertiary education. See this article for some plots. In addition, the rankings are robust to adjusting for self selection in using Coursera through propensity score weighting.

"The Fourth Industrial Revolution is upon us, foreshadowing massive changes to the nature of work. Without a concerted focus on skill development, the dislocations will be widespread and felt most acutely by the poorest and least educated. Keeping pace with the fundamental market shifts will demand coordinated investments in skill development — not just by individuals, but also by companies and governments around the world." —
Emily Glassberg Sands and Vinod Bakthavachalam (Coursera Data Science)
Harvard Business Review

Coursera

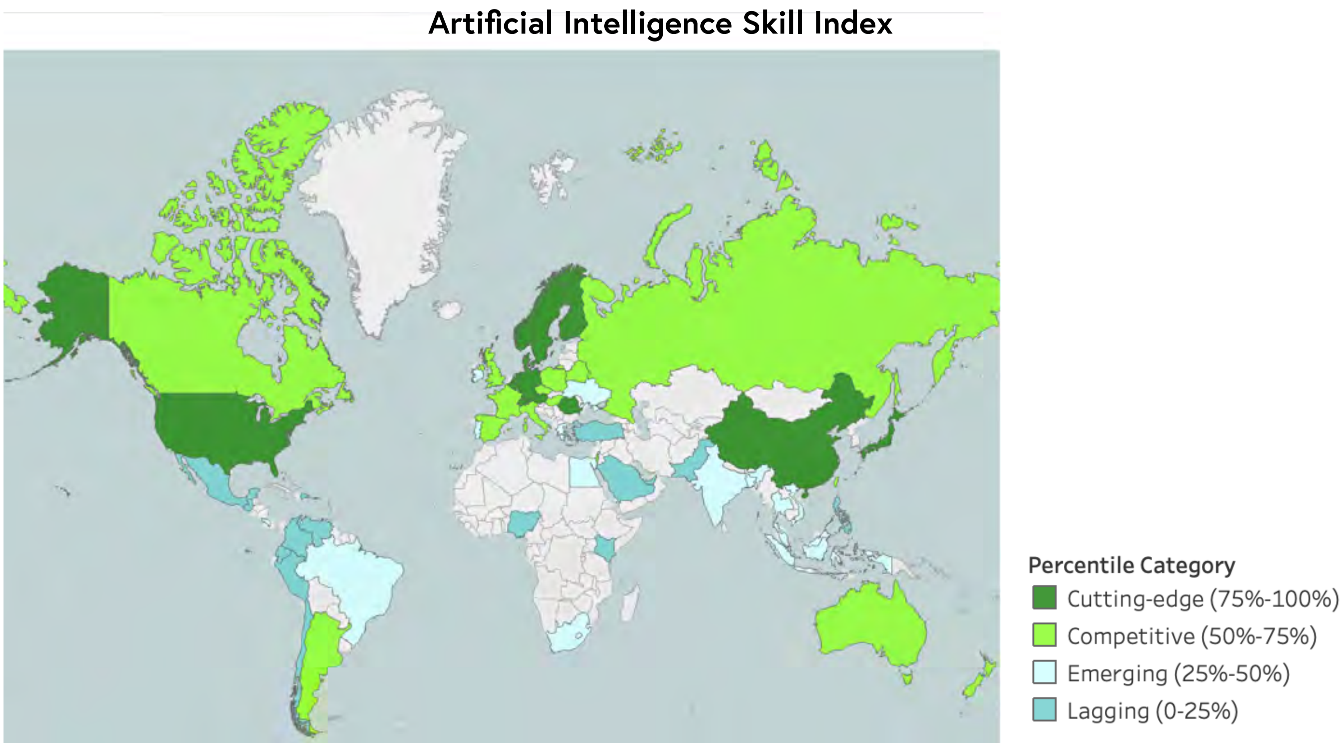

Fig. 5.1.

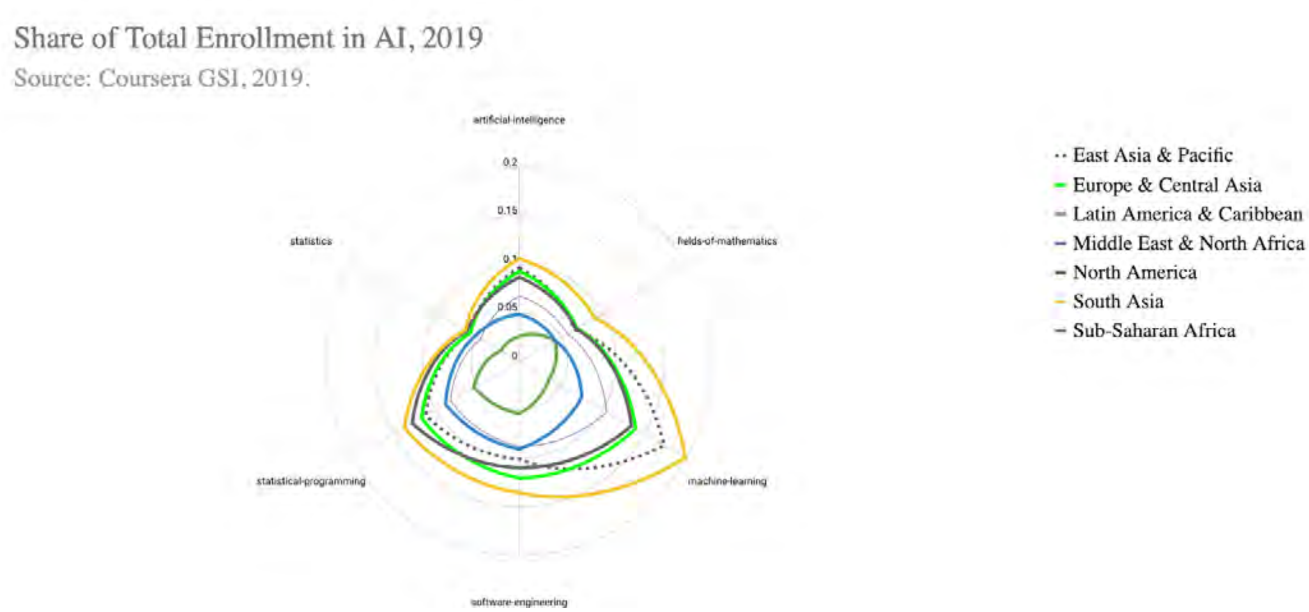

Fig 5.2.

Udacity

The enrollment in different AI specialization courses on Udacity is presented next (Figure 5.3). The chart shows the running total enrollment in the various AI specializations for Udacity AI specialization courses. *Introduction to TensorFlow for Deep Learning* has maintained the highest total enrollment till mid-2019. However, *Introduction to Machine Learning* has cumulatively the highest enrollment number in later 2019, with over 125,000 cumulative global enrollment. *Introduction to AI* is close behind, followed by more computer systems engineering topics such as *Introduction to Hadoop and MapReduce*.

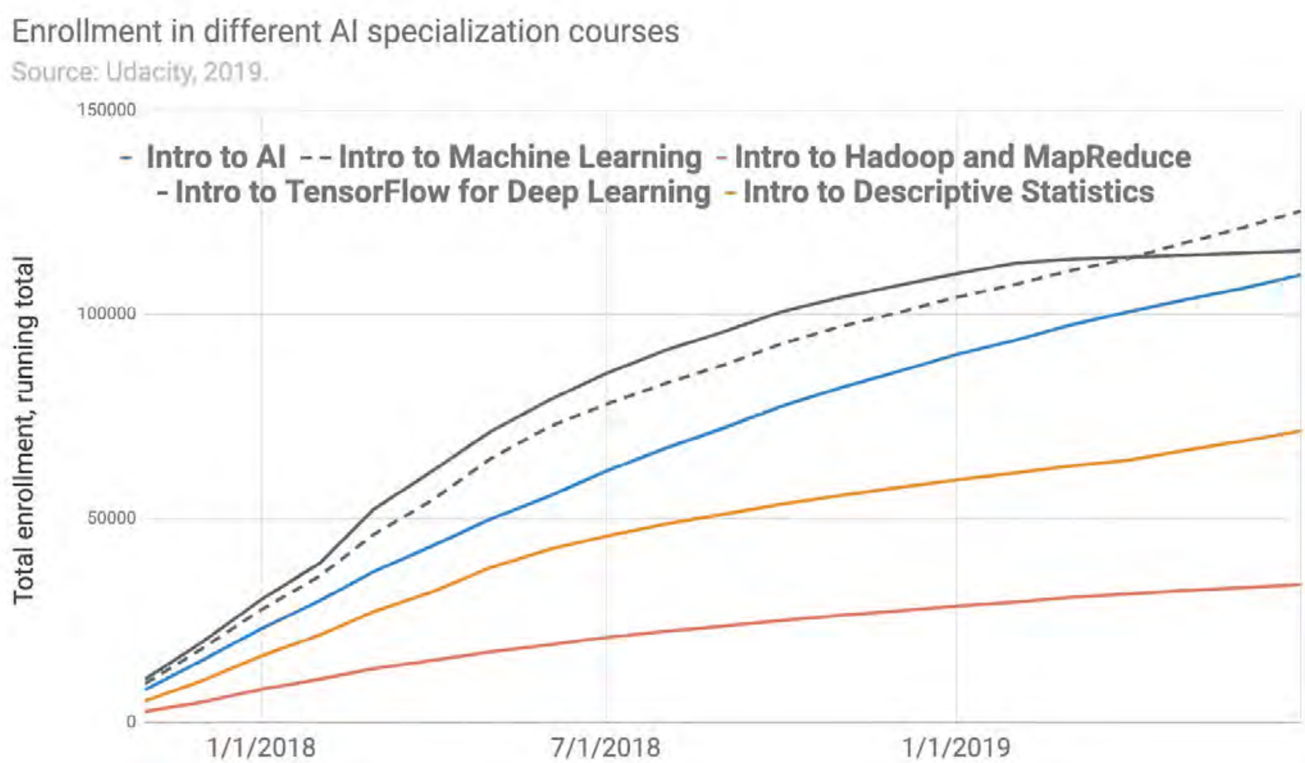

Fig. 5.3.

US Course Enrollment

The graphs below (Figures 5.4a & 5.4b) show the number of students enrolled in introductory AI and ML courses in a number of US universities. School selection criteria, actual enrollment numbers, and full university names can be found in the appendix. Enrollment in *Introduction to Artificial Intelligence* grew five-fold between 2012 and 2018 at Stanford University. Enrollment in *Introduction to Machine Learning* grew 12-fold between 2010 and 2018 at the University of Illinois at Urbana-Champaign (Figure 5.4c & Figure 5.4d). Some schools indicated that growth in enrollment was limited by availability of classes, so these graphs may underrepresent the real demand for these courses.

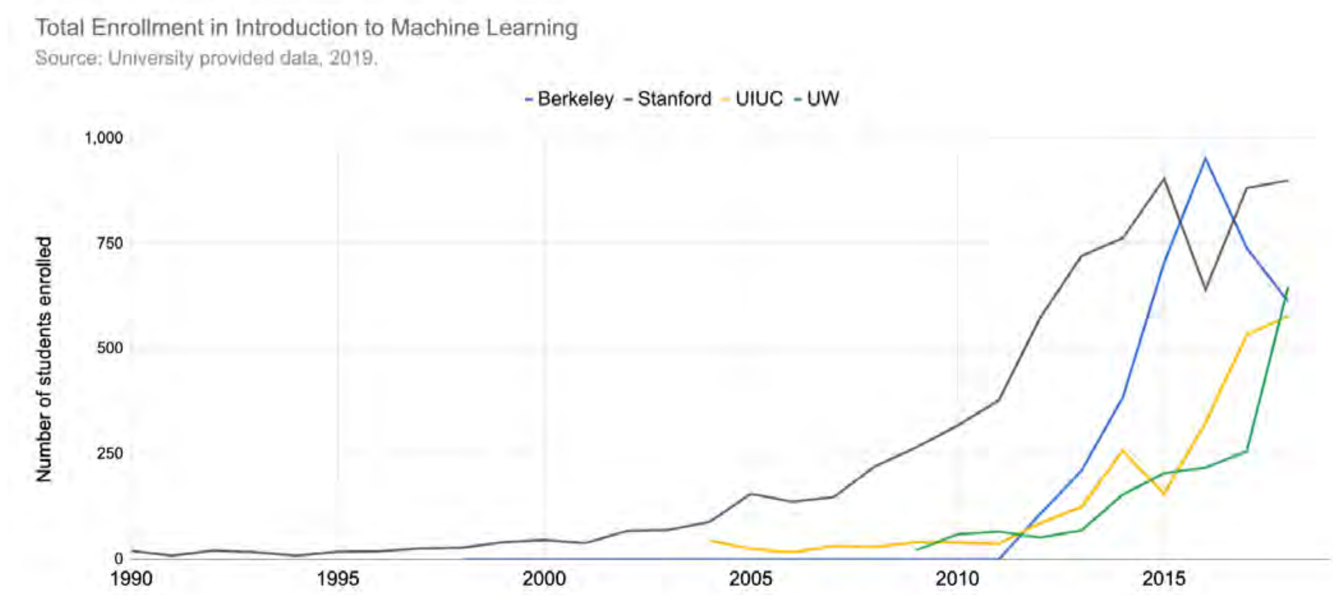

Fig. 5.4a.

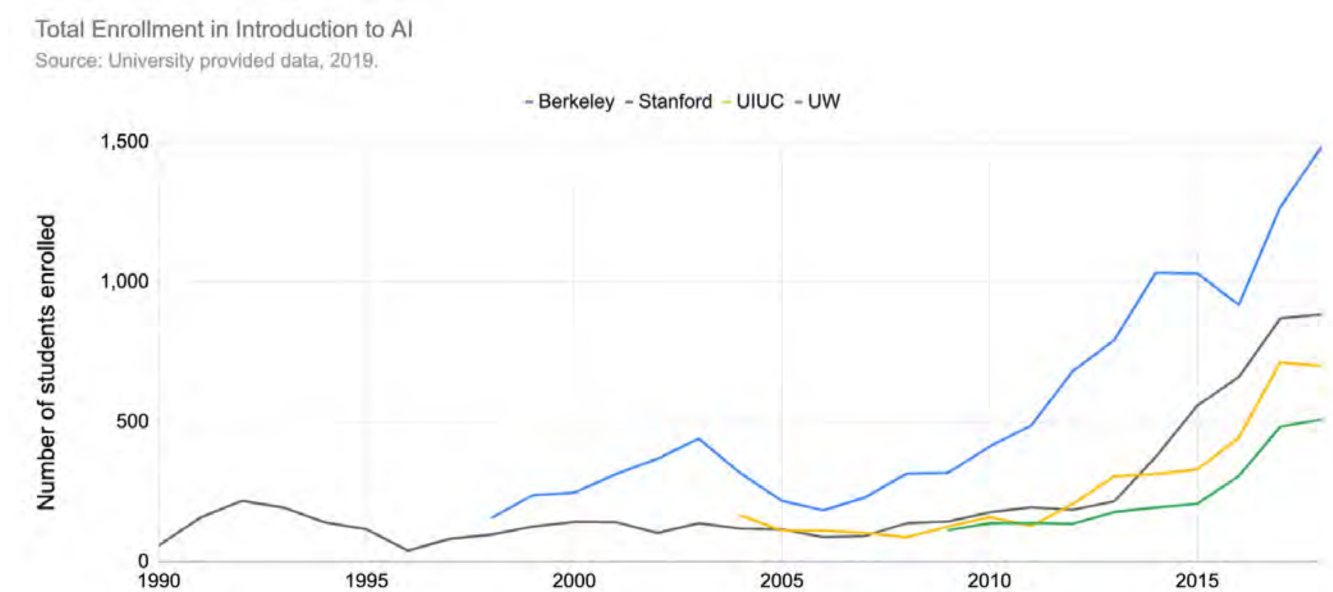

Fig. 5.4b.

US Course Enrollment

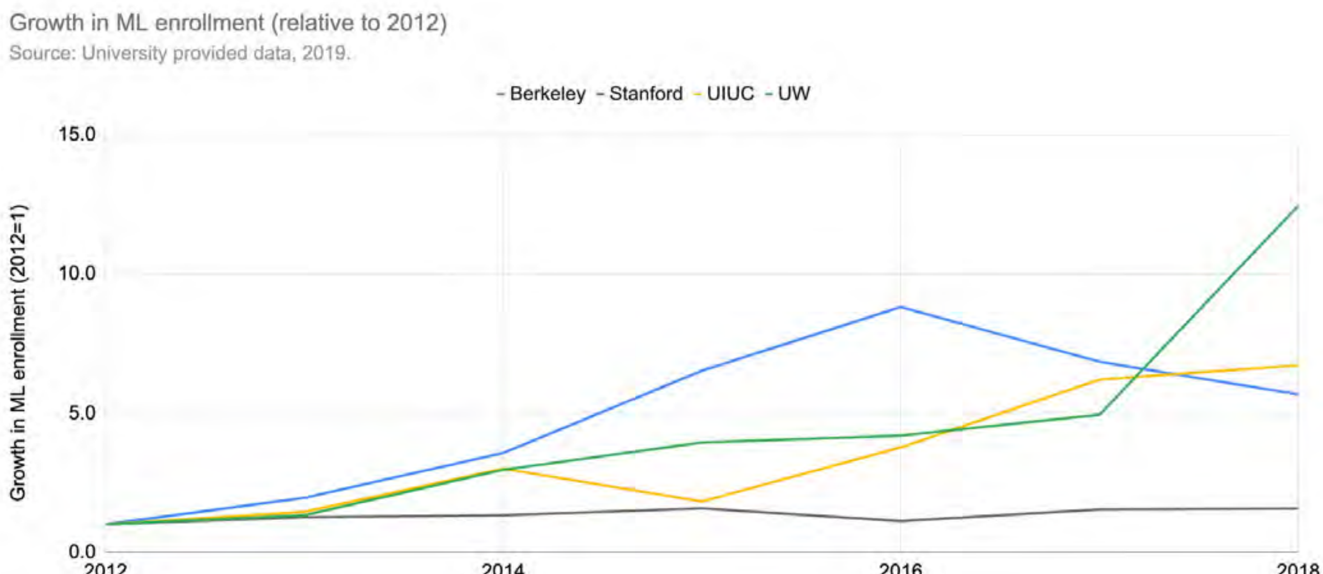

Fig. 5.4c.

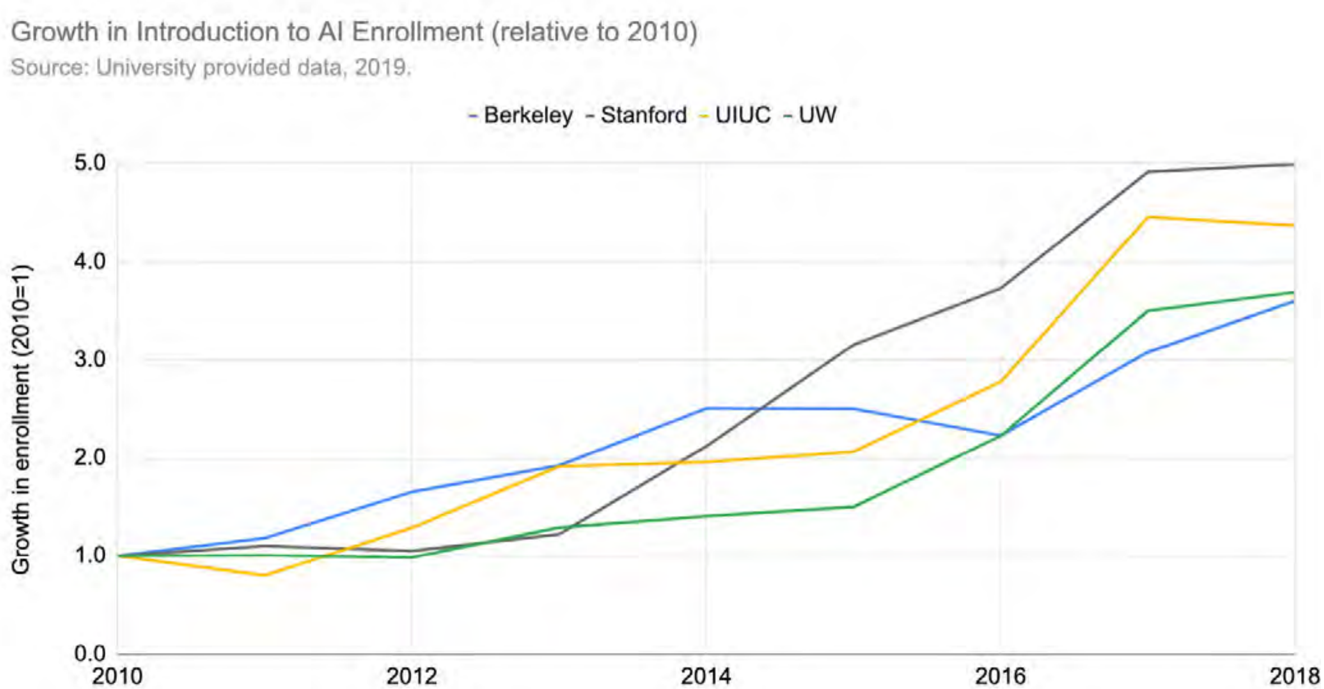

Fig. 5.4d.

International Courses

The graphs below (Figure 5.5a) show AI and ML course enrollment at several leading computer science universities outside of the US The graph shows relative growth for international schools that provided data for academic years 2010 — 2019. School selection criteria, actual enrollment numbers, and full university names can be found in the appendix. In the given sample, the University of Toronto (Canada) has the highest number of registered students for Introduction to AI+ML, followed by High School of Economic (Russia), and Tsinghua University (China) in 2018. Relative to 2015, enrollment has grown four-folds at Tsinghua University, three-folds at University of Toronto, and doubled at University of Melbourne (Figure 5.5b).

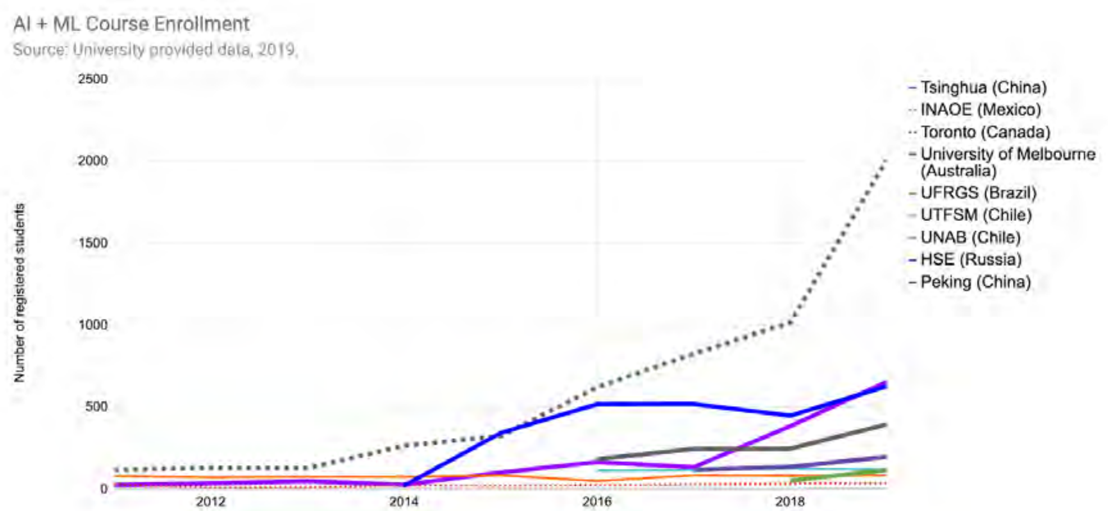

Fig. 5.5a.

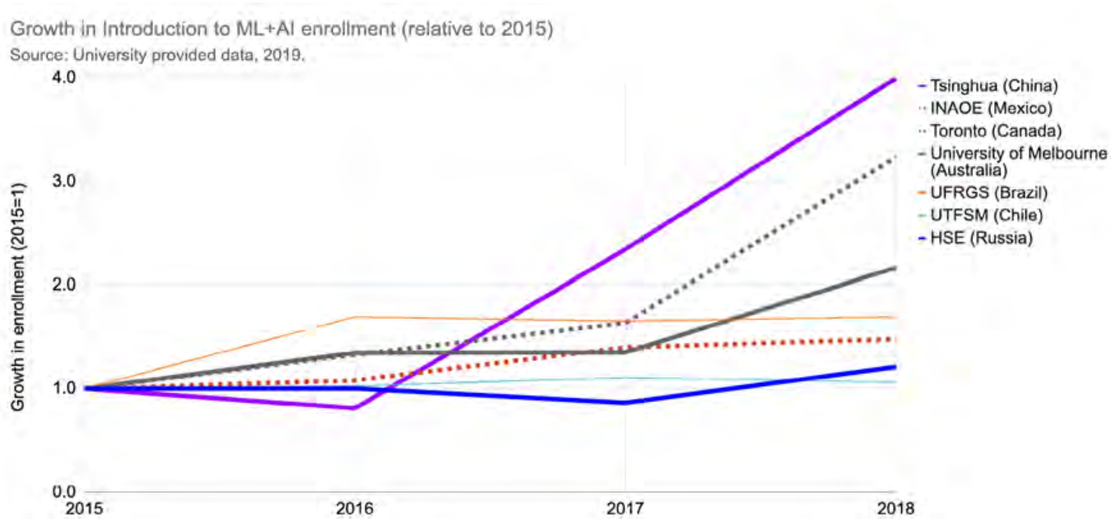

Fig. 5.5b.

Across the schools studied, we found that growth in AI course enrollment was relatively school dependent, and was not particularly influenced by geography. The AI Index looks forward to refining this hypothesis in future reports.

Trends From Europe

Text mining and machine learning techniques were applied to all universities across Europe that have a website (as listed by the Webometrics initiative). The data related to the programs of study address the domains that have been identified by the Joint Research Centre (JRC), the science and knowledge service of the European Commission (EC). The data collection effort identified a suitable term of comparison when considering third party sources, to measure strengths and weaknesses of a (semi) automatic classification system for program content. Readers can refer to Academic offer and demand for advanced profiles in the EU for more technical details.

This data (Figure 5.5c) identified a total number of 2,054 programs covering the domain of Artificial Intelligence to differing extents. The vast majority of AI academic offerings in Europe are taught at the masters level, as the MS is the expected terminal degree and generally perceived as the most appropriate to acquire the needed advanced skills. The graph (Figure 5.5d) shows that there are 197 European universities offering a total of 406 specialized masters in AI; 84 of the universities, or 43%, offer at least 2 specialized masters in AI. Programs have been classified, depending on the level, into bachelors and masters. Though not exhaustive, the selected data source offers a perspective on the academic offerings targeting the selected domains in EU28.[12]

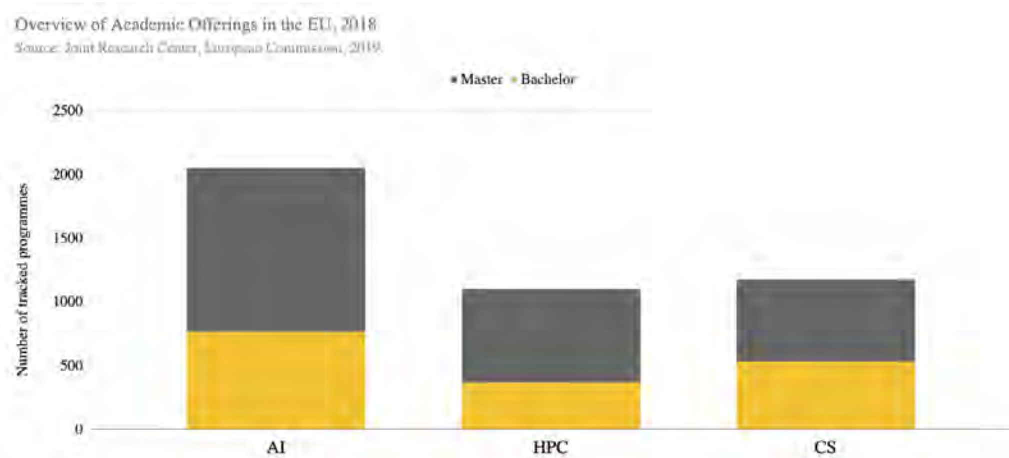

Fig. 5.5c.
Note: The total number of programmes in the selected domains does not correspond to the sum of programmes in each domain due to the fact that a programme may correspond to more than one domain.

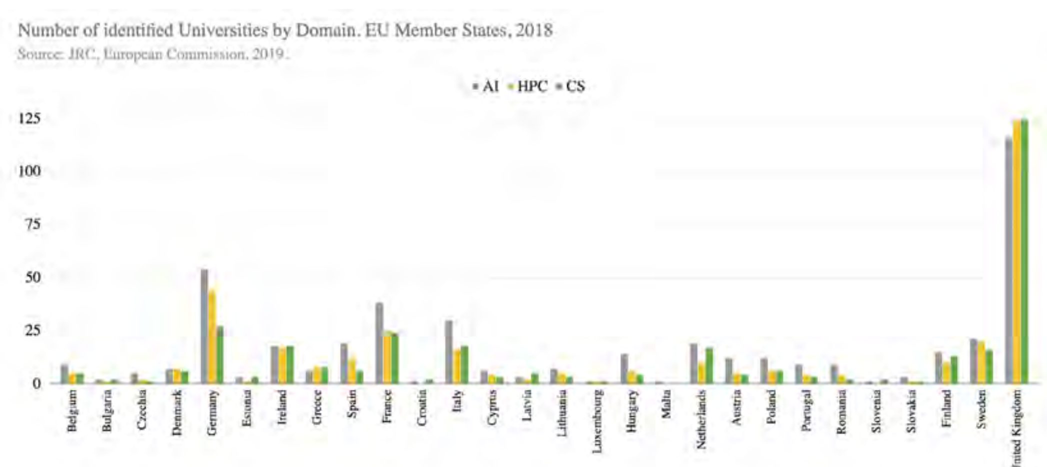

Fig. 5.5d.

[12] United Kingdom leads both in number of companies and of programmes offered by universities, hosting one third of AI companies and more than half of AI programmes. In 2016, countries employing highest number of ICT specialists were United Kingdom (1.7 million persons), Germany (1.5 million), France (1.0 million), Italy (721 thousands) and Spain (632 thousands).

PhD Specialization in AI

The Computing Research Association's (CRAs) Taulbee Survey is conducted annually to document trends in student enrollment, degree production, employment of graduates, and faculty salaries in academic units in the US and Canada that grant the Ph.D. in computer science (CS), computer engineering (CE), or information (I). Only doctoral departments of computer science and computer engineering are included. Historically, Taulbee has covered 1/4 to 1/3 of total BS CS recipients in the US. The categorization of specialty areas changed in 2008 and was clarified in 2016. From 2004-7, AI and Robotics were grouped; since 2008, AI has been separate; in 2016 AI also included ML.

The first chart (Figure 5.6a) shows AI/ML PhD grad specializations as a percent of computing PhD graduates in the US (and the number of AI/ML graduating PhDs). It is more difficult to estimate the growth in AI/ML undergraduate specialization, but the appendix chart shows undergraduate enrollment in CS is over 130,000 in 2018.[13] The specialization of computing PhDs is presented next. The bar chart (Figure 5.6b) shows (a) the share of computing PhD grads in 2018 by areas of specialization, and (b) the changes in share of each specialization between 2010-18. AI is the most popular PhD specialization for computing PhD grads and continues growing the fastest. In 2018, over 21 percent of graduating computing PhDs specialize in Artificial Intelligence/ Machine Learning.

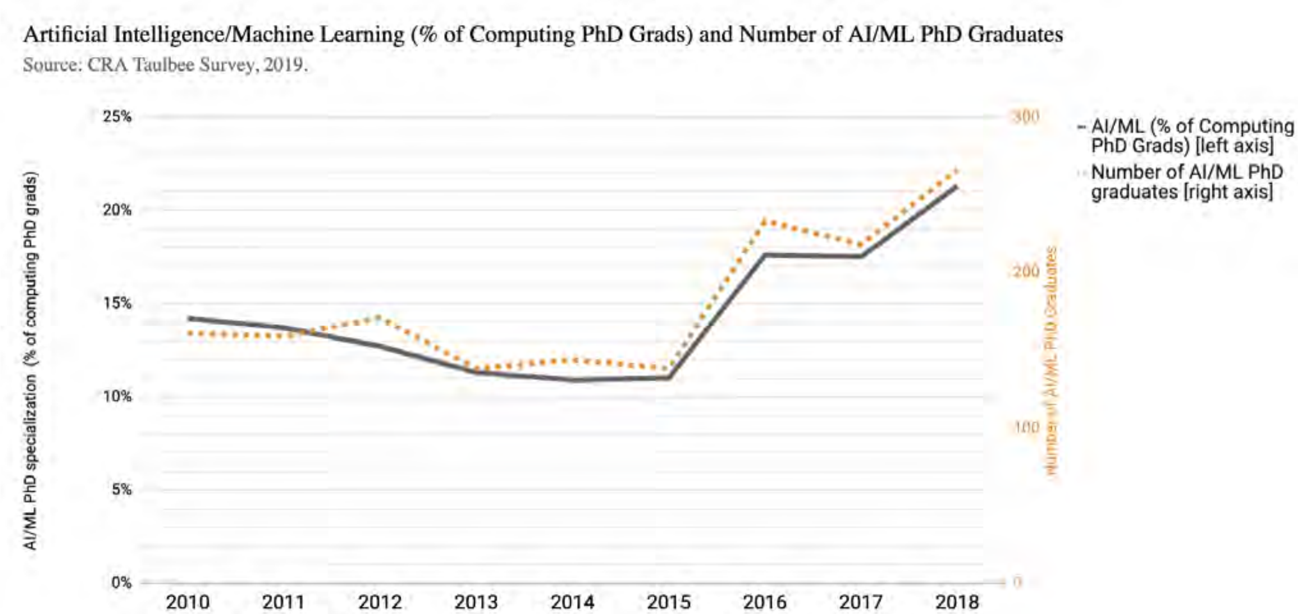

Fig. 5.6a.

AI is the most popular area for CS PhD Specialization. In 2018, over 21 percent of graduating computing PhDs specialize in Artificial Intelligence/Machine Learning.

[9] The number of students entering undergraduate enrollment (~34,000) exceed the number of undergraduates graduating (~27,000) in 2018. The growth in the number of students starting undergraduate studies in CS is growing the fastest, growing 4-fold since 2006.

PhD Specialization in AI

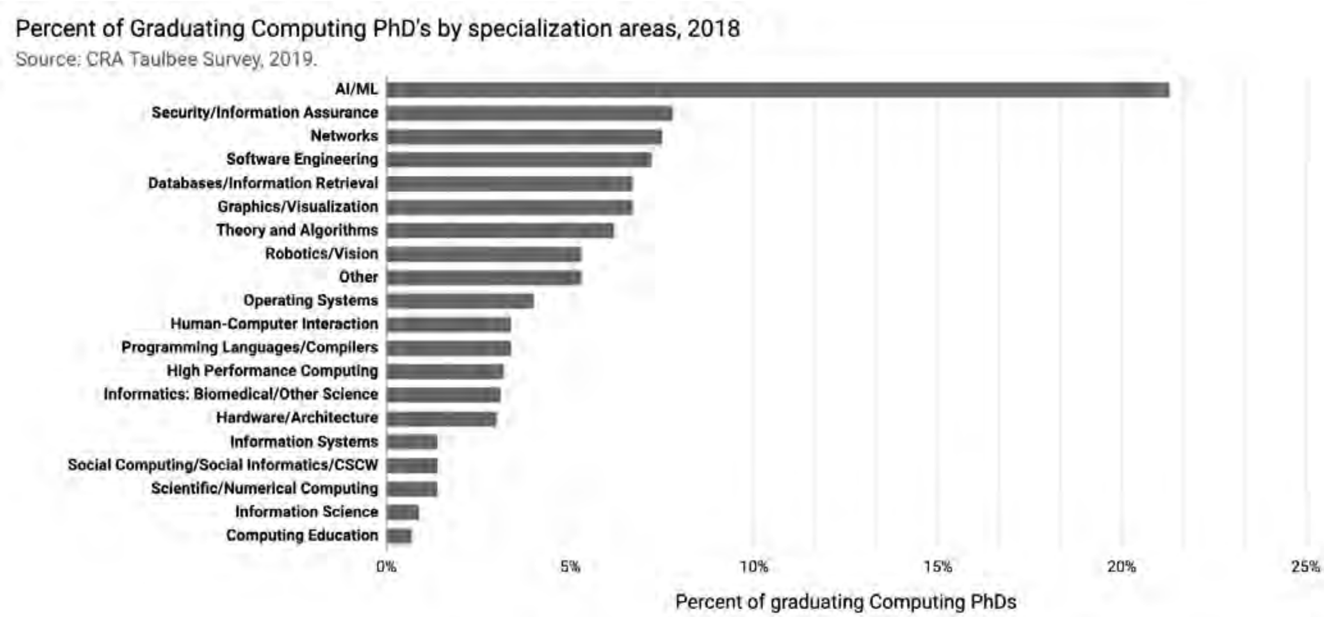

Fig. 5.6b.

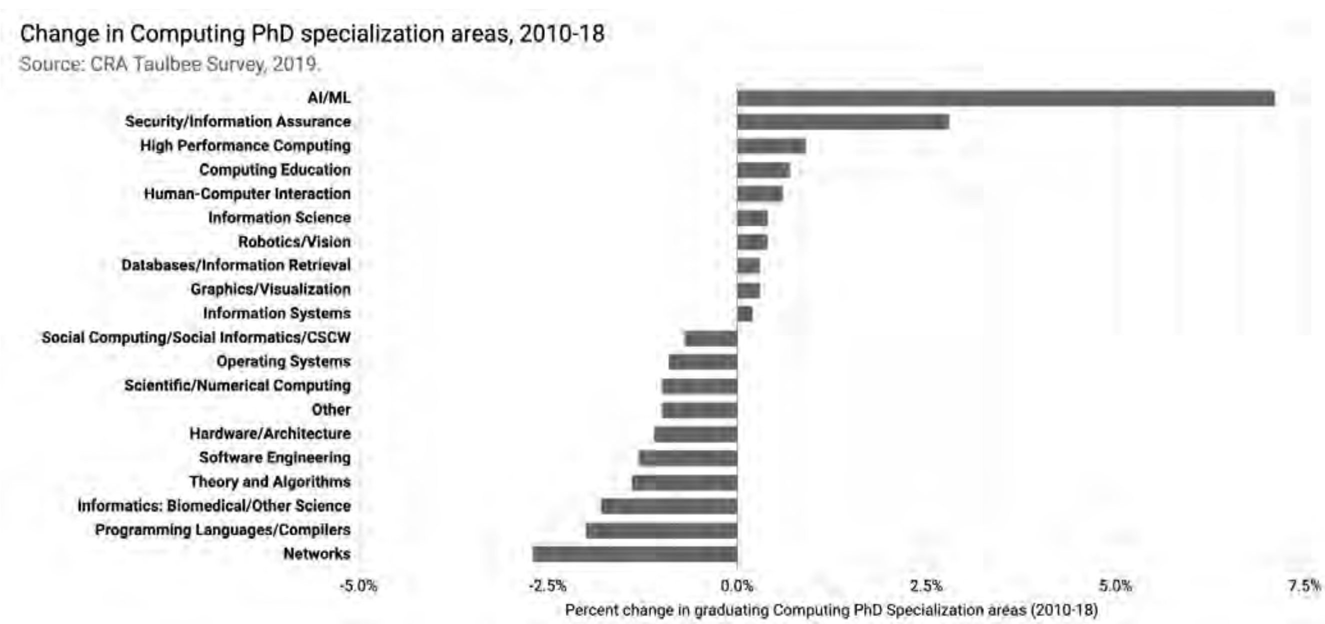

Fig. 5.6c.

PhD to Industry

Over 150 new AI PhDs went to industry in 2018, and this number represents a percentage of new graduates three times as large as 2004 (Figures 5.7a & 5.7b). The percent of graduating AI PhDs going to industry increased from 21% in 2004 to over 62% in 2018. It should be noted that in many fields in academia there is no expectation that every PhD student goes on to get an academic job. For example, in the life and health sciences, the fields that award the most Ph.Ds, only 23% of PhDs held a tenured or tenure-track position in academia in 2017 (see Science, 2019).

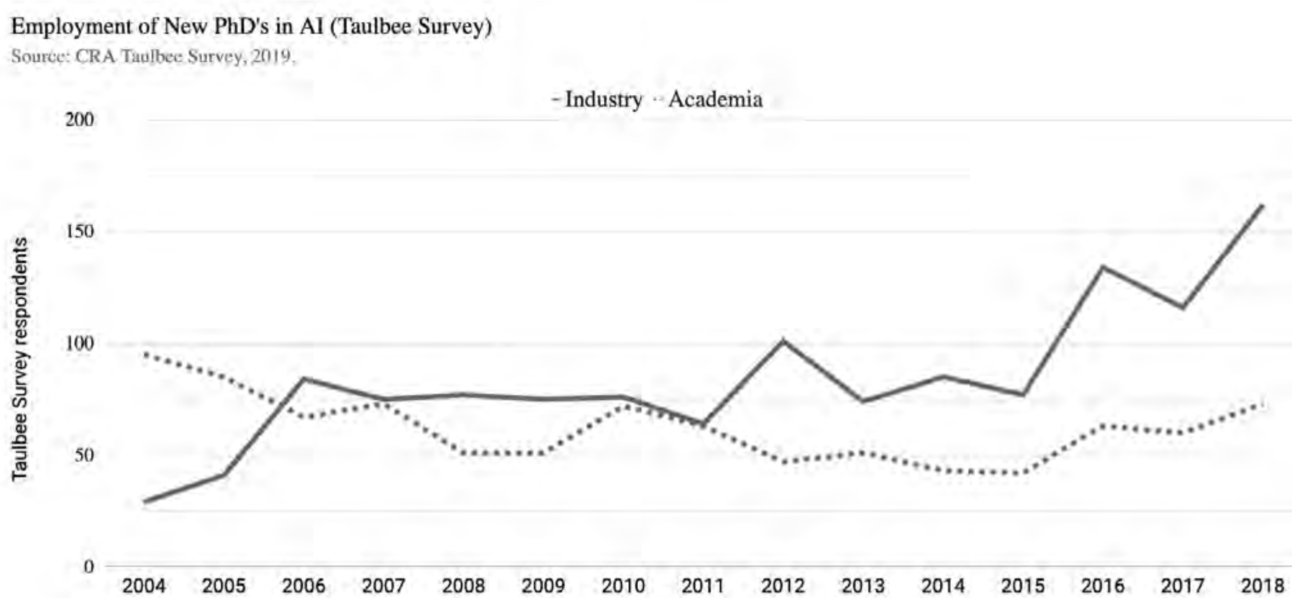

Fig. 5.7a.

Note: Categorization of specialty areas changed in 2008 and was clarified in 2016. 2004-7, AI and Robotics were grouped; 2008-present AI is separate; 2016 clarified to respondents that AI included ML.

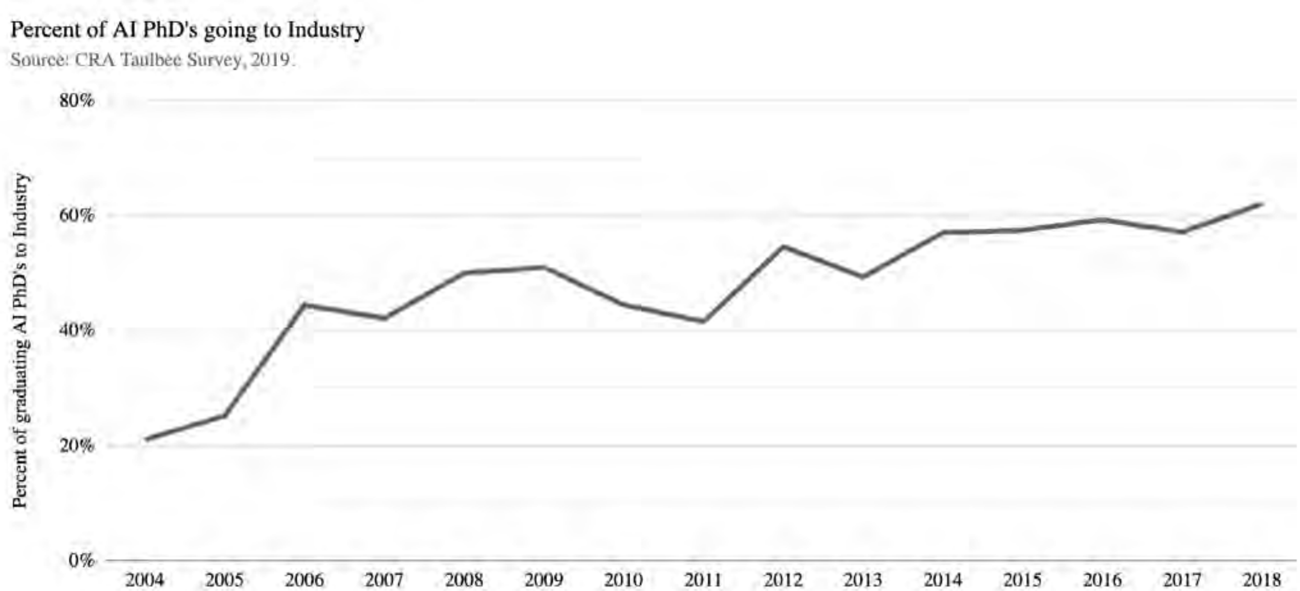

Fig. 5.7b.

Faculty Hires

The trends in new faculty hires are presented next (Figures 5.8a, 5.8b & 5.8c). The 2018 Taulbee survey asked for the first time how many new faculty hires came from the following sources: new PhD, postdoc, industry, and other academic. 29% of new faculty hires came from another academic institution. Some may have been teaching or research faculty previously rather than tenure-track, and there is probably some movement between institutions. Thus, the total number hired overstates the total who are actually new to academia.[14]

The total number of CS tenure-track faculty has been rising steadily, making up half of the faculty hiring pool (Figure 5.8a). The percent of new female tenure-track faculty has remained largely constant at slightly over 21%. The percentage of new faculty who are international is smaller, at around 18% (Figure 5.8b). The last chart (Figure 5.8c) shows that although most new AI PhDs do a postdoc, the portion going directly tenure-track positions is increasing.

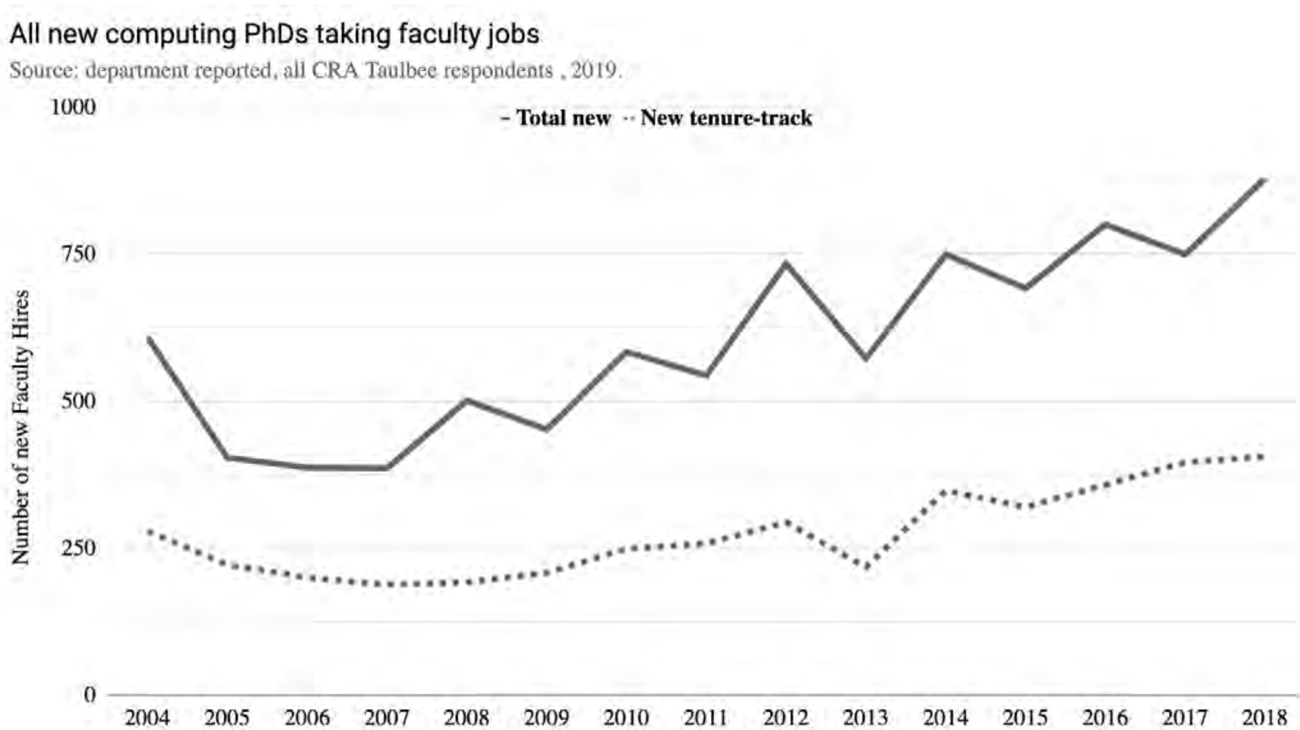

Fig. 5.8a.

[14] If Professor Q leaves institution A for Institution B, and A hires his replacement from Institution C, who hires a replacement from Institution D, who hires a new PhD, 4 institutions will report new hires but there's only a total increase of 1 new faculty member.

AI Faculty Hiring

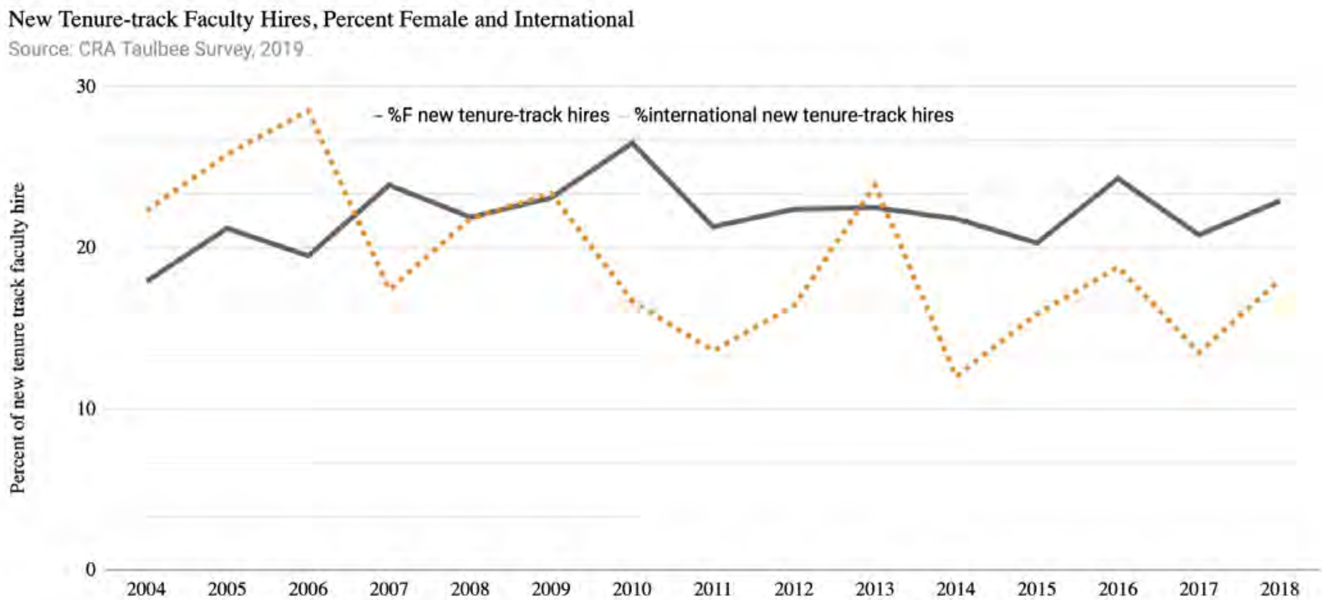

Fig. 5.8b.

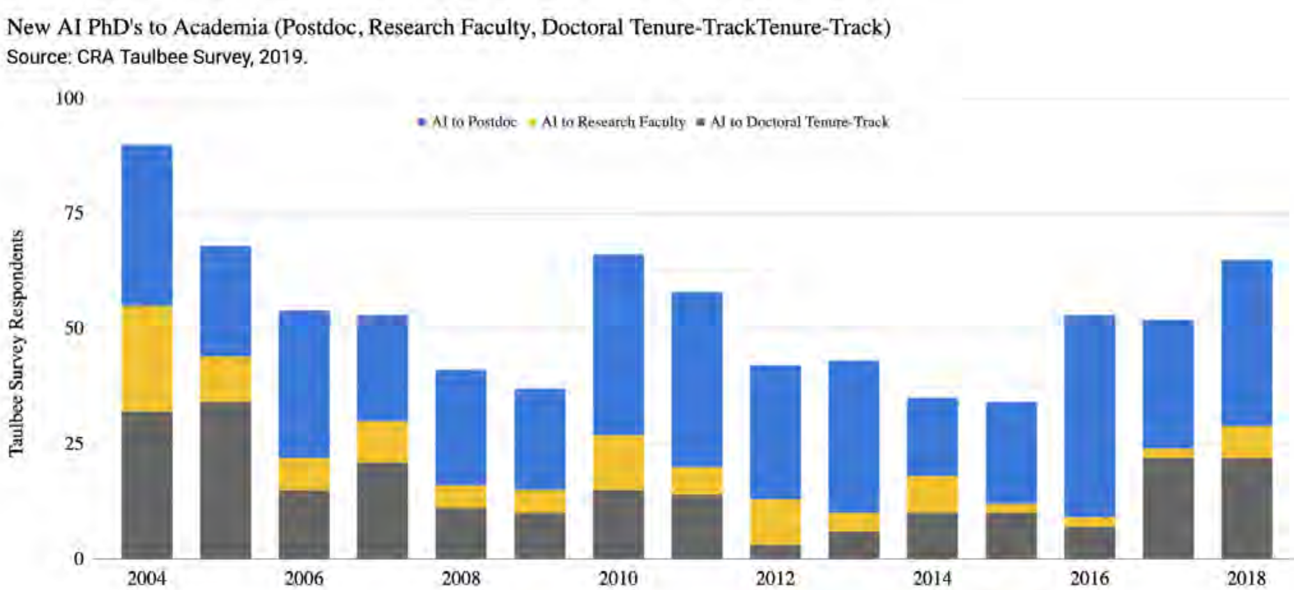

Fig. 5.8c.

Faculty Departures

Goffman and Jin (2019) document the brain drain of AI faculty to industry.[15] The first graph (Figure 5.9a) below shows the number of North American tenure-track professors in AI leaving each year for an industry job. The movement affects both tenured and untenured faculty. This next figure (Figure 5.9b) shows the 18 North American universities with the largest losses of AI-related tenure-track or tenured professors between 2004 and 2018. Some of them left the university completely and some still keep university affiliations while working for companies. The three universities that lost the most AI faculty are Carnegie Mellon University (CMU), the University of Washington, and UC Berkeley. CMU lost 17 tenured faculty members and no untenured faculty, and the University of Washington lost 7 tenured and 4 assistant professors. For Canadian universities in the sample, the University of Toronto lost the most AI professors, 6 tenured faculty and 3 assistant professors.

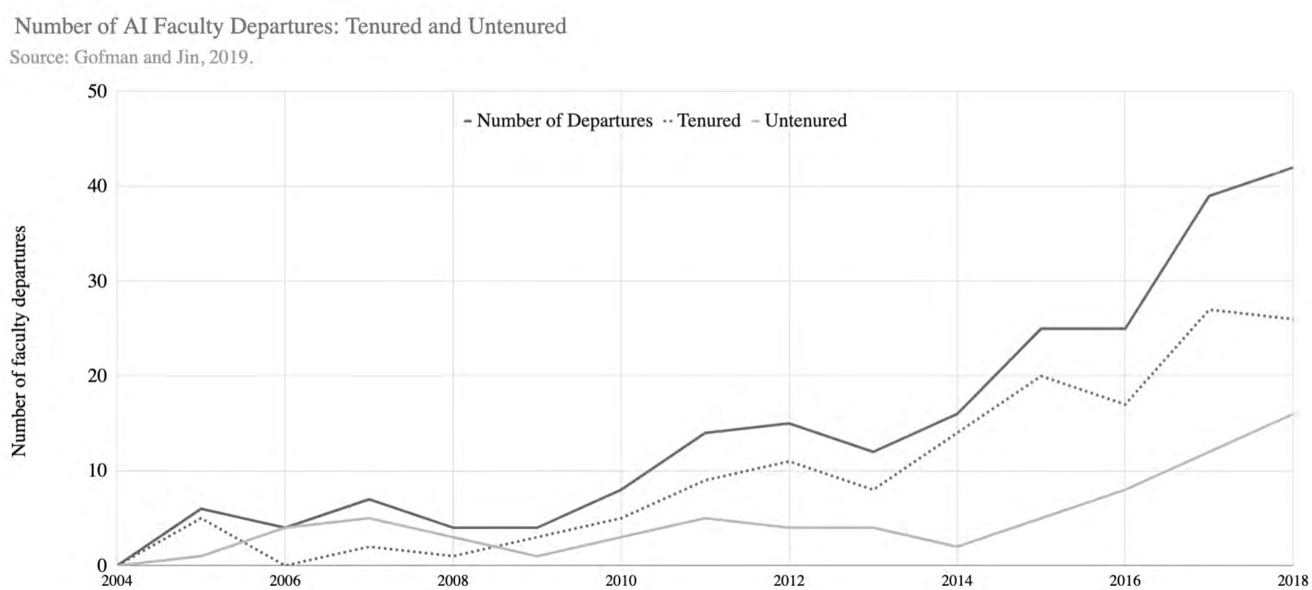

Fig. 5.9a.

"AI's emergence as a general-purpose technology has resulted in an unprecedented brain drain of AI professors from academia to industry. What are the consequences of this brain drain is an important policy question."
Michael Gofman, Assistant Professor of Finance, University of Rochester

[15] Gofman, M., and Z. Jin, (2019) "Artificial Intelligence, Human Capital, and Innovation", University of Rochester Working paper. This paper combines data from LinkedIn, CSRanking.com, CrunchBase, and Google Scholar. For AI professors leaving for an industry job is based on hand-collected sample from LinkedIn. The second method is to search in LinkedIn using reviewers' and program committee members' names of AI related conferences. Researchers also hand-collect data on faculty size at the top 100 universities' computer science departments from CSRankings.org, which provides the number of full-time, tenure-track and tenured CS faculty for each year based on data from DBLP Entrepreneurs'. Startups' information is based on a sample from the CrunchBase database Finally, hand-collected citation data from Google Scholar are used as a proxy for quality of research of AI faculty. Readers are referred for further technical details to the paper. The most updated AI brain drain index can be downloaded at http://www.aibraindrain.org

Faculty Departures

The Gofman and Jin paper also documents trends in AI startups founded by graduates from North American universities. Figure 5.9c shows the North American universities that produced the most AI entrepreneurs who received their highest degrees from these universities from 2004 - 2018 and who established AI startups thereafter.[16] In the sample, 77 MIT graduates, 72 from Stanford and 39 from Carnegie Mellon University established AI startups. The Canadian university with the most AI entrepreneur alumni is the University of Waterloo, with 21 such graduates.

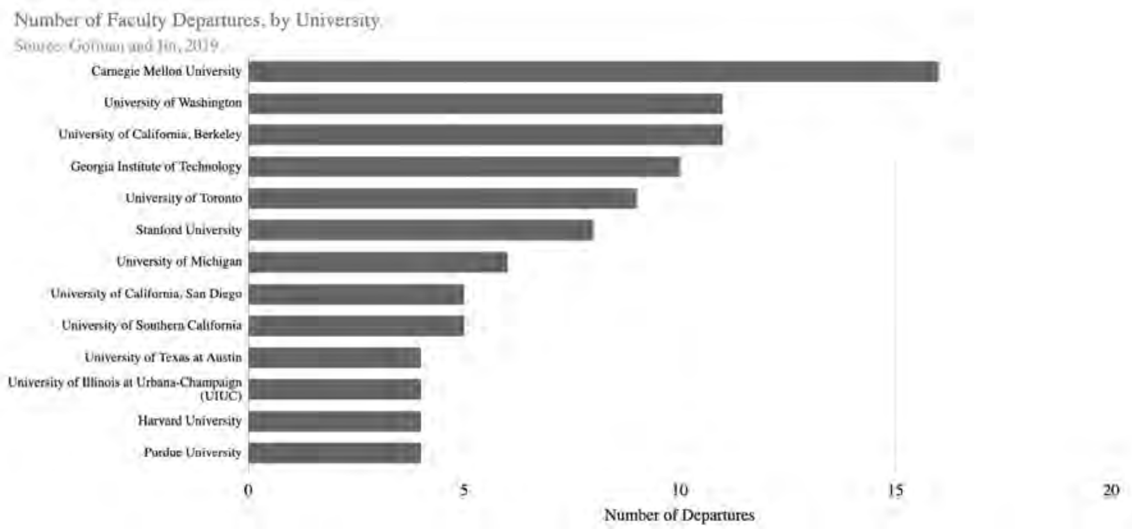

Fig. 5.9b.

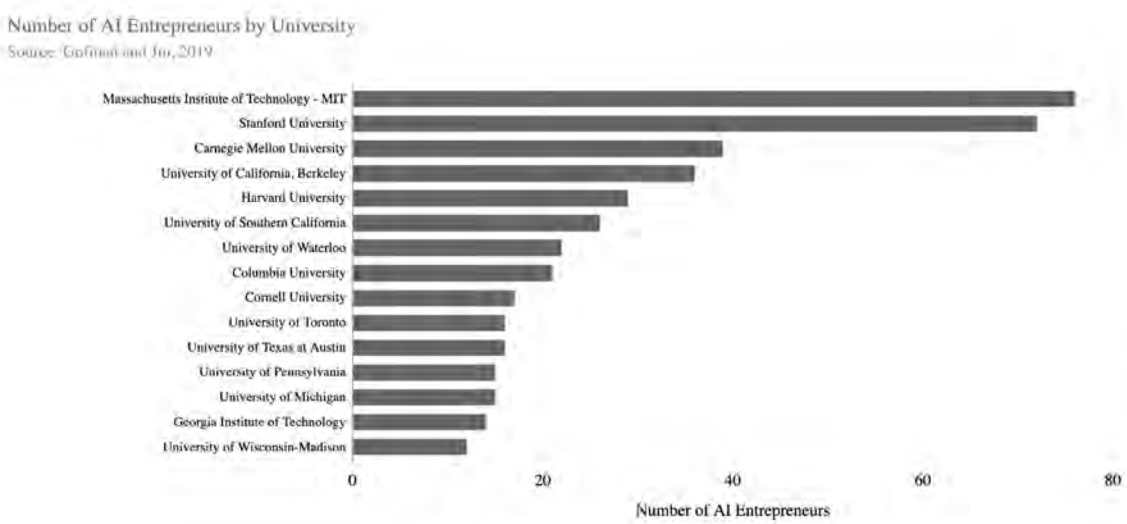

Fig. 5.9c.

"AI startups require significantly more domain-specific knowledge than non-AI startups. AI brain drain negatively affects students' ability to gain the essential knowledge they need to be successful AI entrepreneurs."
Zhao Jin, Finance PhD Candidate, University of Rochester

[16]An AI entrepreneur is identified if they start an AI startup after receiving their highest degree. AI startups are defined as startups that their business description includes one of the following fields: face recognition, neural networks, image processing, computer vision, semantic web, speech recognition, machine learning, natural language processing, artificial intelligence, deep learning, autonomous driving, autonomous vehicle, and robotics.

Women in AI

Figure 5.10a plots the percent of female AI PhD recipients in the US between 2010-18, which has remained stable at around 20%. Figure 5.10b shows that in 2018, the percentage of new women faculty hire in computation fields is slightly higher than the proportion of female graduating with AI or CS PhD.

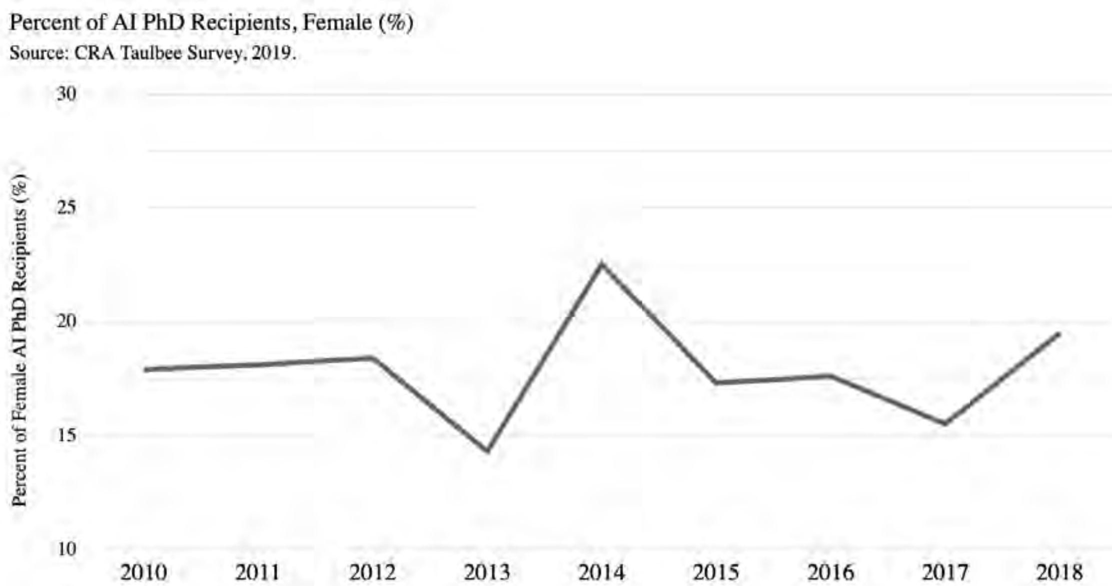

Fig. 5.10a.

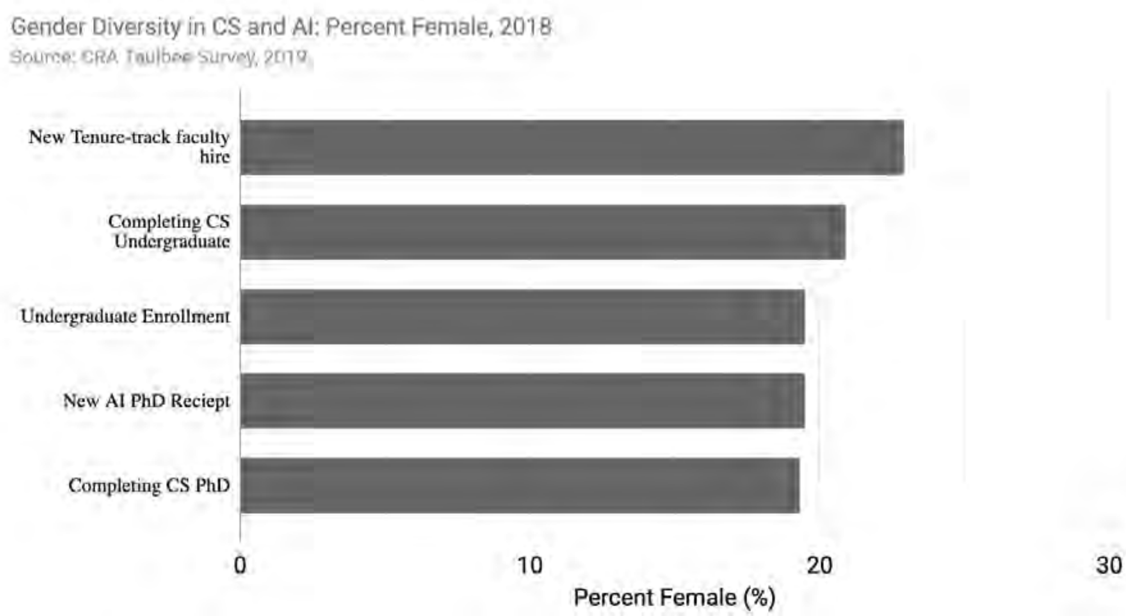

Fig. 5.10b.

Between 2010 and 2018, the percent of female AI PhD recipients has remained stable at around 20%.

[Education_Technical_Appendix]
[Access_Data]

International Academic Presence

As shown in Figure 5.11a, the proportion of new AI PhD recipients from abroad has increased from below 40% in 2010 to over 60% in 2018. This remarkable trends indicates that the production of AI doctorates in the US is largely driven by international students.

Only a small portion of these graduates go to academia (around 18%) and an even smaller portion leave the US for jobs after graduating (around 10%) (Figure 5.11b).

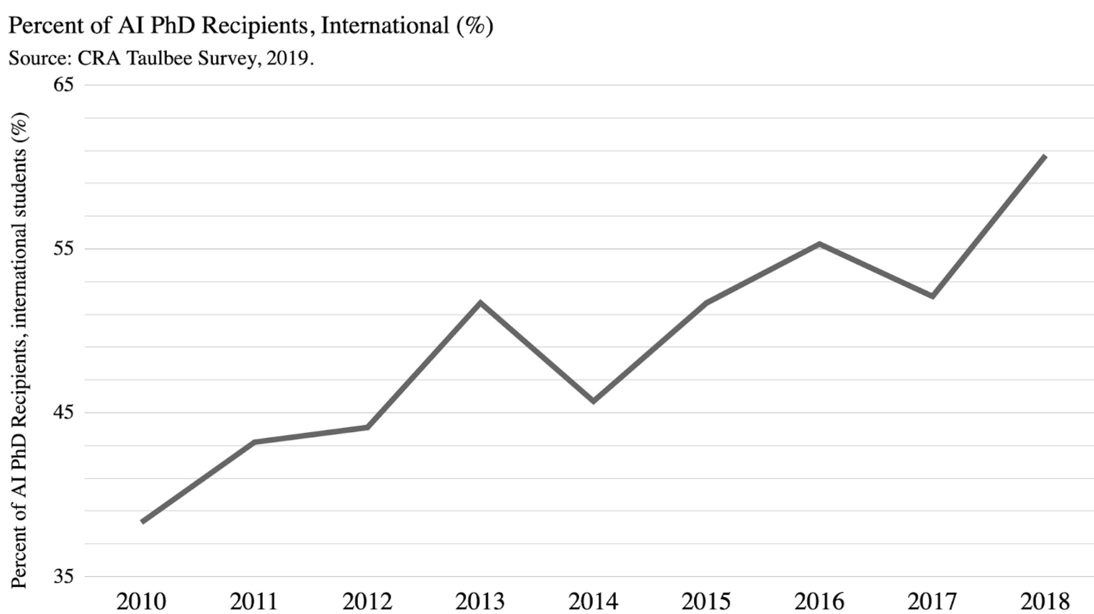

Fig. 5.11a.

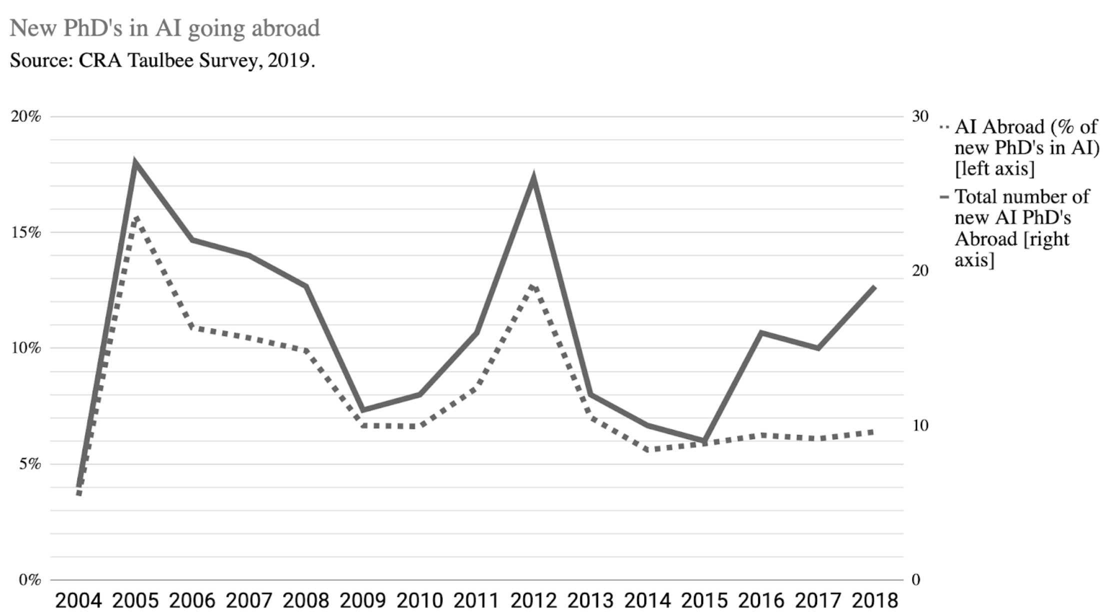

Fig. 5.11b.

Between 2010 and 2018, the number of international doctoral recipients has increased from below 40% to over 60%.

Gender Diversity

The graph below (Figure 5.12) shows the gender breakdown of AI professors at several leading computer science universities around the world. Data was collected using faculty rosters on September 21, 2019.[17] Schools with easily accessible AI faculty rosters were selected. Due to the limited number of schools studied, these findings are a small view onto a much larger picture.

Across all educational institutions examined, males constituted the clear majority of AI department faculty, making up 80% of AI professors on average.

Within the institutions examined, ETH Zurich had the most female AI faculty as a percentage of the total department at 35%, while IIT Madras had the lowest percentage at 7%. There were no discernible differences in gender split across different regions of the globe, nor was there any correlation between the faculty gender split and department size.

There remains a lack of data on diversity statistics in industry and in academia. See Appendix for data and methodology.

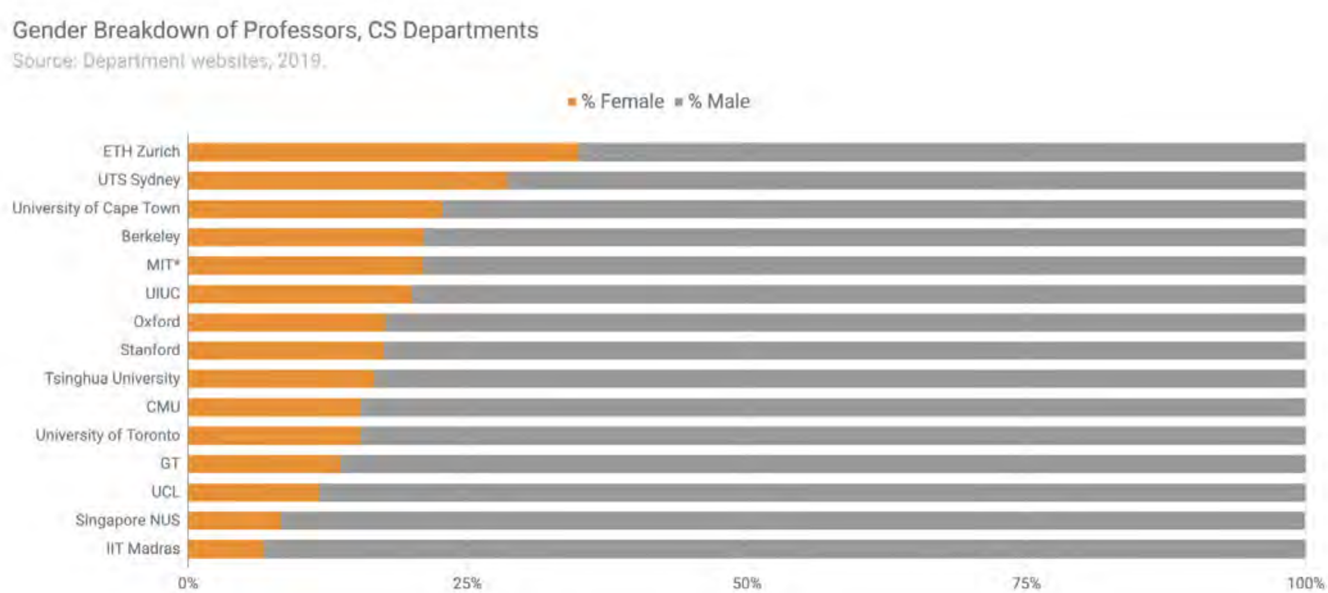

Fig. 5.12.

A significant barrier to improving diversity is the lack of access to data on diversity statistics in industry and in academia.

[17] "Female" and "male" are the terms used in the data. The Index aims to include options beyond binary in future data collection efforts.

Ethics Courses

With the rise of AI, there has been an increased urgency to reimagine approaches to teaching ethics within computer science curricula. Currently, there are two approaches: (1) stand-alone ethics courses, which are individual courses that combine ethics and policy, and (2) program-wide efforts to integrate ethics into courses in the core computer science curriculum, like Harvard's Embedded EthiCS and other efforts in the Responsible CS Challenge. Fiesler et al., 2019 and Grosz et al., 2019 discuss these models.[18] (Figures 5.13).[19] The first approach includes broad "CS and Ethics" courses, like Stanford's CS181 and Berkeley's CS 195, which include AI topics, and more specific "AI and Ethics" courses, like Harvard's CS 108 and Cornell's CS 4732, which typically examine ethical challenges from several different areas of AI. The second approach adds ethics modules to the full range of individual AI and ML courses (as well as to courses in other areas of CS). Both approaches are important, and some universities are working to integrate both.

Tech Ethics Courses, by Department

Sample includes 235 courses from universities around the world.

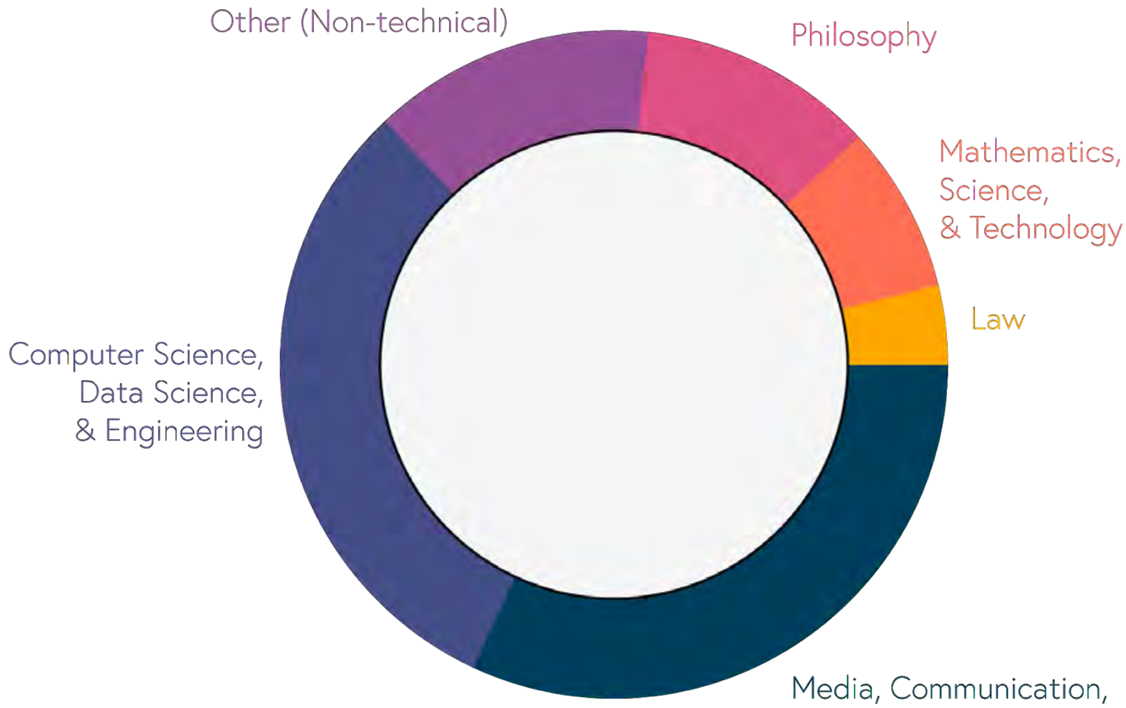

Source: Tech Ethics Curriculum (Casey Fiesler), 2019

Fig. 5.13a.

"In addition to encouraging contribution to this growing research space, we also hope that this work can serve as a call to action that can encourage and assist instructors at all educational levels who are interested in including ethics as part of their class, as well as computing programs with a goal towards increasing the reach of ethics across a curriculum."
Casey Fiesler, Natalie Garrett, Nathan Beard
What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis

[18] B.J. Grosz, D.G. Grant, K.A. Vredenburgh, J. Behrends, L. Hu, A. Simmons, and J. Waldo, (2019) "Embedded EthiCS: Integrating ethics broadly across computer science education." Communications of the ACM.
[19] The dataset downloaded from the Tech Ethics Curriculum spreadsheet had 238 courses listed. At the time of analysis 235 courses had the department listed. Included are what the instructor (or crowdsourced additions) would have deemed appropriate to add to a list of "tech ethics courses". In this dataset, the authors did not make any judgments about the character of the course beyond its inclusion in the crowdsourced list. It should be noted that by no means this analysis is a representative sample.

Measurement Questions

- A common definition of AI skills is required to assess AI education outcomes in a comprehensive manner.
- Likewise, there needs to be a survey (either annual or real-time) to accurately estimate AI course enrollment and graduation for undergraduate, masters, and PhD programs that are nationally representative and comparable across countries and regions.
- Innovative methods to scrape web data of university courses and programs could also be an invaluable resource for tracking AI learning. It is also important to get a sense of the generation of AI-trained workforce, in the US and globally.

Chapter Preview

Chapter 6: Autonomous Systems

[Table_of_Contents]
[Autonomous_Systems_Technical_Appendix]

Introduction

AI is a key component of Autonomous Systems. This chapter presents data on Autonomous Systems divided in two sections: Autonomous Vehicles (AV's) and Autonomous Weapons (AW's). The AV section shows the countries (AI Index web survey) and cities (Bloomberg Philanthropies) testing AV's. This is followed by US state policy on AV from the National Conference on State Legislation (NCSL). Data from the State of California presents metrics on total AV miles driven and number of companies testing based on the Department of Motor Vehicles (DMV) Disengagement Reports. The results from DMV Collision reports are also analyzed to present safety and reliability metrics related to AVs. The section on AW presents the known types of autonomous weapon deployments and by which country based on expert survey data collected by the Stockholm International Peace Research Institute (SIPRI).

Global

Autonomous Vehicles (AVs) are one of the most visible and potentially disruptive applications of AI. There are prototypes currently being tested around the world. While it is difficult to present a fully comprehensive list of countries where testing is taking place, data from Bloomberg Philanthropy offers insight on the global reach of AV's beyond the United States. The map (Figure 6.1a) below shows at least 25 countries with cities that are testing AV's.

Nordic countries and the Netherlands have made big strides in deploying electric vehicles (EV) charging stations and in using AV's for logistic supply chain management. In cooperation with Germany and Belgium, AV truck platoons will run from Amsterdam to Antwerp and Rotterdam to the Ruhr Valley. Similarly, Singapore has designated test areas in the metropolis for AV's (Figure 6.1b).

World Map of Countries Testing AVs

Source: Online searches on nations testing AV's.

Fig. 6.1a.

Cities Testing Autonomous Vehicles

Source: Bloomberg PhilanthropiesBloomberg Philanthropy, 2019.

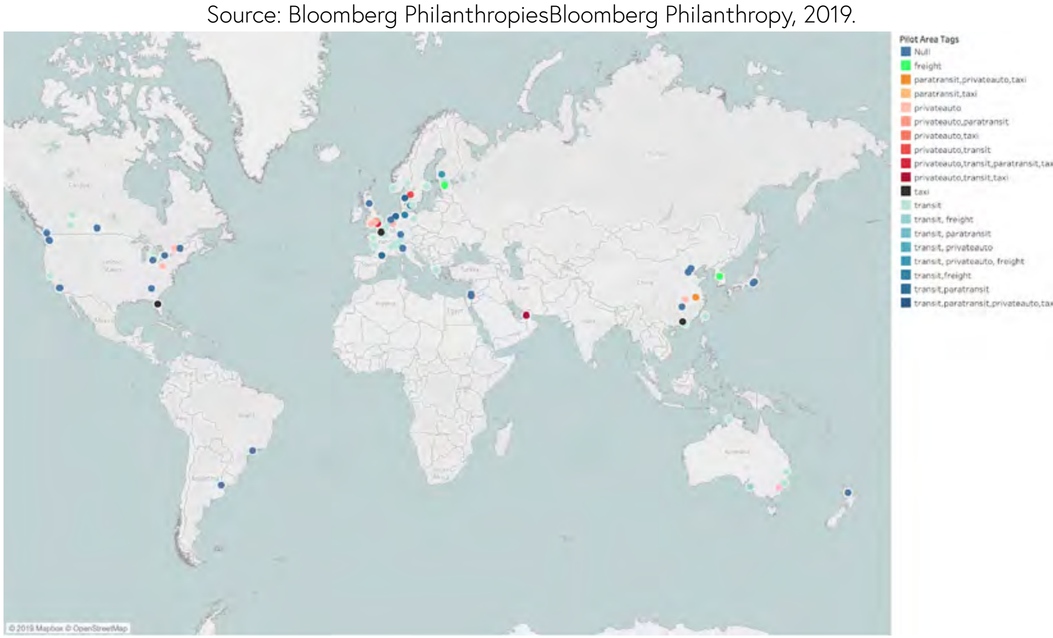

Fig. 6.1b.

US: State Policies for AVs

California was the first state with autonomous vehicle testing regulations. The number of states considering legislation related to autonomous vehicles has been increasing (Figure 6.2). Since 2012, at least 41 states and D.C. have considered legislation related to autonomous vehicles.[21] Ten states authorize full deployment without human operator, including Nevada, Arizona, or Texas, as well as many States on the east coast. Colorado authorized full deployment with a human operator. Many states, such as South Carolina, Kentucky, and Mississippi, already regulate truck platooning.[22]

US State Law on AVs

Source: National Council on State Legislation (NCSL), Governors Highway Safety Association (GHSA), 2019.

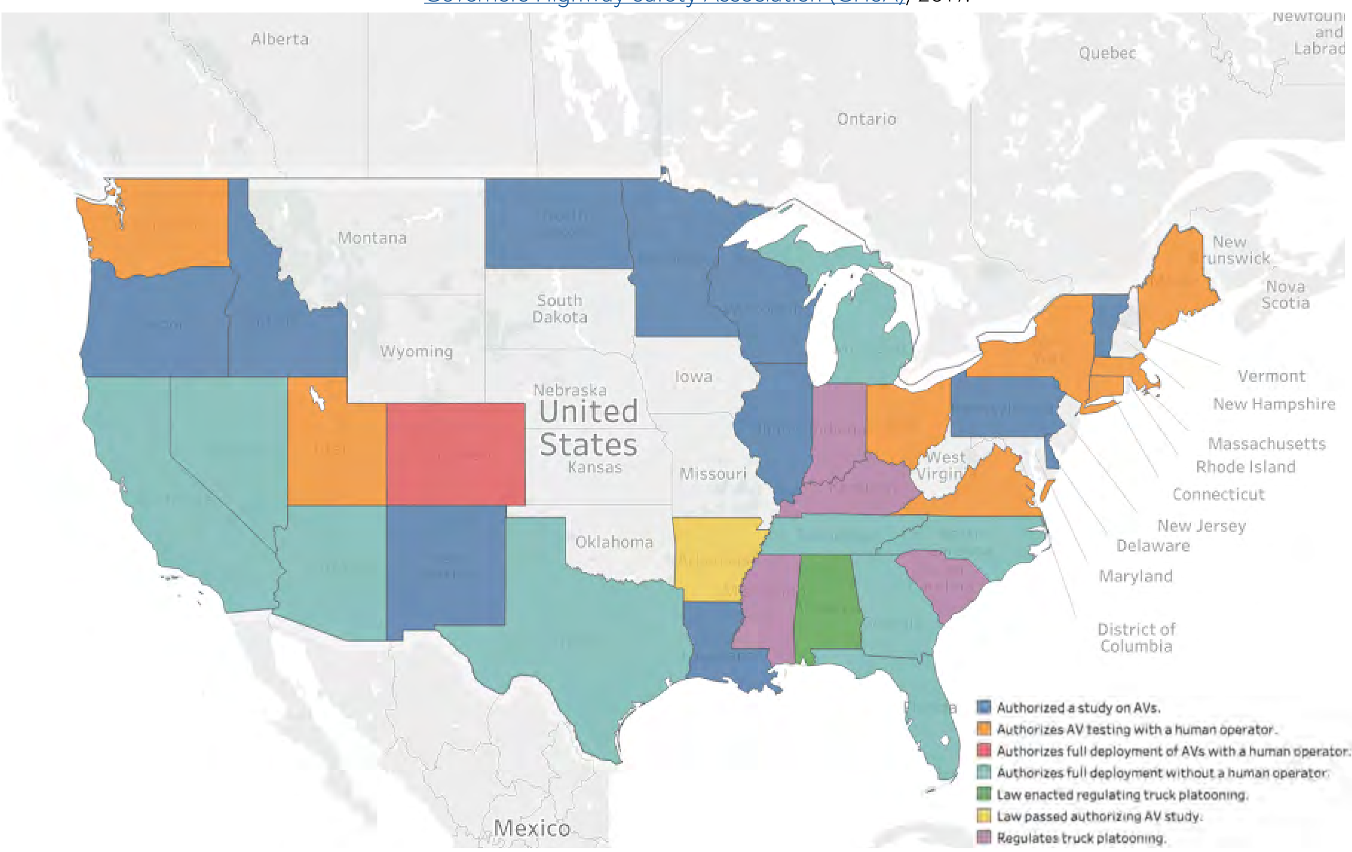

Fig. 6.2.

[21]In 2012, six states, in 2013 nine states and D.C., in 2014 12 states, in 2015 16 states, in 2016 20 states, in 2017 33 states enacted AV related bills. In 2018, 15 states enacted 18 AV related bills. In 2017, 33 states have introduced legislation. In 2016, 20 states introduced legislation. Sixteen states introduced legislation in 2015, up from 12 states in 2014, nine states and D.C. in 2013, and six states in 2012. In total, 29 states have enacted legislation related to autonomous vehicles. Readers can find California DMV Title 13, Division 1, Chapter 1, Article 3.7 –Testing of Autonomous Vehicles which defines the capability and operations that meets the definition of Levels 3, 4, or 5 of the SAE International's Taxonomy and Definitions for Terms Related to Driving Automation Systems.

[22]Truck platooning is the linking of two or more trucks in convoy, using connectivity technology and automated driving support systems. These vehicles automatically maintain a set, close distance between each other when they are connected for certain parts of a journey, for instance on motorways (ACEA, 2019). Multi-brand platooning (up to SAE level 2) with the driver still ready to intervene. By 2023, it should be possible to drive across Europe on motorways (thus crossing national borders) with multi-brand platoons, without needing any specific exemptions. Subsequently, allowing the driver of a trailing truck to rest might come under consideration. Full autonomous trucks will only come later. On 09/2016, NHTSA issued a "Federal Policy for safe testing and deployment of automated vehicles" .

California

In 2018, the State of California licensed testing for over 50 companies and more than 500 AVs, which drove over 2 million miles.[23] Figure 6.3 below shows the number of companies that are testing AV's in California (blue line on the left axis) and the total number of AVs on the road (red line on the right axis). Both metrics grew at an annual compounded growth rate (2015-18) around 90%, increasing sevenfold since 2015. The second chart (Figure 6.4) shows the total number of miles driven and total number of companies testing autonomous vehicles (AVs). This number is calculated by summing the total number of miles driven by individual AV companies, as reported in the Annual DMV Disengagement Reports. 2018 was the year of fastest growth in total miles covered by AVs totaling over 2 million miles. The compounded annual growth (2015-18) for total AV miles driven was 64% growing fourfold since 2015.

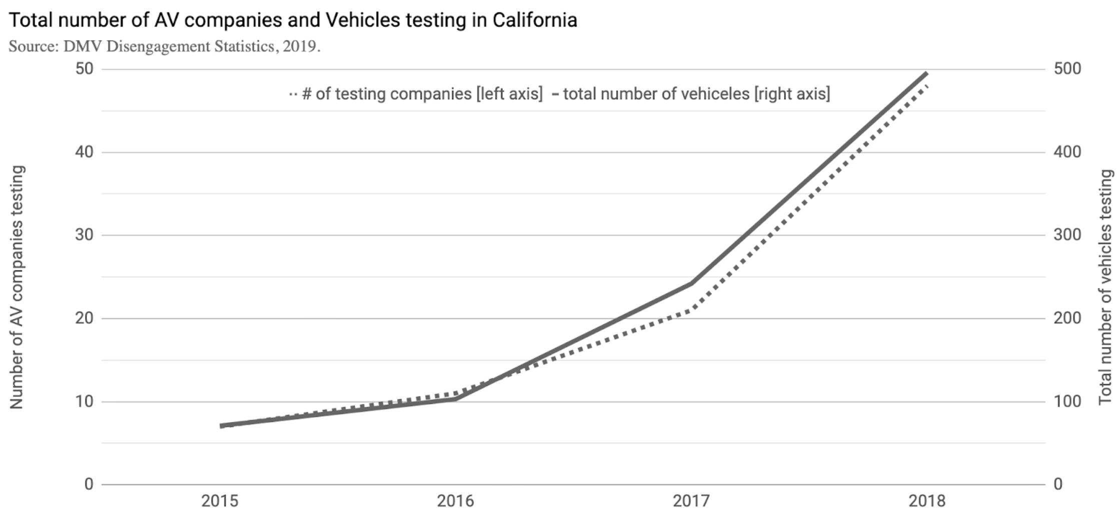

Fig. 6.3.

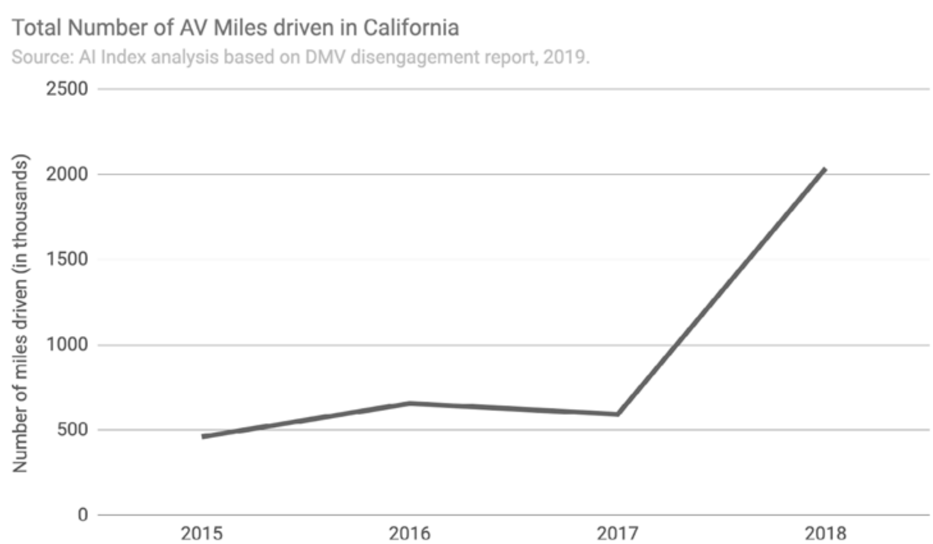

Fig. 6.4.

[23] Effective on September 16, 2014, the autonomous vehicles testing regulations in California require a driver and every autonomous mile, accident, and disengagement to be reported under CA regulation §227.02.

Safety and Reliability

Six times more people have died in traffic related fatalities than the number of fatalities in all wars for the US (Washington Post, 2019). The hope is that AVs can help reduce traffic fatalities in both advanced and developing countries.

Crashes per million miles driven in autonomous mode is the simplest and is the most reliable measure of AV safety (Figure 6.5). In 2018, AV's in CA had 46 crashes coded as being in the autonomous mode in 2018, while driving 2.05 million miles* in the autonomous mode. Or 22.44 crashes per million miles driven. To put this number in perspective below is a table from a 2016 UMTRI report that took an early look at CA AV crash rates. Even adjusting for under-reporting, the 22.44 crashes per million miles for the CA AV fleet is about 5.5x higher than the ADJUSTED rate expected for human-driven vehicles. (see notes on crash rate in Appendix).

In the early stages of development of AV testing, the number of AV related fatalities could be higher than normal traffic fatalities. A higher crash rate may be observed through every mode of automated driving. For example, in 2018 California had 2.05 million AV miles. The point estimate of human driver is at 4.1 (UMTRI) the expected crashes for AV is 8.4 with actual AV crashes in California of 46.

The pie charts summarize the Collision Report of the DMV. In most of the accidents, a car driven in the daytime by a human rear ends an AV that is either stopped or going straight. Studies suggest that these are caused by unexpected behavior by the AV or error by the human driver. Most damages have been minor.

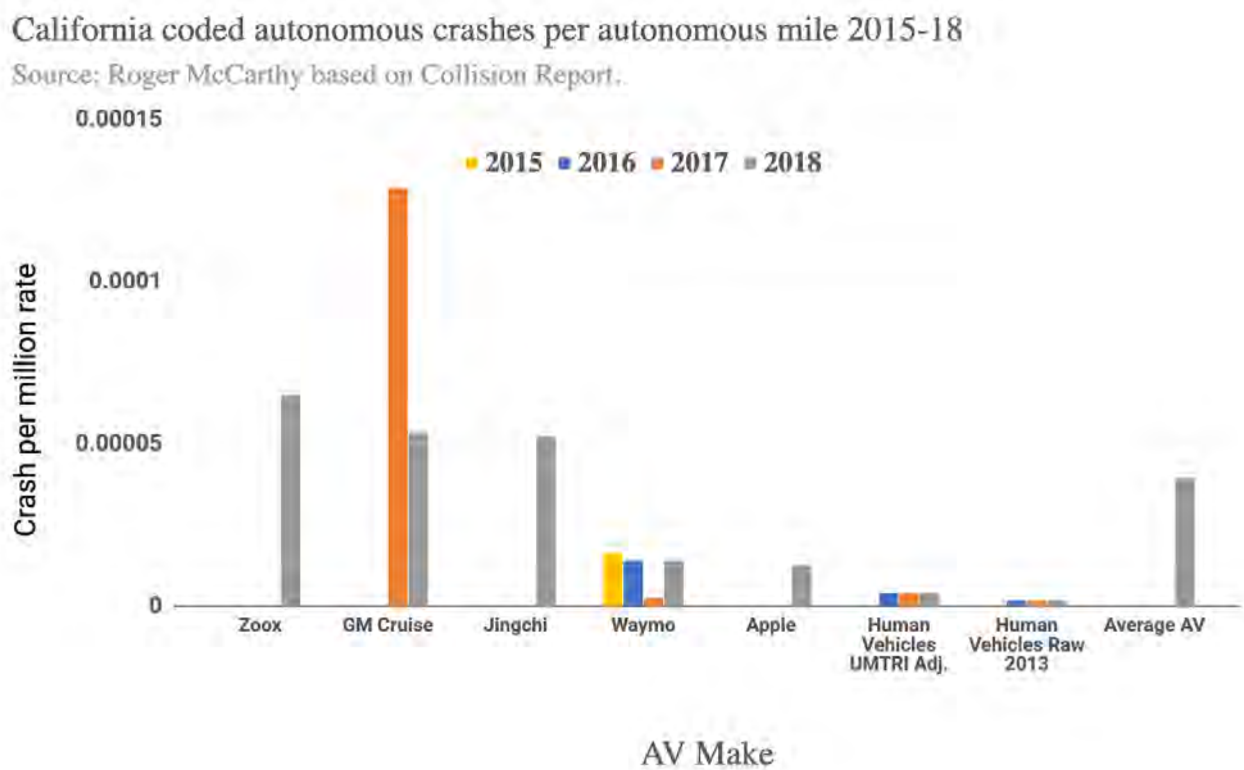

Fig. 6.5.

"I believe the 2018 AV crash rate is an underestimate of the true crash rate, and I expect the AV crash rate to continue rising. The calculated 22.4 2018 crash rate is based on the OL 316 crash form coding, which doesn't capture the effect of the AV driver turning off the AV mode moments before a crash. I believe more accurate coding would move additional crashes into the "autonomous" category. Secondly, AV's are driven, and have their crashes, under virtually ideal daytime driving conditions. When AV's are finally tested in more adverse environments of rain, snow, and fog, I am sure the AV crash performance will degrade, as with human drivers. The technical challenges of keeping sensors clean and operational under such conditions remain."
Roger McCarthy, Principal, McCarthy Engineering

Safety and Reliability

Summary of Collision Report for Autonomous Vehicles in California, 2018

Source: DMV Collision Reports, 2019.

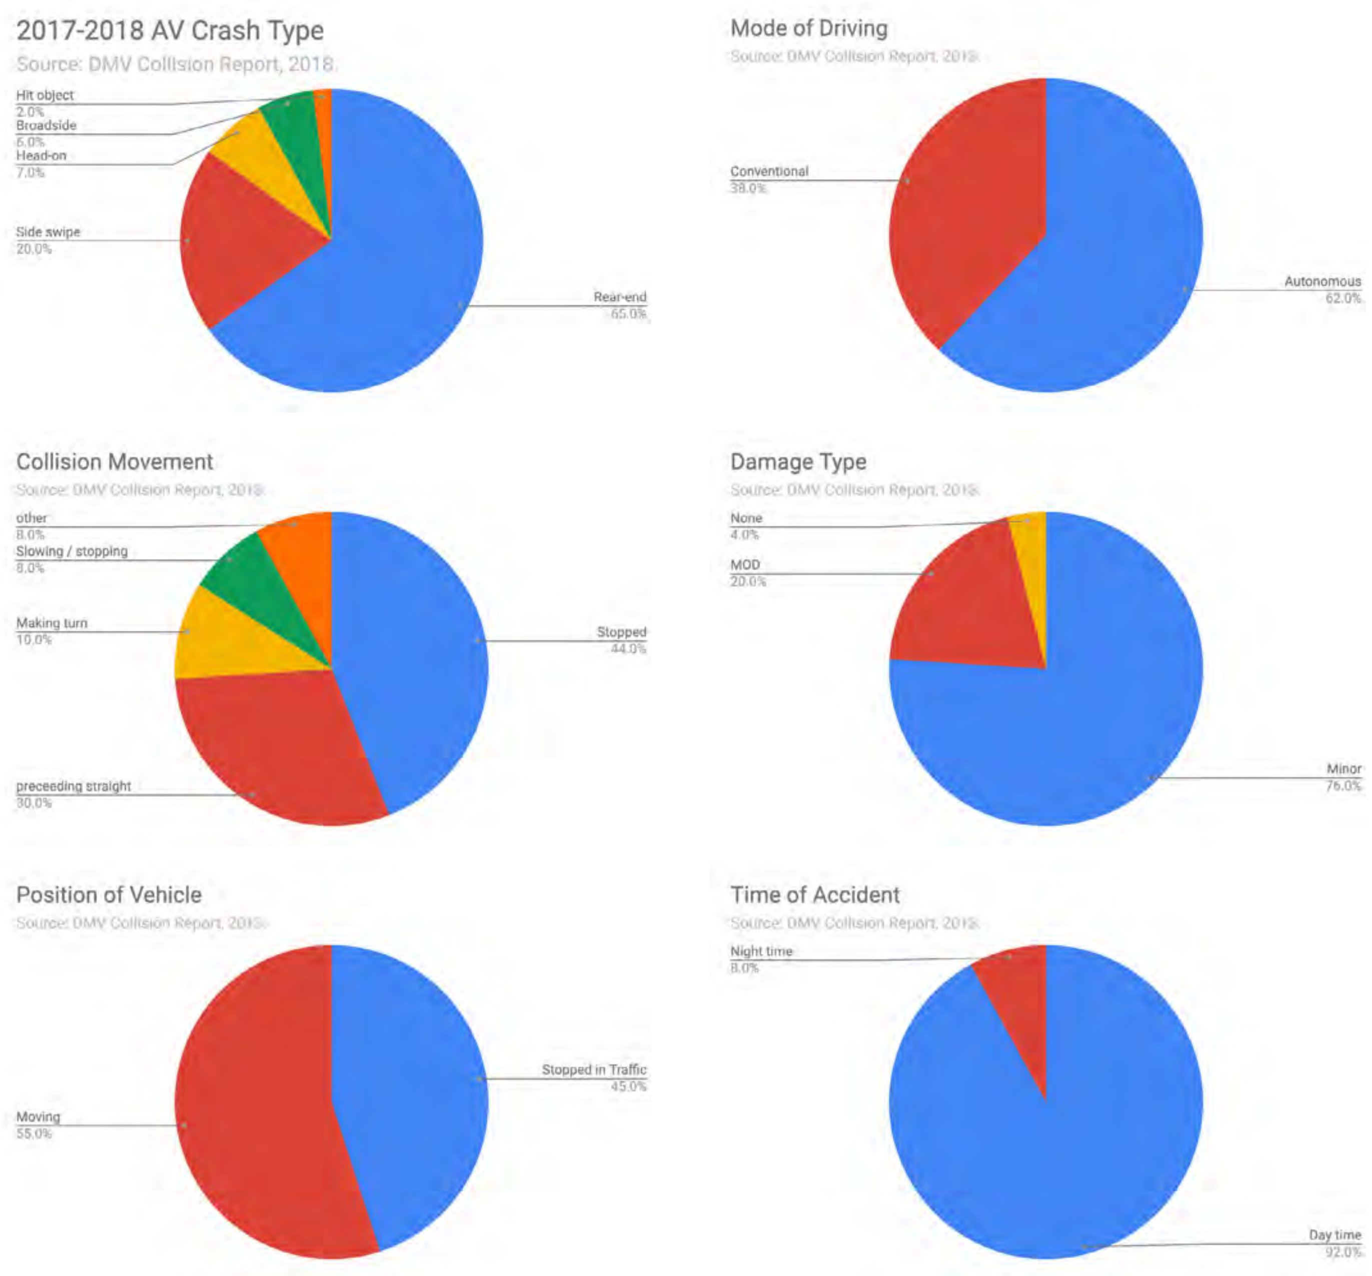

Fig. 6.6.

Measurement Questions

- The data uncertainties related to disengagement reports are well-known. Improvement in fine-grained data collection and intrinsic reporting from AV companies is critical, as is understanding which are the identifiable AI components in AV systems. The failure and incidents report of AV-AI components is industry sensitive information, which nevertheless requires standardized measurement, reporting, and identification of reliability metrics. In particular, diverse approaches to reporting even when using the same measure (for example, disengagement) highlights challenges in standardization. Further, measurement practices from companies could be associated with self-selection bias that accentuate the positive and share selectively (voluntary safety self assessment).

- Risk-informed performance-based approaches could characterize all uncertainties including engineering ones into the operation, policy and regulation of AVs. Adoption of probabilistic risk analysis from other complex engineering domains could help empower innovation and lead to better design, adequate safety features and sound policy (see Summary and Presentation Slides from: Workshop on Risk Analysis for Autonomous Vehicles: Issues and Future Directions).

Autonomous Weapons

Autonomous Weapons (AW) include various systems for either defensive or offensive capabilities. For example, Automated Target Recognition (ATR) systems autonomously acquire targets and have been in existence since the 1970s. Existing systems are largely defensive in nature with humans determining the decisions surrounding the time, location, and category of targets. A recent survey found that at least 89 countries have automatic air defense systems in their arsenal and 63 countries deployed more than one type of air defense system. Active Protection (AP) systems are developed and manufactured by only nine known producing countries. The charts below show the total known number of AW systems known to be deployed globally according to expert-curated data from the Stockholm International Peace Research Institute (SIPRI) (Figure 6.7a). The total number are classified into three labels: combative for military purpose with more than targeting capabilities i.e. machine makes the execution decision, systems with targeting capabilities only, and systems designed for intelligence, reconnaissance, and surveillance purposes including logistics, EODs, etc.. called others. A SIPRI report on [Mapping the Development of Autonomy in Weapon Systems] provides a detailed survey of AW systems. The total number of known AW systems by countries is presented between 1950-2017 (Figure 6.7b).

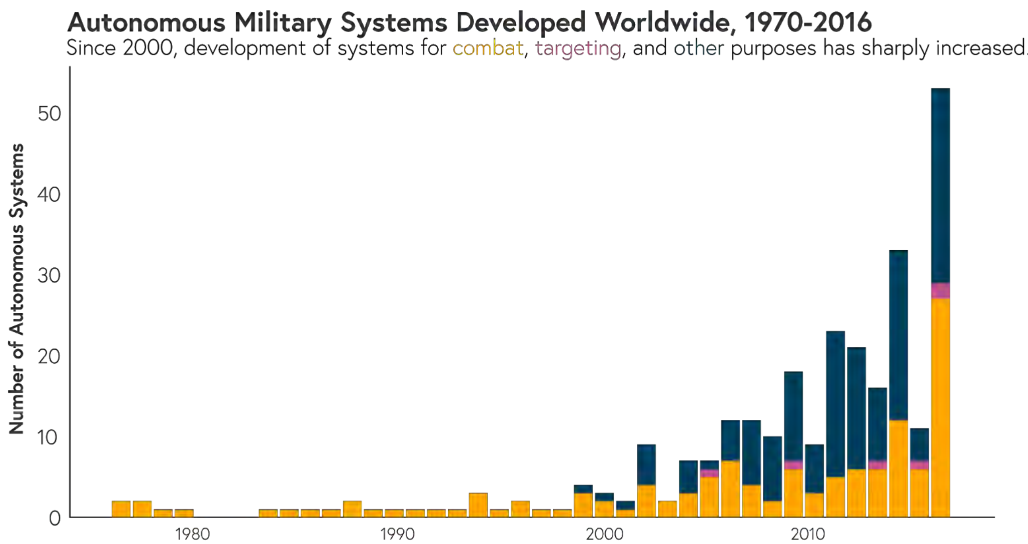

Fig. 6.7a.

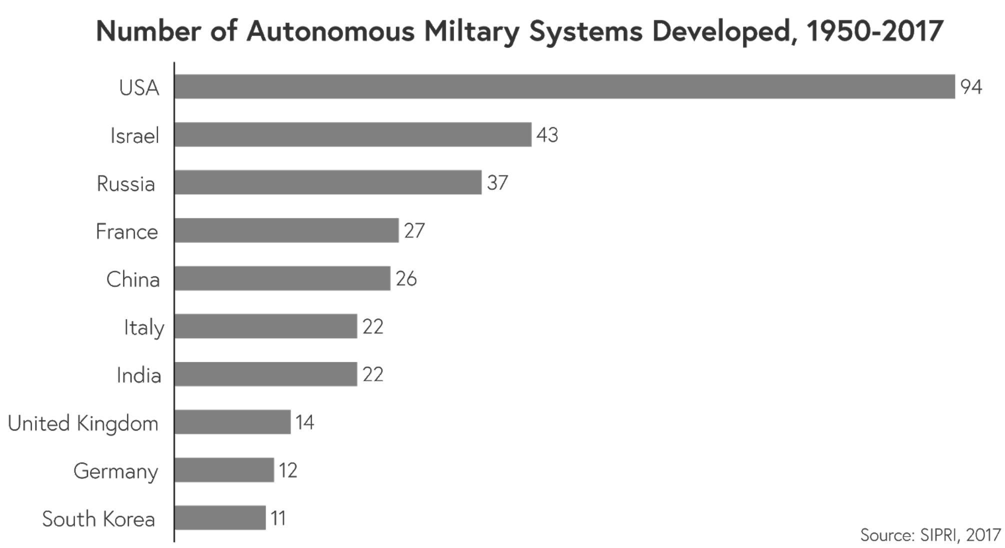

Fig. 6.7b.

Chapter Preview

Chapter 7: Public Perception

Central Banks

Central banks around the world demonstrate a keen interest in AI, especially for its ability to predict geopolitical and macroeconomic conditions, and better understand the regulatory and policy environment. The first chart below plots the global aggregate document types by central banks across 14 central banks (Figure 7.1a).[24] It shows a significant increase in central bank communications mentioning AI, with a shift from other publications to speeches mentioning AI over time. This more intensive communication reflects greater efforts to understand AI and the regulatory environment as it relates to the macroeconomic environment and financial services. The second chart plots the ranking of central banks based on the total number of AI mentions for the last ten years (Figure 7.1b). The Bank of England, the Bank of Japan, and the Federal Reserve have mentioned AI the most in their communication.

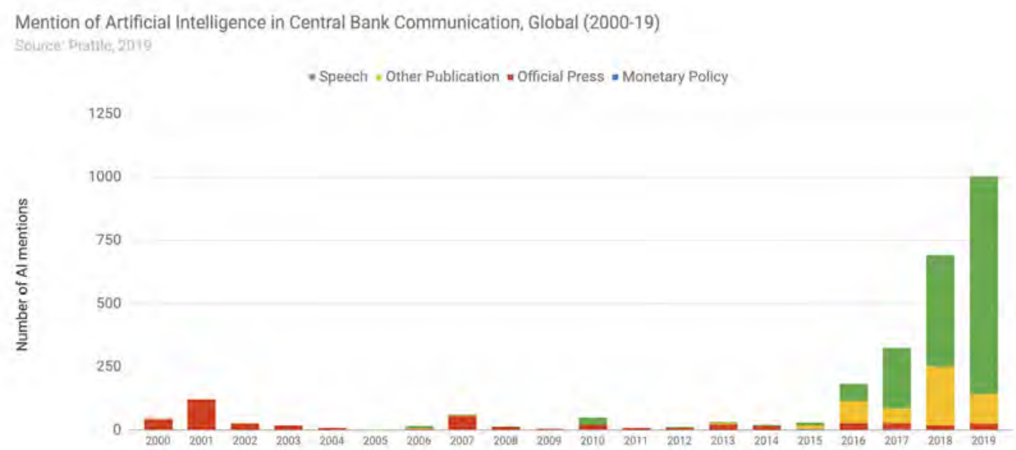

Fig. 7.1a.

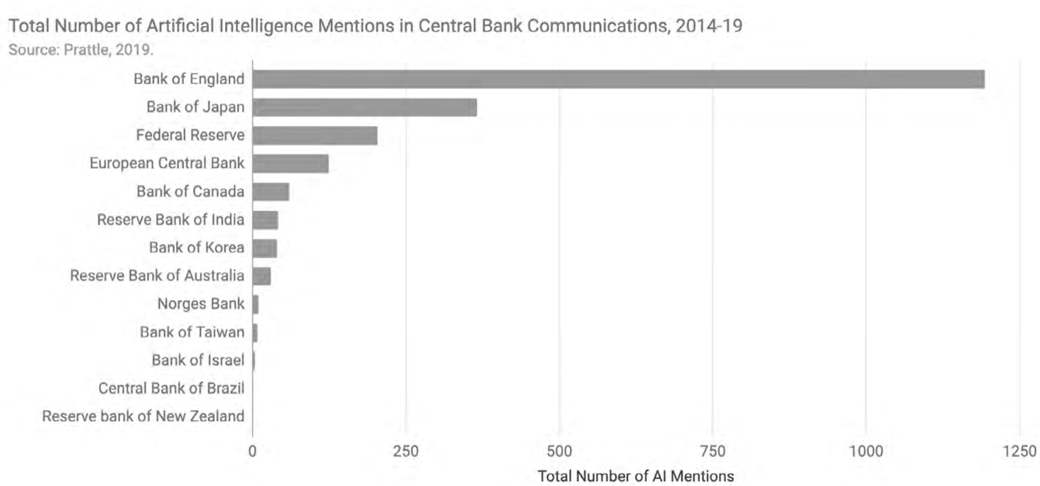

Fig. 7.1b.
Note: The chart represents data with latest data point till Q12019.

"In the last few years, the Bank of England has pursued a clear research agenda around AI as well as the use of blockchain and cryptocurrencies. Other central banks, like the Fed and BOJ, have addressed these topics in speeches, but they are just beginning to structure formal research agendas around AI."
Evan Schnidman, founder and CEO of Prattle

[24] Bank of Canada, Bank of England, Bank of Israel, Bank of Japan, Bank of Korea, Bank of Taiwan, Central Bank of Brazil, European Central Bank, Federal Reserve, Norges Bank, Reserve Bank of Australia, Reserve Bank of India, Reserve Bank of New Zealand, Sveriges Riksbank.

US Government Perception

Government officials are paying more attention to AI. The Index partnered with Bloomberg Government to analyze mentions of AI in the US congress. Each data point on the graph refers to one piece of proposed legislation, one report published by a congressional committee, or one report published by the Congressional Research Service (CRS), which serves as a nonpartisan fact-finding organization for US lawmakers, that explicitly references one or more AI-specific keywords. The data shows a greater than ten-fold increase in activity around AI in the 2017-2018 Congress, compared to prior years. More activity can be expected: our preliminary data for the 2019-2020 congress shows a further increase in activity when compared to prior years. With more than a year remaining in its term, the 116th will undoubtedly become the most AI-focused US Congress in history.

Fig. 7.2.

US, Canada, and the UK Government Perception

The next graphs show mentions of the terms 'Artificial Intelligence' and 'Machine Learning' in transcripts of US Congress (Figure 7.3a), the records of proceedings (known as Hansards) of the Parliaments of Canada (Figure 7.3b) and the United Kingdom (Figure 7.3c). Prior to 2016, there were few mentions of artificial intelligence or machine learning in the parliamentary proceedings of each country. Mentions appeared to peak in 2018, and, while remaining significant, have declined in 2019 for Canada and the United Kingdom. In transcripts of the US Congress, 2019 was year of highest AI mentions to date.

Note that it is difficult to make country-to-country comparisons, due to variations in how remarks and comments are counted between each (see Appendix for methodology). Thus, rather than country-to-country comparisons, it would be better to compare trends over time within a country.

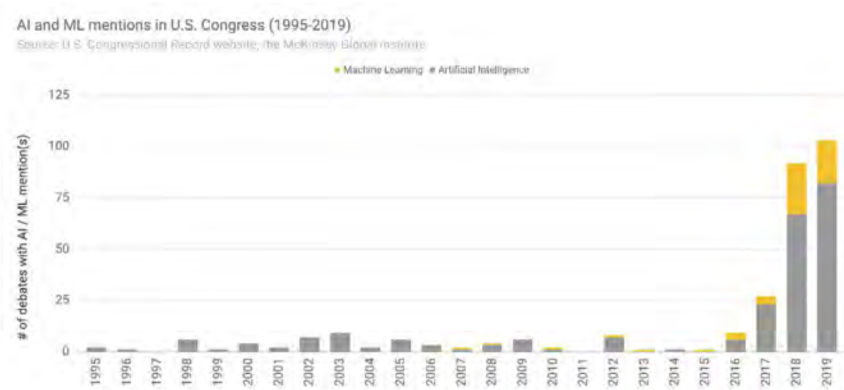

Fig. 7.3a.

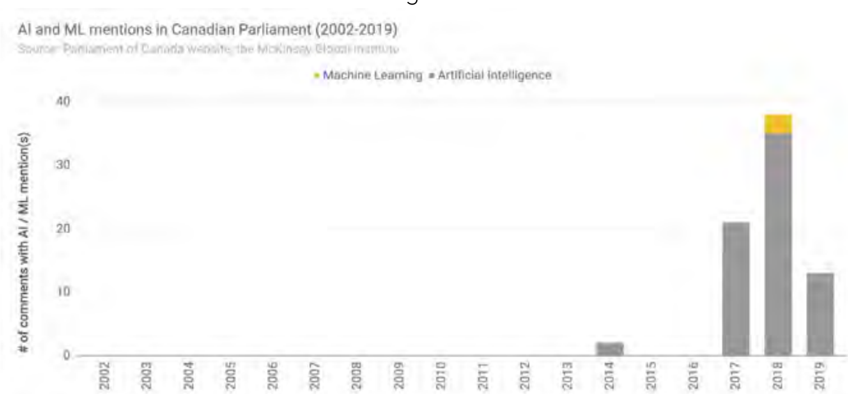

Fig. 7.3b.

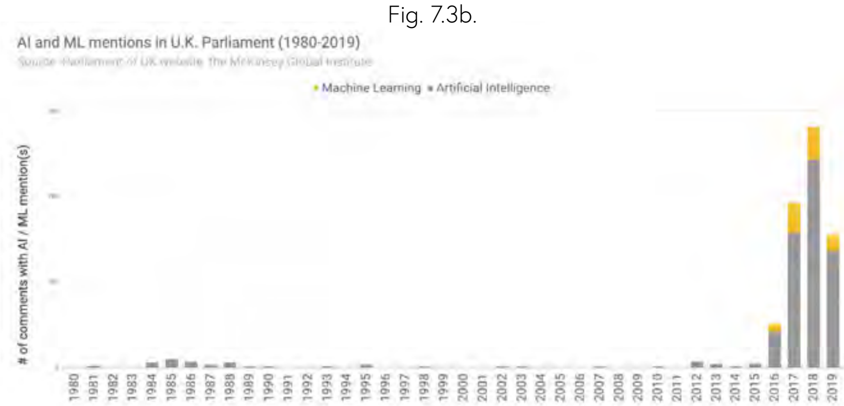

Fig. 7.3c.

Corporate Perception

The following earnings calls data includes all 3000 publicly-traded companies in the US, including American Depositary Receipts (ADRs - foreign-listed companies that also trade on a US exchange). The charts below show the individual instances of AI-related terms mentioned on earnings calls (Figure 7.4a). The share of earning calls where AI is mentioned has increased substantially, from 0.01% of total earnings calls in 2010 to 0.42% in 2018.

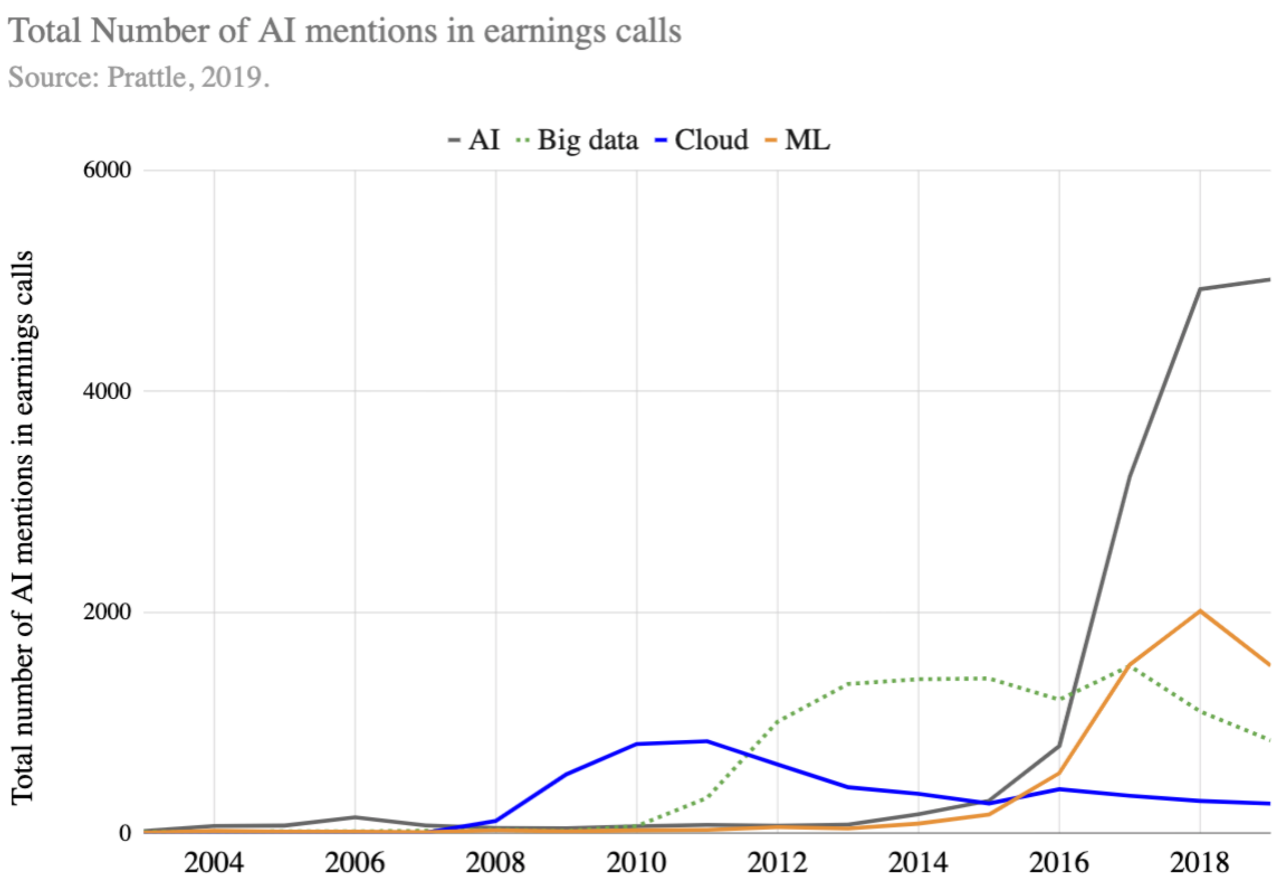

Fig. 7.4a.

Corporate Perception

Among sectors, finance has the largest number of AI mentions in earnings calls from 2018 to Q1 of 2019, followed by the electronic technology, producer manufacturing, healthcare technology, and technology services sectors (Figure 7.4b). A normalized view for the mentions of AI relative to total earnings calls is presented in the Appendix chart.

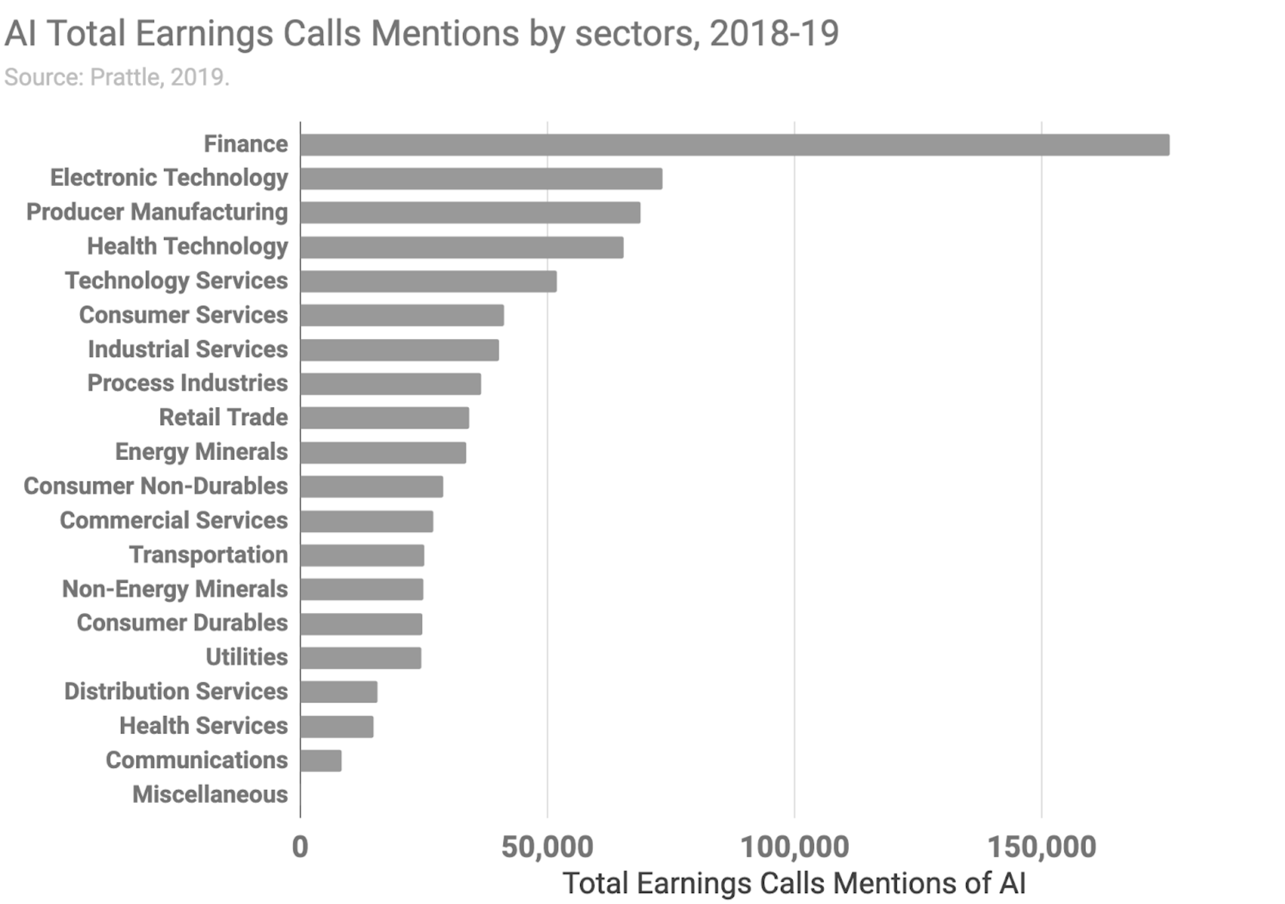

Fig. 7.4b.

Web Search and World News

The timeline below shows the relative search interest by month of web searchers in the United States from January 2004 to August 2019 for the phrases "data science," "big data," "cloud computing," and "machine learning" using Google Trends (Figure 7.5a). Google's methodology calculates the time period with the highest amount of searching, then treats that as 100 and scales the rest accordingly.

In this analysis there is an emergence of cloud computing in 2008, which is replaced as the term of art by "big data" which starts taking off in 2011. Machine learning and data science both take off together in 2013, following technical advances in deep learning like the results on the 2012 ImageNet competition.

US search interest for “data science,” “big data,” “cloud computing” and “machine learning” via Google Trends
Source: Google Trends, GDELT, 2019.

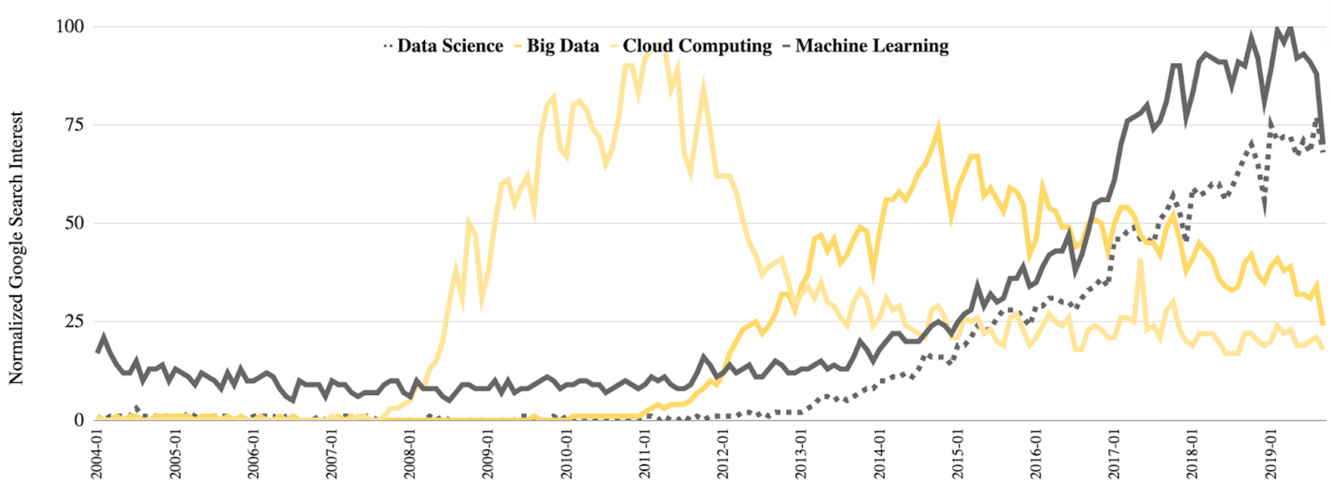

Fig. 7.5a.

Web Search and World News

The timeline below compares some of the terminology used to refer to AI today: "machine learning," "deep learning," "artificial intelligence", as well as the term for the most popular deep learning software, "TensorFlow" (Figure 7.5b). Google's TensorFlow package is now searched just as often as AI and both have been slowly decreasing in search interest since early 2018. After taking off in 2013, deep learning plateaued in late 2017, around the time that searches for machine learning began to slowly level off.

Using data from the GDELT Project, the timeline below shows the percentage of worldwide news coverage in 65 languages monitored by GDELT by day that contain those same four terms since January 1, 2017, using a 7-day rolling average to smooth the data. This graph shows that online news coverage of cloud computing and big data has steadily declined and data science and machine learning have increased. This frequency of queries suggests that "big data" retains its allure as a media term for journalists covering the latest data-driven news, but that in both searches and news coverage, Machine Learning is the term *du jour*.

US search interest for "machine learning," "deep learning," "TensorFlow" and "artificial intelligence" via Google Trends
Source: Google Trends, 2019.

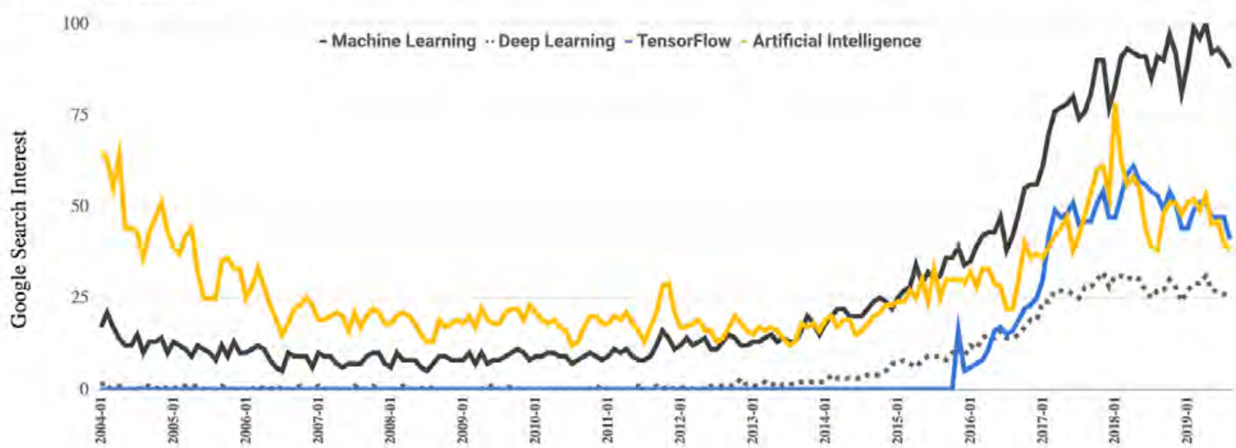

Fig. 7.5b.

Percent of worldwide news coverage monitored by GDELT that mentioned "data science," "big data," "cloud computing" and "machine learning"
Source: GDELT, 2019.

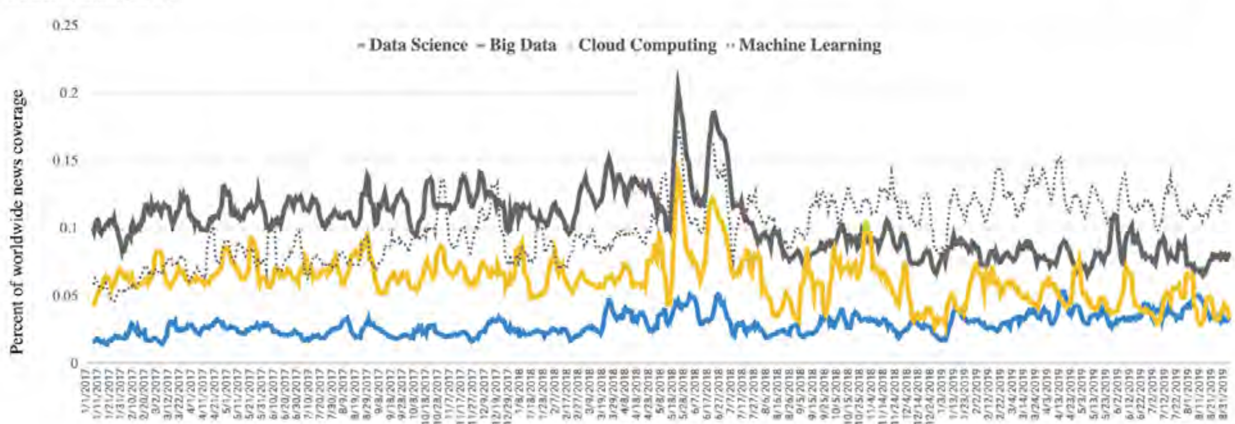

Fig. 7.5c.

Web Search and World News

Looking at online news coverage, the timeline below shows that "Artificial Intelligence" is the clear winner, followed by Machine Learning and deep learning (Figure 7.5d).

When the media covers AI, what does media think AI is influencing? The bar chart below shows the percentage of articles monitored by GDELT containing either "artificial intelligence" or "machine learning" or "deep learning" that also contained either "job" or "jobs" or "employment" or "unemployment," the percentage that contained either "killer robot" or "killer robots" or "autonomous weapon" or "autonomous weapons," and the percentage that contained either "bias" or "biases" or "biased" (Figure 7.5e).

Percent of worldwide news coverage monitored by GDELT that mentioned “machine learning,” “deep learning,” “TensorFlow” and “artificial intelligence”
Source: GDELT, 2019.

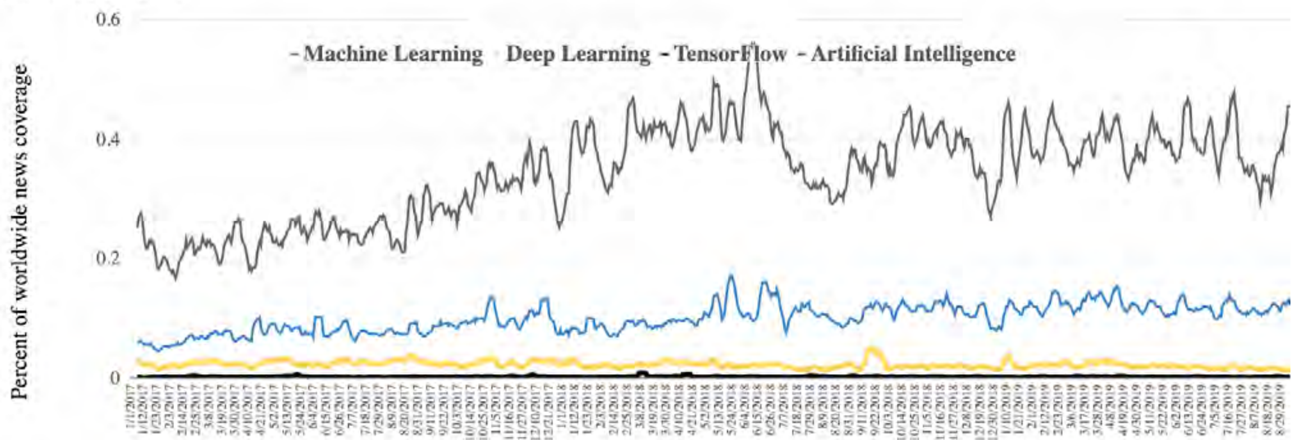

Fig. 7.5d.

Percent of worldwide online news coverage of AI monitored by GDELT that focused on jobs, autonomous weapons or bias.
Source: GDELT, 2019.

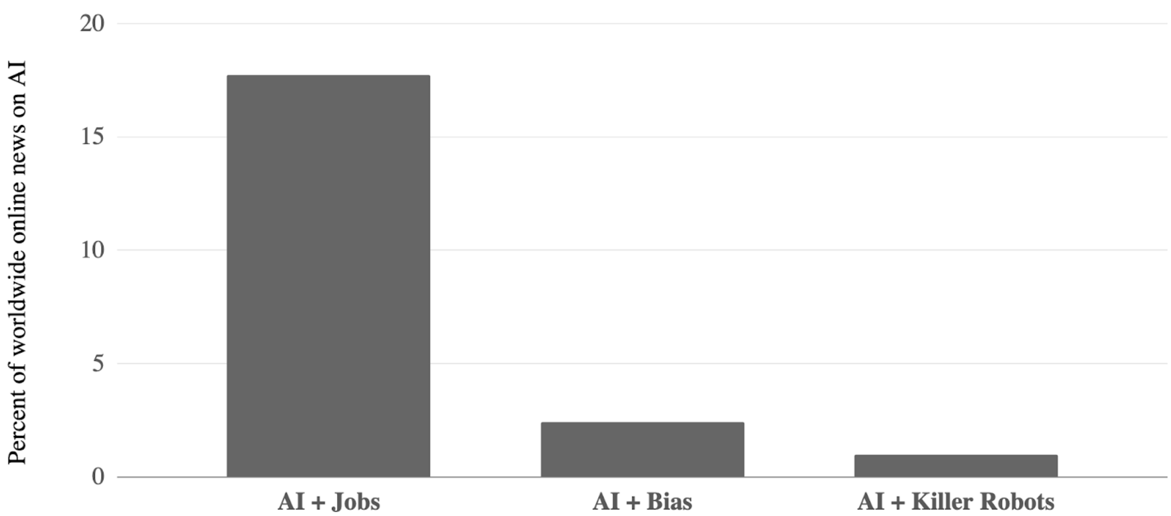

Fig. 7.5e.

Web Search and World News

Articles addressing AI's potential impact on jobs, including concern over the potential for AI to displace human jobs, accounted for 17.7% of all AI-related coverage GDELT monitored over the past two and a half years. Killer robots accounted for just 0.99% and bias issues accounted for just 2.4% of AI discussions (Figure 7.5f).

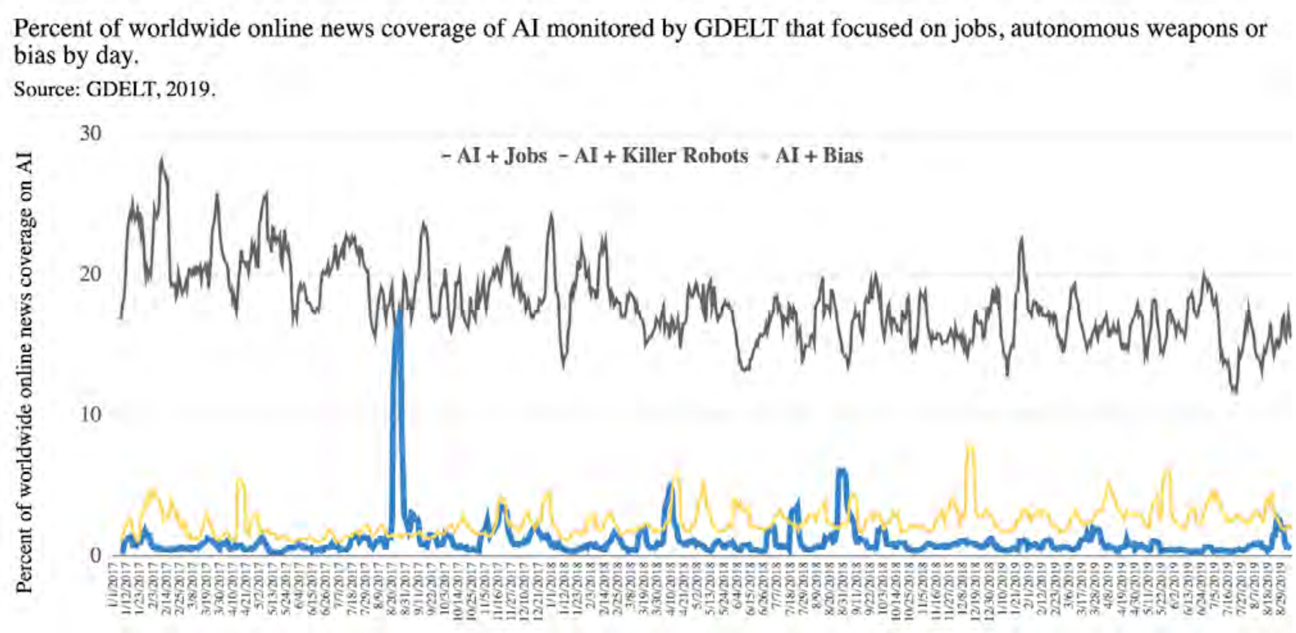

Fig. 7.5f.

Chapter Preview

Chapter 8: Societal Considerations

[Table_of_Contents]
[Societal_Considerations_Technical_Appendix]

Introduction

This chapter begins by identifying the topics in ethical challenges mentioned in 59 Ethical AI Principle documents based on a dataset compiled by PricewaterhouseCoopers (PwC). The chapter also documents the key topics discussed in global news media on AI and Ethics based on LexisNexis data and Quid. AI use cases supporting each of the 17 United Nations (UN) Sustainable Development Goals (SDGs) are identified based on curated data from the McKinsey Global Institute (MGI).

Ethical Challenges

AI systems raise a broad variety of ethical challenges that are now the concern of government, public interest organizations, NGO's, academia, and industry. Efforts to identify these challenges and to develop guiding principles for ethically and socially responsible AI systems are emerging from each of these sectors,. This snapshot of some such t efforts was derived from an analysis of more than 100 documents.

PricewaterhouseCoopers (PwC) compiled a dataset of ethical challenges (based on topic modeling) by looking at ethical AI guidelines across for 110 documents, of which only 59 were deemed to discuss a set of AI principles. Many were simply reviews or recommendations, and were not included in the analysis. [The list of organizational documents](#) and the [list of principles](#) is available in the Appendix.

A view of ethical AI frameworks over time is plotted identifying Associations and Consortiums, Industry and Consultancy groups, Governments, Tech Companies, and Think Tanks/Policy Institutes and Academia (Figure 8.1a). It is interesting to note that initial impetus for Ethical Principles sprang from Associations and Consortiums, with other organizations subsequently releasing their respective AI Principles in 2018 and 2019.

Top 3 Ethical Challenges, Associations and Consortiums, Governments, and Tech Companies

Associations and Consortiums (19 documents)
1.) Interpretability & Explainability is cited in 95% of frameworks.
2.) Fairness is cited in 89% of frameworks.
3.) Transparency is cited in 84% of frameworks.
Governments (13 documents)
1.) Interpretability & Explainability, Fairness, and Transparency are each cited in 92% of frameworks..
Tech Companies (11 documents)
1.) Fairness is cited in 100% of frameworks.
2.) Transparency is cited in 81% of frameworks.
3.) Accountability is cited in 72% of frameworks.
Think Tanks/Policy Institutes and Academia (8 documents)
1.) Fairness is cited in 100% of frameworks.
2.) Human Control is cited in 88% of frameworks.
3.) Interpretable & Explainable Model is cited in 88% of frameworks.
Industry and Consultancy (8 documents)
1.) Transparency is cited in 88% of frameworks.
2.) Fairness, Data Privacy, and Reliability, Robustness, and Security are each cited in 75% of frameworks.

Number of Ethical AI Frameworks Produced 2016-2019, by Type of Organization

Source: PwC based on 59 Ethical AI Principle documents.

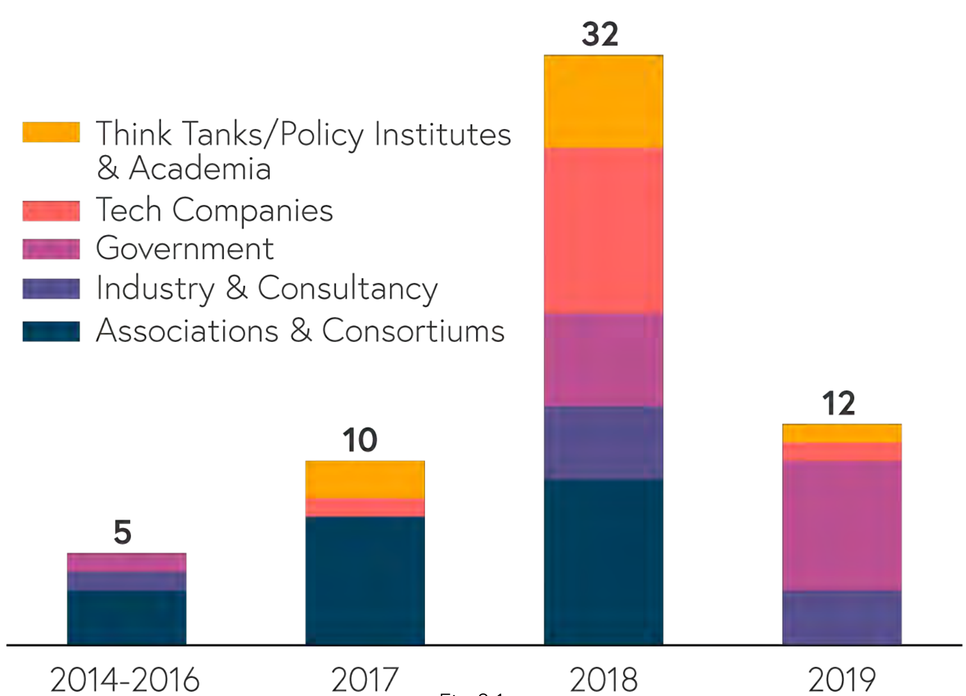

Fig 8.1a.

Ethical Challenges

Twelve ethical challenges were mentioned across many ethical AI framework documents. This list is non-exhaustive, and many important ethical issues -- including justice, economic development, poverty reduction, and inequality, are missing. Even so, these 12 ethical challenges indicate where attention has been focused:

- Accountability
- Safety
- Human Control
- Reliability, Robustness, and Security
- Fairness
- Diversity and Inclusion
- Sustainability
- Transparency
- Interpretability and Explainability
- Multi Stakeholder engagement
- Lawfulness and Compliance
- Data Privacy

To communicate the thrust of the ethical AI issues to the general public, the bar graph shows the incidence of identified ethical challenges across 59 AI Principles documents (Figure 8.1b). It shows that Fairness, Interpretability and Explainability, Transparency are most mentioned across all documents studied.

Ethical Challenges covered across AI Principle Documents

Source: PwC based on 59 Ethical AI Principle documents.

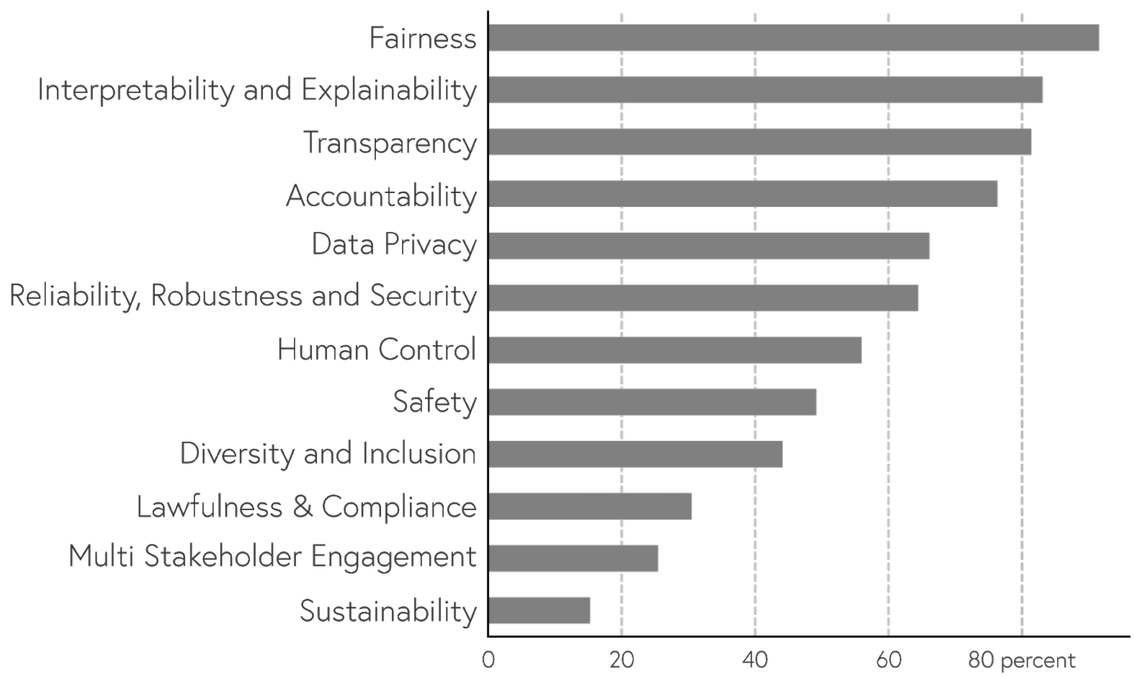

Fig 8.1b.

"Research around Ethical AI, especially on fairness, accountability, and transparency (FAT) of machine learning models has grown significantly in the past couple of years. While there is a broad consensus emerging on the core set of principles associated with ethics and AI, the contextualization of these principles for specific industry sectors and functional areas is still in its infancy. We need to translate these principles into specific policies, procedures, and checklists to make it really useful and actionable for enterprise adoption."

Anand Rao, Global AI Lead, PwC

Ethics and AI: Global News Media

Global news coverage of Artificial Intelligence has increasingly shifted toward discussions about its ethical use. To better understand how these narratives are taking shape, we leveraged Quid to search the archived news database of LexisNexis for news articles from 60,000 global English news sources and over 500,000 blogs on AI ethics from August 12, 2018 to August 12, 2019 (see Appendix for more detail on search terms).

Based on keywords defined by Harvard (seen here), Quid included search terms such as human rights, human values, responsibility, human control, fairness, discrimination or non-discrimination, transparency, explainability, safety and security, accountability, and privacy related to AI technology. Then, we selected the 10,000 most relevant articles using the platform's NLP algorithm and visualized unique articles.

Each node (or dot) on a Quid network map represents a single news article. Links connecting these articles denote articles that share similar language. When a large number of similar articles are identified and linked, clusters form to reveal unique topics. The Quid algorithm classified the resulting media narratives into seven large themes based on language similarity: **Framework and Guidelines** (32%), **Data Privacy Issues** (14%), **Facial Recognition** (13%), **Algorithm Bias** (11%), **Big Tech Advisory on Tech Ethics** (11%), **Ethics in Robotics and Driverless Cars** (9%), and **AI Transparency** (6.7%).

Quid network with 3,661 news articles on AI Ethics from August 12, 2018 to August 12, 2019. Colored by theme. Labeled by theme.

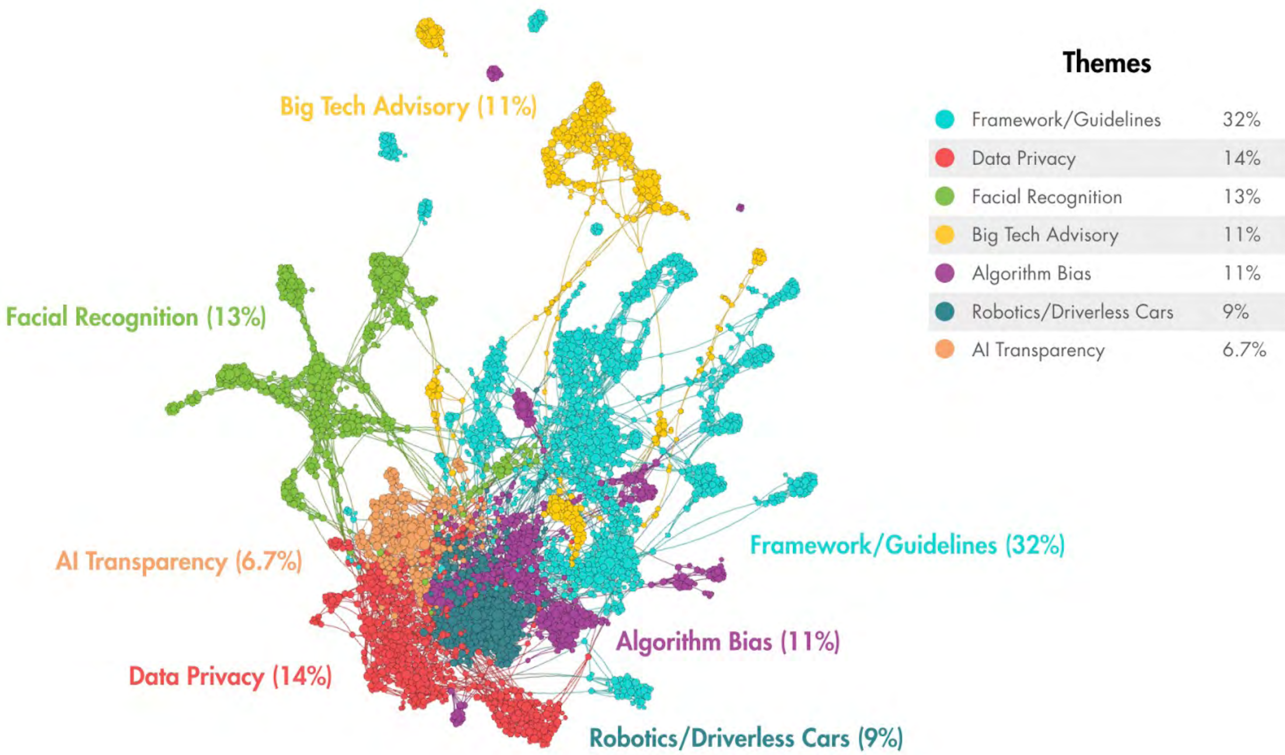

Fig. 8.2a.
Appendix: How to Red a Quid Network

Ethics and AI: Global News Media

These results indicate that the global media conversation on AI Ethics in 2019 is largely about AI ethics frameworks or guidelines led by governments, intergovernmental organizations, and research institutes (Figure 8.2a). Within the last year, nearly a third (32%) of all news articles covered AI guidelines proposed by governments or other large policy institutes, including those by the European Union and the Organisation for Economic Co-operation and Development (OECD). A smaller, but not an insignificant chunk of the conversation (11%) also included commentary from advisory groups attached to tech giants such as Google, Facebook, and Microsoft.

When filtering for ethics discussions around specific AI technologies, facial recognition dominated the attention of the news media, with 13% of all articles (Figure 8.2a). his cluster's position on the periphery of the larger AI ethics narrative indicates a high degree of uniqueness from the rest of the conversation. Public concerns over the technology's threat to data privacy have grown over time, driven by news of mistaken identities during crime surveillance, biometric scans that can be applied to videos or photos without consent, and the idea of data ownership as it relates to social media platforms that utilize the technology.

Countries differ significantly with respect to which AI ethical issues (as defined by Harvard here) they give most news coverage. While media sources based in the US or UK had more balanced coverage between categories, others reflected specific focus areas (Figure 8.2b). In Switzerland, for example, 45% of all articles covered guidelines and frameworks on AI development, while 44% of Chinese news focused on safety and security, and 48% of articles in Singaporean sources explored transparency and explainability.

Most mentioned ethics categories by Source Country

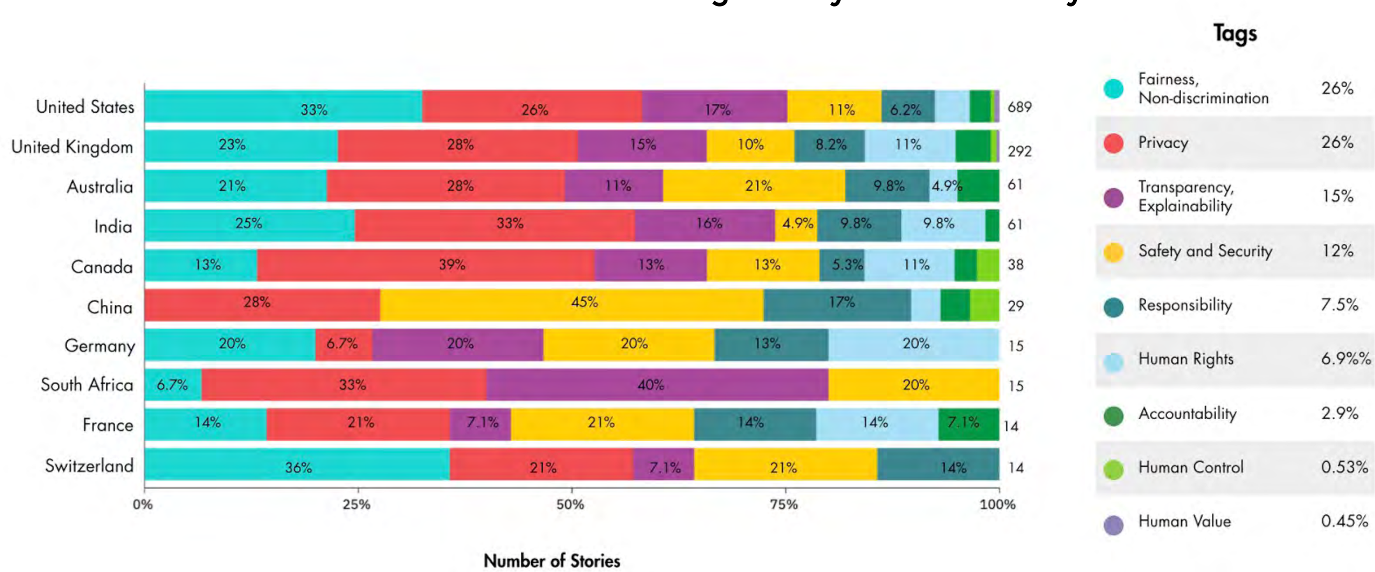

Fig. 8.3b.

Applications of AI for Sustainable Development

Artificial intelligence, while not a silver bullet, has the potential to help contribute to multi-pronged efforts to address some of society's most pressing challenges.

The mapping of AI use cases to the UN Sustainable Development Goals (SDGs) that follows are derived from a library of approximately 160 AI for social good use cases collected by the McKinsey Global Institute and Noble Intelligence, McKinsey's initiative to use AI for humanitarian purposes. The library of use cases is not comprehensive, but reflects a selection of use cases, typically in domains with initial evidence of possible applications. AI deployments in some form were identified for about one-third of use cases in the library; in about three-quarters of use cases, deployments of solutions employing some level of advanced analytics were observed, most (if not all) of which could further benefit from using AI.

To build the use case library, MGI took a two-pronged approach: from a societal point of view, MGI sought to identify key problems known to the social sector community and determine where AI could aid efforts to resolve them; from a technological point of view, MGI took a curated list of 18 AI capabilities and sought to identify which types of social problems they could best contribute to solving. Each use case highlights a meaningful problem that can be solved by an AI capability or some combination of AI capabilities. The library is not comprehensive, but it nonetheless showcases a wide range of problems where AI can be applied for social good. MGI's full discussion paper can be found at Notes from the AI frontier: Applying AI for social good.

Applications of AI for Sustainable Development

Artificial intelligence has applicability across all 17 of the United Nations Sustainable Development Goals

The UN SDGs are a collection of 17 global goals set by the United Nations for the year 2030, for poverty alleviation, improving health and education, reducing inequality, preserving the environment, and boosting economic growth, amongst other priorities. AI use cases have the potential to support some aspect of each of the UN SDGs. The chart below indicates the number of AI use cases in MGI's library that could support each of the UN SDGs (Figure 8.3a).

SDG 3, "Ensure healthy lives and promote well-being for all at all ages", could be supported by the highest number of use cases in MGI's current library. A number of use cases that leverage AI support medical diagnoses: for example, researchers at the University of Heidelberg and Stanford University have created an AI system to visually diagnose skin cancer that outperformed professional dermatologists. There are also potential cases where AI can be used to monitor, track and predict outbreaks of communicable diseases. For instance, Data Science for Social Good and McKinsey's Noble Intelligence initiative developed an algorithm to identify children most at risk of not receiving the measles vaccination, allowing physicians to spend more time educating and following up with these families.

There are also a number of AI use cases that could support SDG 16, "Promote peaceful and inclusive societies for sustainable development, provide access to justice for all and build effective, accountable and inclusive institutions at all levels." The use cases cover domains ranging from helping individuals verify and validate information, providing improved security through detection and prediction of violence, addressing bias to ensure fair and equal access to justice, to optimizing the management of public and social sector institutions. For example, AI could be used to automate question response or provision of services through digital channels, helping to improve government interactions with citizens.

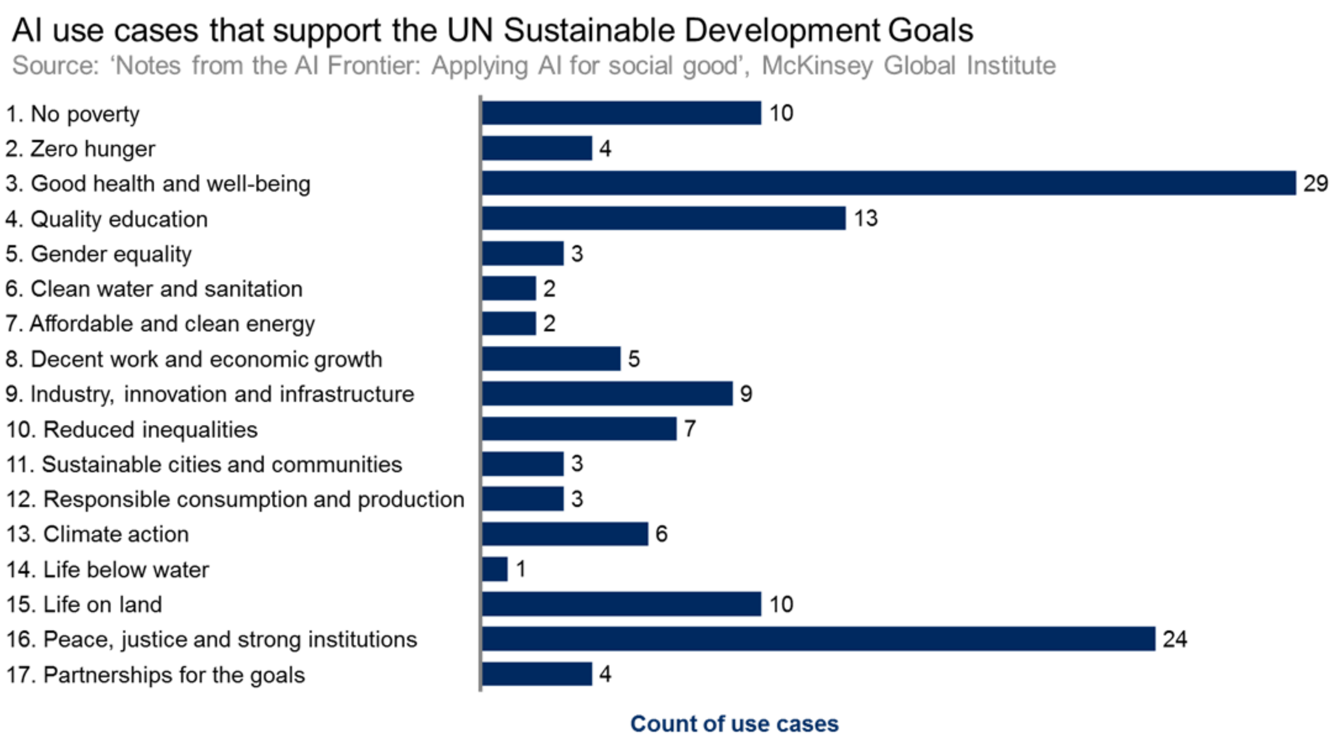

Fig. 8.3a.
NOTE: This chart reflects the number and distribution of use cases and should not be read as a comprehensive evaluation of AI potential for each SDG; if an SDG has a low number of cases, that is a reflection of our library rather than of AI applicability to that SDG.

Applications of AI for Sustainable Development

AI is applicable to driving a subset of targets across the UN SDGs

Each UN SDG is broken down into a list of targets, which are measured with indicators. There are 169 targets across the 17 UN SDGs. While AI use cases can be topically aligned to the SDGs, as displayed in the previous chart, further focus should be directed to the use cases that can directly drive impact towards achieving specific UN SDG targets and indicators.

By mapping AI use cases to the specific target(s) that they could contribute to achieving, MGI identified the subset of targets for which AI has some applicability to address. This analysis builds upon the ~160 use cases in MGI's library and others to identify which targets could be addressed by a solution in which AI is applied, recognizing that AI alone cannot solve any of the targets. The following chart displays the number of targets which AI could contribute to addressing, out of the total number of targets within each SDG (Figure 8.3b).

Some AI for sustainable development use cases are being piloted, although bottlenecks exist

A number of organizations globally are piloting applications of AI for sustainable development, although there are currently few examples of deployments of AI for sustainable development at scale. For example, AI has been piloted for several applications in disaster relief by a number of organizations, including Google, Facebook, Microsoft, Planet Labs, Airbus, SAP, and others. Still, there is more to be done to sustainably adopt these AI applications for widespread use in disaster relief across multiple partners and regions.

Some AI-specific bottlenecks will need to be overcome for AI to reach its potential for social impact. These range from challenges with data (including availability, accessibility, quality, volume, labelling, and integration), accessing to computing capacity, availability and accessibility of AI talent, and the receptiveness and capabilities of organizations deploying solutions. Some efforts are underway to address this, especially to address accessibility of data for social good, including the Global Data Commons and UN Global Pulse.

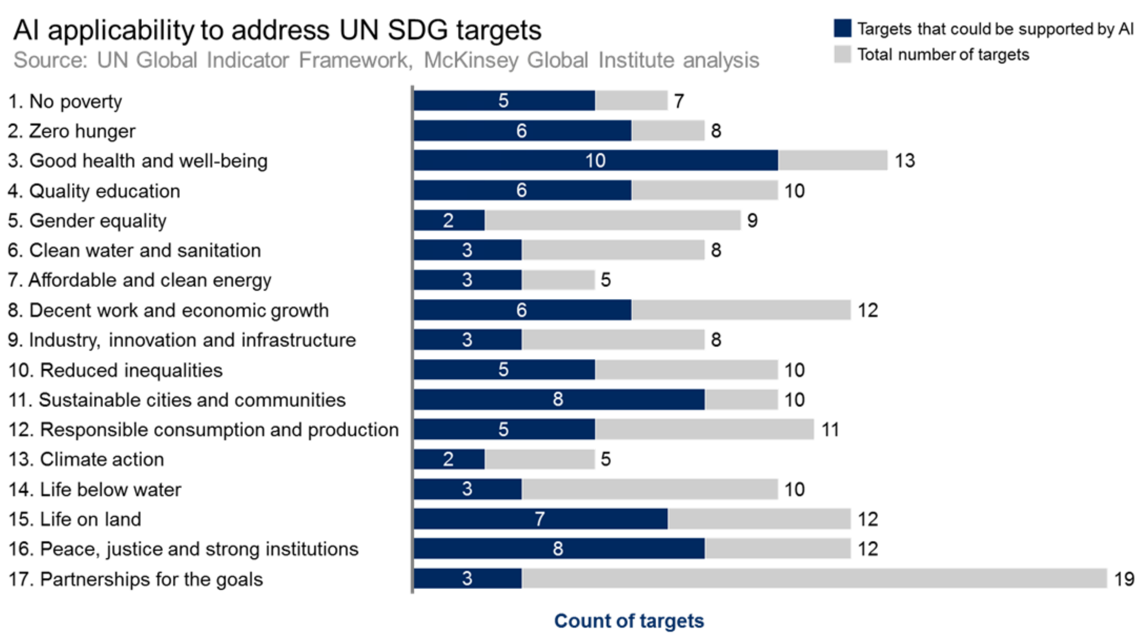

Fig. 8.3b.

Measurement Questions

- How can standardized granular data on AI use cases that impact fairness, human rights, and human dignity be generated?
- How can AI development be integrated into frameworks with social goals, to better plan AI technical development alongside social impacts?
- What measurements can be developed to assess how AI might generate societal threats as well as opportunities?

Chapter Preview

Chapter 9: National Strategies and Global AI Vibrancy

[Table_of_Contents]
[National_Strategies_AI_Vibrancy_Technical_Appendix]

Introduction

This chapter begins by identifying the topics mentioned in official National AI Strategy Radar (NAISR) documents from PricewaterhouseCoopers (PwC). The Appendix documents detailed policy milestones and links to country specific policy documents. The Global AI Vibrancy Tool - a country weighting tool is introduced to aid comparison of countries' global activities, including both a cross-country perspective as well as an intra-country drill down. The tool allows the reader to set the parameters and obtain the perspective they find most relevant. Country pages document key policy milestones accompanied by a country data page for select nations.

There are limitations to overcome in future years' reports. For example, it would be important to know how many official government documents on AI have been published by governments that haven't been translated into English, to help understand what is missing. Similarly, the Global AI Vibrancy will improve with feedback from the community, but also (a) diverse new metrics, (b) more coverage for more developing countries, (c) deeper understanding of causal relationship to inform data-driven decision-making on AI at the national or sub-national level.

National Strategies

The number of official AI strategy documents (both global and national reports) has been increasing over the last few years (Figure. 9.1a). There are several efforts to track and collate national AI strategy documents, including those from UNICRI-FutureGrasp and Future of Life Institute. Other publications have been released by global think tank and thought leadership institutions mentioning the priorities of various nations. These documents can be long and difficult to distill. To support this effort, understand the commonalities and differences of these strategy and overview documents, and observe changes over time, PricewaterhouseCoopers (PwC) has created the National AI Strategy Radar (NAISR) that utilizes natural language processing (NLP) rather than relying on humans to read through the documents. Topic modelling on the documents is conducted to understand the major themes and topics in these documents. Details on country AI policy milestones and methodology can be found in the NAISR Appendix. The non-exhaustive list of global AI reports, strategies and country strategies documents used in the analysis is available here.

Based on 37 analyzed documents, the bar chart shows the percentage of documents mention the topic clusters identified by the topic model. Academic Partnership is present in 94% of the documents, AI R&D in 48% and AI Governance mentioned in over 42% of the documents. Consumer Protection and Fairness is mentioned the fewest times, appearing in 2% of the documents (Figure 9.1b).

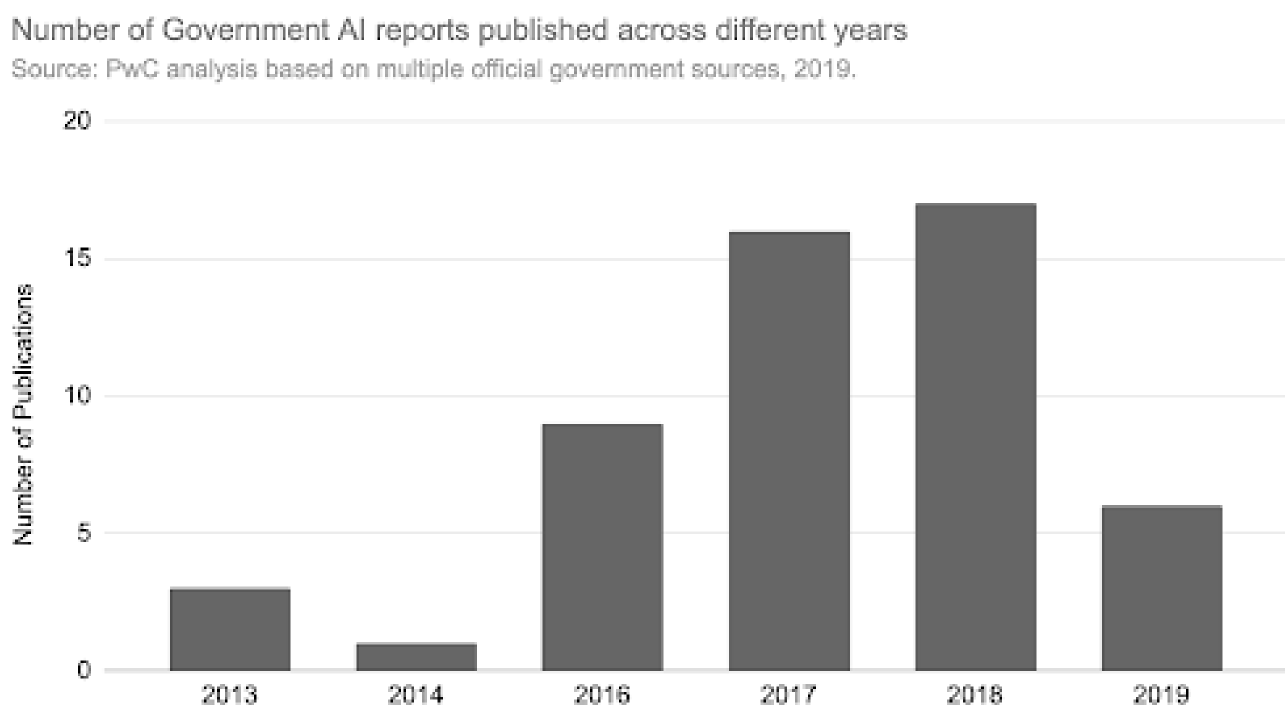

Fig. 9.1a.
Note: Data as of August 2019

National Strategies

Percent of Global and National AI strategy documents mentioning Topics (%)
Source: PwC based on 48 AI Strategy documents.

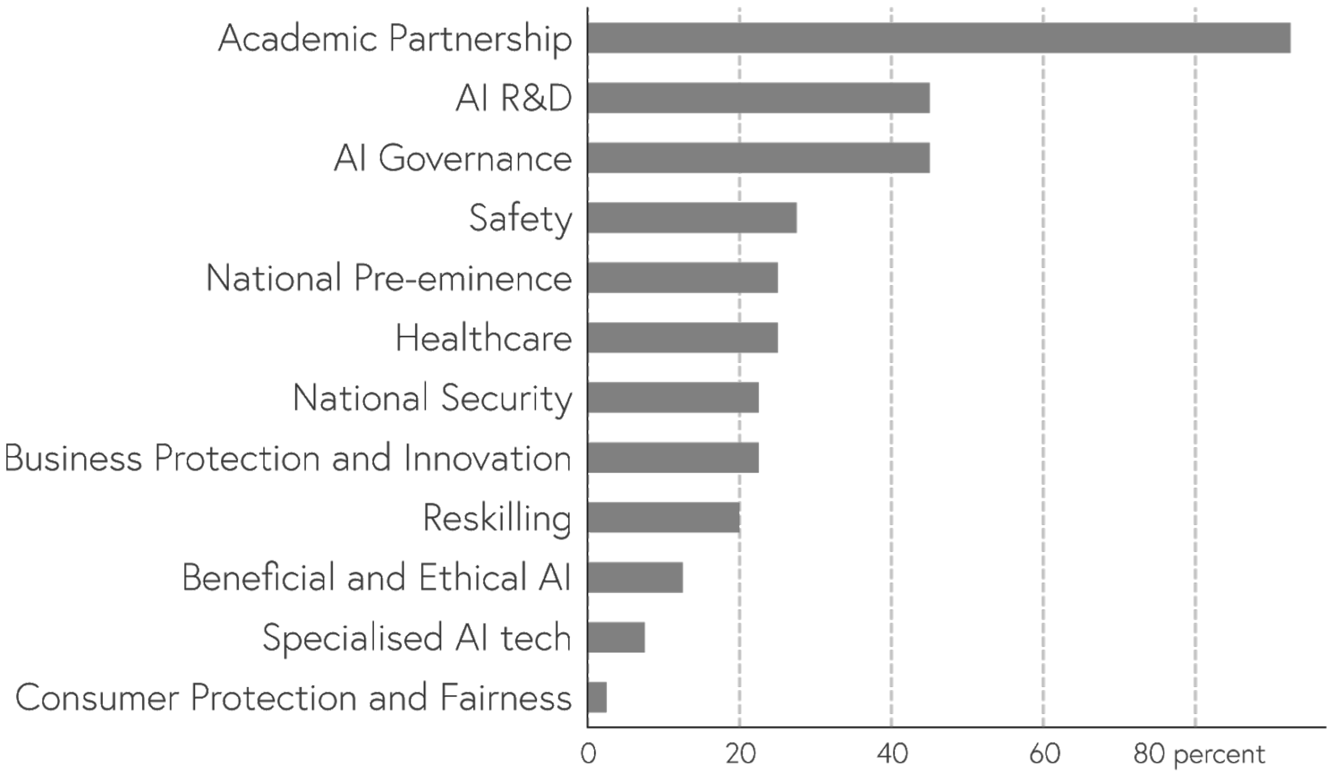

Fig. 9.1b.
Note: Data as of August 2019

National Strategies

A world heatmap shows the number of mentions of countries across the globe in the global sample of AI strategy documents (Figure. 9.2). Countries are developing new strategies constantly. Limitations will exist in sampling official documents until the Index builds an automated crawler for official government AI agencies. Official national strategies documents mentioning Latin America, Africa, and Central Asia are still being acquired, as many countries in these areas are actively exploring AI strategies. The traceability matrix showing the coverage of topics for all documents in the sample (see Appendix Graph). Due to current language limitations, only reports in English or translated to English were considered in this analysis. The 2020 report is building greater translation capacities.

World Map of Countries mentioned in AI documents (official and from major institutions)

Source: PwC NAISR, data as of August 2019 refresh; multiple strategies have been released since

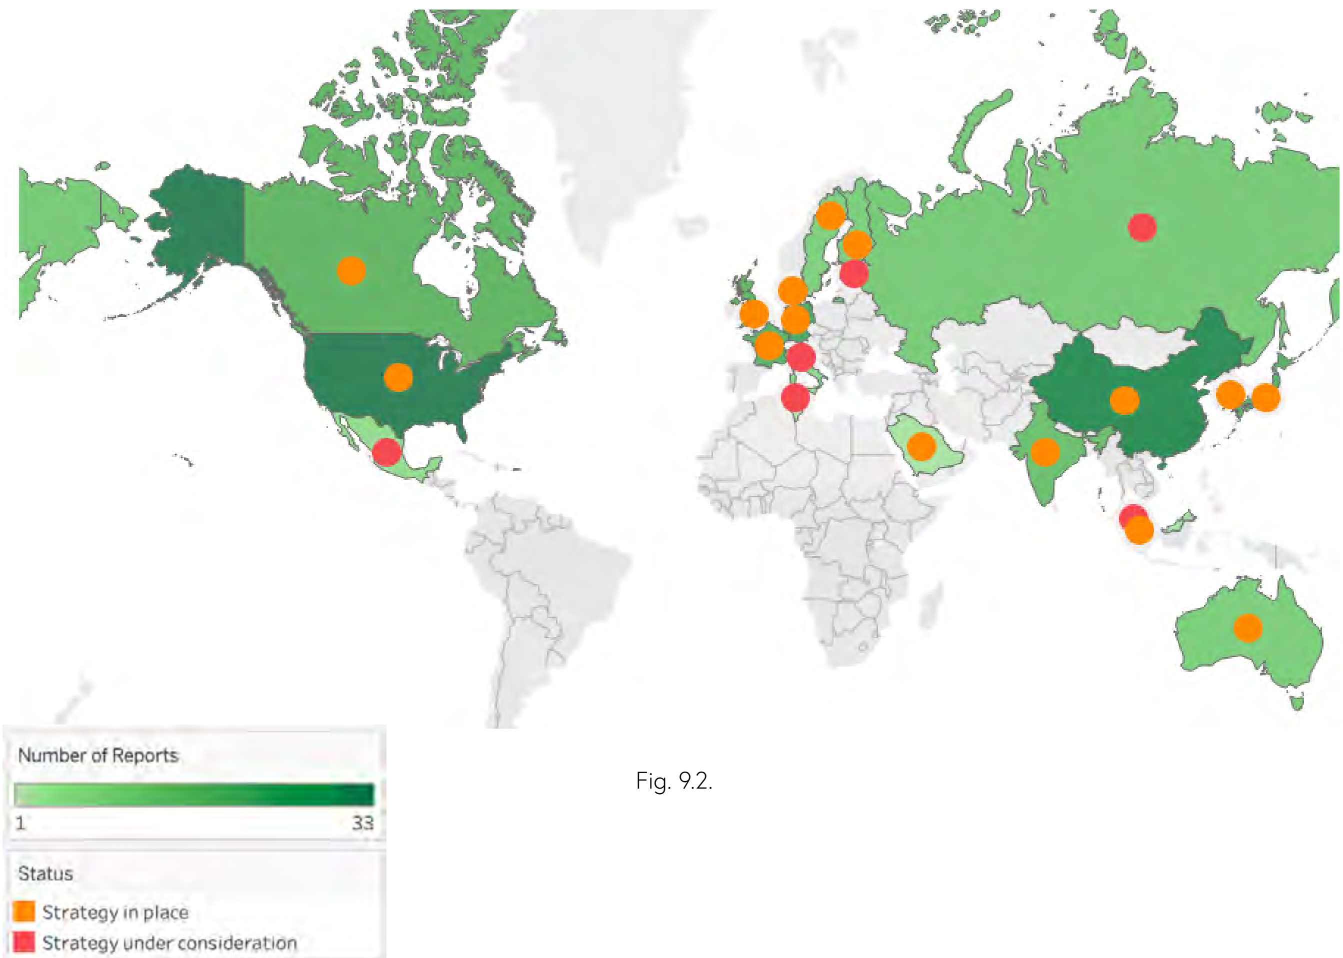

Fig. 9.2.

Global AI Vibrancy Tool

This section summarizes the methodology of the Global AI Vibrancy Tool. The Global AI Vibrancy Tool covers over 28 countries across 34 metrics grouped into three high-level pillars of AI starting in 2015: Research and Development, Economy, and Inclusion. The aggregate indicators are based on several million individual underlying variables, taken from a wide variety of datasources. The data reflect the views on AI from primary data sources and survey from private, public, and NGO sectors worldwide. The metrics are scaled between (0-100) to indicate the relative position of a given country in the global distribution specific to each metric. The Global AI Vibrancy Tool permits meaningful cross country and over time comparisons based on the readers' weighting preference. The underlying source data with detailed description for each indicator are available at vibrancy.aiindex.org.

Country Coverage
The 28 countries covered in the Global AI Vibrancy Tool were selected based on an aggregate data availability threshold of at least 70% (24 out of 34 variables) at the sub-pillar level data availability. The most recent data points for each country were considered in the calculation between 2015 and 2018 as a cutoff year. Meanwhile, each variable had to pass a country-based availability threshold of 50% (28 out of 123 countries). In order to provide transparency and replicability, there was no imputation effort to fill in missing values in the data set. Missing values were noted with 'n/a' and were not considered in the calculation of sub-pillar scores.

Data Sources and Definitions
The abstraction below shows the high-level pillar and sub-pillars covered currently by the Global AI Vibrancy Tool. Each sub-pillar is composed of individual indicators reported in the Global AI Vibrancy codebook. The sub-pillar highlighted in a color denote that metrics about these dimensions are not available (or have not been incorporated) for this version of the Global AI Vibrancy Tool.

The details on data, sources and definition are available in the Appendix. There are 21 metrics used under Research and Development, 10 metrics under Economy, and 5 metrics available under Inclusion.

Global AI Vibrancy

[topics_covered]

Research and Development

- Publication
- Patent
- Conferences
- Education
- Technical Performance

Economy

- Startup Investment
- Corporate Activity
- Public Investment
- Jobs and labor
- Robotic Sales and Trade
- Skill Penetration
- National Strategies

Inclusion

- Gender Diversity
- Public Perception
- Threats

Note: The sub-pillar highlighted in a color denote that metrics about these dimensions are not available (or have not been incorporated) for this version of the Global AI Vibrancy Tool.

Global AI Vibrancy: Country Weighting Tool

To aid data-driven decision-making and policy strategies, the Global AI Vibrancy is available as a web tool. The detailed datasets are available here and on **vibrancy.aiindex.org**.

The webtool allows users to adjust weights to each metric based on their individual preference. The **default settings** of the tool allow the user to select between three weighting options:

All weights to midpoint

This button assigns equal weights to all indicators.

Only absolute metrics

This button assigns maximum weights to absolute metrics. Per capita metrics are not considered.

Only per capita metrics

This button assigns maximum weights to per capita metrics. Absolute metrics are not considered.

The user can adjust the weights to each metric based on their preference.

The charts automatically update when any weight is changed.

The user can select "Global" or "National" view to visualize the results. The "Global" view offers a cross country comparative view based on the weights selected by the user. The "National" view offers country deep dive to assess which AI indicators (or attributes) a given country is relatively better at. The country-metric specific values are scaled (0-100), where 100 indicates that a given country has the highest number in the global distribution for that metric and conversely small numbers like 0 or 1 indicates relatively low values in the global distribution This can help identify areas for improvement and identify national policy strategies to support a vibrant AI ecosystem.

The heatmap below shows 28 countries against 34 metrics in 2018 (Figure 9.4). The color spectrum is between scaled values between 0-100 for each metric (light blue to dark blue spectrum). For example, 100 (blue) for Singapore in AI journal publications in per capita terms represents that Singapore has the highest number. Similarly, black indicates "NA" to denote that data is unavailable for a given country.

AI Vibrancy: Normalized Distribution (0-100) for 28 Countries on 34 Metrics, 2018

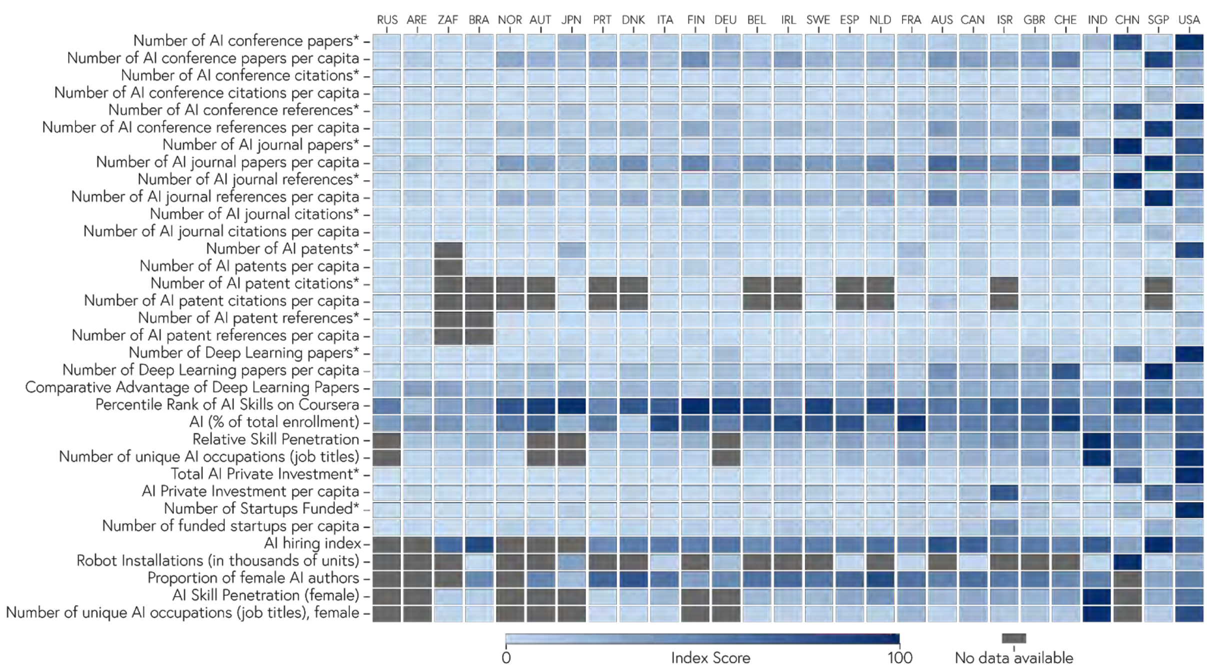

Fig. 9.4.

Country Pages

Country pages provide succinct details on country policy milestones followed by a data page on the respective country. Here, the country policy details are limited to eight countries (key advanced economies and emerging markets) in addition to stock taking of multilateral and regional AI policy developments. Detailed policy milestones with links to official national AI documents are available for over 26 countries is available in the Appendix. The short country policy discussion is followed by country data page so readers can easily lookup available indicators for 2018 to inform country decisions grounded in data.

Country Page: Brazil

In Brazil, broader innovation or government transformation strategies include, but do not focus on, AI. Brazil has not yet published a dedicated artificial intelligence strategy, but the Brazlian government has addressed AI through related initiatives:

2017. Brazil launched the Internet of Things (IoT) National Action Plan. The plan is aimed at positioning the country in the forefront of technology development within the next five years, largely by utilizing AI advancements. Emphasis will be made on health, smart cities, industrial, and rural areas.

2018. The Brazlian government launched the E-Digital strategy. The strategy addresses digital transformation, including AI, while protecting its citizens rights and maintaining privacy, developing an action plan for new technologies, and working with other countries to develop new technologies.

To date, Brazil has most notably implemented AI in facial recognition systems (mainly in criminal establishment and airports). Courts are also being increasingly helped by artificial intelligence technologies, with a focus on automated decision-making, identifying inconsistencies in legal data, analyzing hiring processes, national trading and investments.

Brazil

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	7
2. Number of AI conference papers per capita	4
3. Number of AI conference citations*	0
4. Number of AI conference citations per capita	0
5. Number of AI conference references*	6
6. Number of AI conference references per capita	2
Journal Publications	
7. Number of AI journal papers*	5
8. Number of AI journal papers per capita	4
9. Number of AI journal citations*	1
10. Number of AI journal citations per capita	0
11. Number of AI journal references*	5
12. Number of AI journal references per capita	2
Innovation > Patents	
13. Number of AI patents*	0
14. Number of AI patents per capita	0
15. Number of AI patent citations*	NA
16. Number of AI patent citations per capita	NA
17. Number of AI patent references*	NA
18. Number of AI patent references per capita	NA
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	2
20. Number of Deep Learning papers per capita	1
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	22

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	36
23. AI (% of total enrollment)	25
24. Relative Skill Penetration	22
25. Number of unique AI occupations (job titles)	14
Labor	
26. AI hiring index	84
Investment	
27. Total Amount of Funding*	0
28. Total per capita Funding	0
29. Number of Startups Funded*	1
30. Number of funded startups per capita	0
Robot Installations	
31. Robot Installations (in thousands of units)	1

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	50
33. AI Skill Penetration (female)	9
34. Number of unique AI occupations	2

Country Profile: China

- Prior to the **1980s**, China's interest in AI was focusing more on the theoretical underpinnings of AI and its possible links with contemporary political ideology. AI research in China remained fairly academic until the turn of the millennium, when large Chinese technology firms like Tencent and Baidu began to emerge, offering the opportunity for the government to collaborate with corporations on AI solutions. Since then, this link has grown, as the Chinese government works ever closer with local corporations in the collection and analysis of data for further AI development.

- **June 2017. Launch of the Next Generation AI Development Plan**
China makes one of the biggest pushes towards AI world dominance after announcing "A Next Generation AI Development Plan." For the first time, China announced its plan to become the global leader in AI by 2030.

China

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	80
2. Number of AI conference papers per capita	6
3. Number of AI conference citations*	7
4. Number of AI conference citations per capita	0
5. Number of AI conference references*	76
6. Number of AI conference references per capita	4
Journal Publications	
7. Number of AI journal papers*	100
8. Number of AI journal papers per capita	12
9. Number of AI journal citations*	28
10. Number of AI journal citations per capita	1
11. Number of AI journal references*	100
12. Number of AI journal references per capita	6
Innovation > Patents	
13. Number of AI patents*	8
14. Number of AI patents per capita	0
15. Number of AI patent citations*	0
16. Number of AI patent citations per capita	0
17. Number of AI patent references*	1
18. Number of AI patent references per capita	0
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	49
20. Number of Deep Learning papers per capita	3
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	46

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	83
23. AI (% of total enrollment)	43
24. Relative Skill Penetration	60
25. Number of unique AI occupations (job titles)	36
Labor	
26. AI hiring index	33
Investment	
27. Total Amount of Funding*	77
28. Total per capita Funding	7
29. Number of Startups Funded*	21
30. Number of funded startups per capita	1
Robot Installations	
31. Robot Installations (in thousands of units)	99

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	NA
33. AI Skill Penetration (female)	NA
34. Number of unique AI occupations	NA

Country Profile: France

March 2018. President Emmanuel Macron unveiled France's €1.5B plan to transform France into a global leader in AI. The plan draws heavily from the report, "For a Meaningful Artificial Intelligence: Towards a French and European Strategy," in which Cédric Villani, France's famed mathematician and Deputy for the Essonne, outlined a number of policies and initiatives for the government to consider.

The plan consists of four components: (1) the launch of the National Artificial Intelligence Programme, which will create a network of four or five research institutes across France; (2) an open data policy to drive the adoption and application of AI in sectors where France already has the potential for AI excellence, such as healthcare; (3) a regulatory and financial framework to support the development of domestic "AI champions;" (4) regulations for ethics.

In total, the government will invest €1.5 billion in AI by the end of the current five-year term. Details for the following have not been released, but €700 million will go towards research, €100 million this year to AI startups and companies, €70 million annually through France's Public Investment Bank, and $400 million to industrial projects in AI.The Villani report recommended focusing on four sectors (healthcare, transportation, environment, and defence).

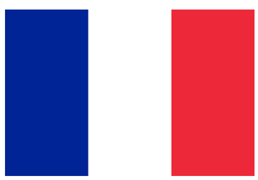

France

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	11
2. Number of AI conference papers per capita	17
3. Number of AI conference citations*	1
4. Number of AI conference citations per capita	1
5. Number of AI conference references*	11
6. Number of AI conference references per capita	13
Journal Publications	
7. Number of AI journal papers*	9
8. Number of AI journal papers per capita	21
9. Number of AI journal citations*	3
10. Number of AI journal citations per capita	2
11. Number of AI journal references*	12
12. Number of AI journal references per capita	16
Innovation > Patents	
13. Number of AI patents*	19
14. Number of AI patents per capita	9
15. Number of AI patent citations*	0
16. Number of AI patent citations per capita	1
17. Number of AI patent references*	2
18. Number of AI patent references per capita	4
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	7
20. Number of Deep Learning papers per capita	10
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	22

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	64
23. AI (% of total enrollment)	95
24. Relative Skill Penetration	34
25. Number of unique AI occupations (job titles)	31
Labor	
26. AI hiring index	55
Investment	
27. Total Amount of Funding*	4
28. Total per capita Funding	7
29. Number of Startups Funded*	8
30. Number of funded startups per capita	6
Robot Installations	
31. Robot Installations (in thousands of units)	4

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	62
33. AI Skill Penetration (female)	35
34. Number of unique AI occupations	22

[National_Strategies_AI_Vibrancy_Technical_Appendix]
[Access_Data]

Country Profile: Germany

2017. The Federal Ministry of Education and Research launched a government aid campaign in the field of machine learning. Subsequently, it funded The Platform Learning Systems (an expert AI platform running from 2017 to 2022) and the Automated and Networked Driving Project. The Federal Ministry of Transport and Digital Infrastructure also published "Ethics Commission: Automated and Connected Driving," with 20 ethical guidelines for self-driving cars.

November 2018. Germany launched its Artificial Intelligence Strategy and allocated €3B for investment in AI R&D. The strategy was developed by the Economic Affairs Ministry, the Research Ministry, and the Labour Ministry. The strategy focuses on three objectives: (1) making Germany and Europe global leaders in AI; (2) developing AI which serves the good of society; (3) integrating AI into society in the active political context.

Previously, the German Institute for Innovation and Technology within the Federal Ministry for Economic Affairs and Energy found that AI will add approximately €32 billion to Germany's manufacturing output over the next five years.

Germany

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	18
2. Number of AI conference papers per capita	23
3. Number of AI conference citations*	2
4. Number of AI conference citations per capita	2
5. Number of AI conference references*	16
6. Number of AI conference references per capita	15
Journal Publications	
7. Number of AI journal papers*	13
8. Number of AI journal papers per capita	26
9. Number of AI journal citations*	4
10. Number of AI journal citations per capita	2
11. Number of AI journal references*	17
12. Number of AI journal references per capita	17
Innovation > Patents	
13. Number of AI patents*	9
14. Number of AI patents per capita	4
15. Number of AI patent citations*	1
16. Number of AI patent citations per capita	1
17. Number of AI patent references*	1
18. Number of AI patent references per capita	2
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	16
20. Number of Deep Learning papers per capita	17
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	26

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	95
23. AI (% of total enrollment)	53
24. Relative Skill Penetration	NA
25. Number of unique AI occupations (job titles)	NA
Labor	
26. AI hiring index	59
Investment	
27. Total Amount of Funding*	2
28. Total per capita Funding	3
29. Number of Startups Funded*	4
30. Number of funded startups per capita	2
Robot Installations	
31. Robot Installations (in thousands of units)	17

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	49
33. AI Skill Penetration (female)	NA
34. Number of unique AI occupations	NA

Country Profile: India

- **February 2018**. A Task Force was assigned by MoD to study the strategic implementation of AI for National Security and Defense.

- **June 2018**. The Indian government's think-tank NITI Aayog defined a national policy on AI in a working paper titled National Strategy for AI (#AIforAll). India has taken a unique approach to its national AI strategy by focusing on how it can leverage AI not only for economic growth, but also for social inclusion. The strategy aims to (1) enhance and empower Indians with the skills to find quality jobs, (2) invest in research and sectors that can maximize economic growth and social impact, and (3) scale Indian-made AI solutions to the rest of the developing world. The government wants to establish India as an "AI Garage," meaning that if a company can deploy an AI in India, it will then be applicable to the rest of the developing world.

  The strategy clarifies five major sectors that AI research in India will focus on – healthcare, agriculture, education, smart cities and infrastructure, and smart mobility and transportation. To pave the way for these advancements, the Indian government has doubled its allocation to the 'Digital India' program to $480m (₹3,073 crore) in 2018-19.

India

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	20
2. Number of AI conference papers per capita	2
3. Number of AI conference citations*	1
4. Number of AI conference citations per capita	0
5. Number of AI conference references*	13
6. Number of AI conference references per capita	1
Journal Publications	
7. Number of AI journal papers*	28
8. Number of AI journal papers per capita	3
9. Number of AI journal citations*	5
10. Number of AI journal citations per capita	0
11. Number of AI journal references*	19
12. Number of AI journal references per capita	1
Innovation > Patents	
13. Number of AI patents*	1
14. Number of AI patents per capita	0
15. Number of AI patent citations*	0
16. Number of AI patent citations per capita	0
17. Number of AI patent references*	0
18. Number of AI patent references per capita	0
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	6
20. Number of Deep Learning papers per capita	0
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	31

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	41
23. AI (% of total enrollment)	50
24. Relative Skill Penetration	100
25. Number of unique AI occupations (job titles)	99
Labor	
26. AI hiring index	73
Investment	
27. Total Amount of Funding*	1
28. Total per capita Funding	0
29. Number of Startups Funded*	5
30. Number of funded startups per capita	0
Robot Installations	
31. Robot Installations (in thousands of units)	3

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	54
33. AI Skill Penetration (female)	100
34. Number of unique AI occupations	100

[National_Strategies_AI_Vibrancy_Technical_Appendix]
[Access_Data]

Country Profile: The Netherlands

In **2018**, AINED*, the public-private partnership on AI, has formulated AI Voor Nederland — a first draft for a Dutch National AI strategy. The setup will provide a concrete action plan to make AI a national priority, with the Netherlands seeing potential for AI development in the areas of health, agriculture, mobility, and decarbonization. AINED is currently working in a public-private context to turn the report into a concrete action plan, which should be launched soon.

The report includes a wide range of measures that governments and businesses can take to help the Netherlands further its excellent standing in this field, and provides an interesting focus on education. A shortage of talent, for instance, can be obviated by making it easier for international students to extend their stay in the Netherlands after graduating. The Netherlands could also improve its collaboration in existing chains, develop a national AI research centre of high repute, serve as a catalyst for new businesses, and make better use of available data. Universities are already conducting good technical research; for instance, the University of Amsterdam collaborating with the municipality and other businesses to create Amsterdam's AI Hub.

The central government is, partly in response to the AINED report, also preparing an action plan.

*AINED was founded to map the position of the Netherlands in AI development and is a public-private partnership between TopTeam ICT, Dutch employer federation VNO-NCW, business group MKB Nederland, Innovation Center for Artificial Intelligence, Netherlands Organisation for Scientific Research (NWO) and Netherlands Organisation for Applied Scientific Research (TNO).

The Netherlands

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	4
2. Number of AI conference papers per capita	22
3. Number of AI conference citations*	1
4. Number of AI conference citations per capita	3
5. Number of AI conference references*	4
6. Number of AI conference references per capita	17
Journal Publications	
7. Number of AI journal papers*	5
8. Number of AI journal papers per capita	47
9. Number of AI journal citations*	2
10. Number of AI journal citations per capita	5
11. Number of AI journal references*	7
12. Number of AI journal references per capita	34
Innovation > Patents	
13. Number of AI patents*	0
14. Number of AI patents per capita	1
15. Number of AI patent citations*	NA
16. Number of AI patent citations per capita	NA
17. Number of AI patent references*	0
18. Number of AI patent references per capita	0
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	5
20. Number of Deep Learning papers per capita	23
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	29

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	86
23. AI (% of total enrollment)	42
24. Relative Skill Penetration	23
25. Number of unique AI occupations (job titles)	13
Labor	
26. AI hiring index	62
Investment	
27. Total Amount of Funding*	0
28. Total per capita Funding	2
29. Number of Startups Funded*	1
30. Number of funded startups per capita	4
Robot Installations	NA
31. Robot Installations (in thousands of units)	

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	82
33. AI Skill Penetration (female)	42
34. Number of unique AI occupations	7

Country Profile: Singapore

AI has been identified as one of four frontier technologies which are essential to growing Singapore's economy. Singapore aims to advance its vision to be a leading Digital Economy and Smart Nation, continually embracing digital transformation and reinventing itself to remain globally competitive. In doing so, Singapore focuses on the technical capabilities, technology investments, and regulatory requirements through the following core initiatives:

May 2017. The Singaporean government launched AI Singapore (AISG) with $150 million in funding to catalyse, synergise and boost Singapore's AI capabilities. Today, AISG is Singapore's premier national research and innovation programme in AI.

2018. The Singaporean government established an Advisory Council on the Ethical Use of AI and Data, an industry-led initiative to examine legal and ethical issues raised by commercial deployment of AI. Members comprise international leaders in AI such as Google, Microsoft and Alibaba. The Research Programme on the Governance of AI and Data was also set up with the Singapore Management University.

November 2019. Singapore's National AI Strategy (NAIS) was unveiled by the Deputy Prime Minister. The full NAIS is available publicly.

Davos 2019. At Davos the Singaporean government announced it is working with the World Economic Forum's Centre for Fourth Industrial Revolution (WEF C4IR) to help drive the ethical and responsible deployment of artificially intelligent technologies. Singapore's Model AI Governance Framework is the first of its kind to exist throughout Asia and provides detailed guidance to private sector organizations to address key ethical and governance issues when building, deploying and investing in AI solutions. Singapore has long been pushing to become a global leader in AI, and this Model Framework will be welcomed by those who work with this emerging technology.

Singapore

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	5
2. Number of AI conference papers per capita	87
3. Number of AI conference citations*	1
4. Number of AI conference citations per capita	19
5. Number of AI conference references*	7
6. Number of AI conference references per capita	92
Journal Publications	
7. Number of AI journal papers*	3
8. Number of AI journal papers per capita	100
9. Number of AI journal citations*	3
10. Number of AI journal citations per capita	20
11. Number of AI journal references*	7
12. Number of AI journal references per capita	100
Innovation > Patents	
13. Number of AI patents*	1
14. Number of AI patents per capita	4
15. Number of AI patent citations*	NA
16. Number of AI patent citations per capita	NA
17. Number of AI patent references*	0
18. Number of AI patent references per capita	2
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	6
20. Number of Deep Learning papers per capita	100
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	39

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	93
23. AI (% of total enrollment)	34
24. Relative Skill Penetration	32
25. Number of unique AI occupations (job titles)	11
Labor	
26. AI hiring index	100
Investment	
27. Total Amount of Funding*	3
28. Total per capita Funding	65
29. Number of Startups Funded*	3
30. Number of funded startups per capita	26
Robot Installations	
31. Robot Installations (in thousands of units)	3

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	33
33. AI Skill Penetration (female)	27
34. Number of unique AI occupations	9

Country Profile: The United States

February 2019. Launch of the American AI Initiative
In February 2019, the President signed an Executive Order launching the American AI Initiative, which will take a multipronged approach to accelerating America's national leadership in AI. The Executive Order states that the Federal Government will have a central role not only in facilitating AI R&D, but also in promoting trust, training people for a changing workforce, protecting national security, enhancing collaboration with foreign partners and the private sector.

June 2019. Launch of the US AI R&D Strategic Plan
In June 2019, the White House's AI R&D Strategic Plan defines several key areas of priority focus for the Federal agencies that invest in AI. These areas of strategic AI R&D focus include: (1) continued long-term investments in AI (2) effective methods for human-AI collaboration (3) understanding and addressing the ethical, legal, and societal implications for AI (4) ensuring the safety and security of AI (5) developing shared public datasets and environments for AI training and testing (6) measuring and evaluating AI technologies through standards and benchmark (7) better understanding the National AI R&D workforce needs, and (8) expanding public-private partnerships to accelerate AI advances.

2019 marked the biggest year in funding, both federal and private, for artificial intelligence ventures yet. For 2020, the President's Budget prioritizes AI as one of four key Industries of the Future to invest in. Annual federal spending on non-defence-related AI research is set to jump to nearly $1 billion. That figure represents an increase, given that agencies including the US defence department and non-defence related entities spent about US$1 billion on AI research in 2016.

September 2018. DARPA announced the "AI Next" campaign, a multi-year investment $2b+ in new and existing programs. Key areas of the campaign include automating critical DoD business processes. AI Next builds on DARPA's five decades of AI technology creation to define and to shape the future, always with the Department's hardest problems in mind.

October 2019. The Defense Innovation Board, a panel of 16 prominent technologists advising the Pentagon, voted to approve AI ethics principles for the Department of Defense. The report includes 12 recommendations for how the US military can apply ethics in the future for both combat and non-combat AI systems.

November 2019. The interim report was released by the National Security Commission on AI.

The United States

Research and Development

Conference Publications	Scaled (0-100)
1. Number of AI conference papers*	100
2. Number of AI conference papers per capita	33
3. Number of AI conference citations*	21
4. Number of AI conference citations per capita	6
5. Number of AI conference references*	100
6. Number of AI conference references per capita	24
Journal Publications	
7. Number of AI journal papers*	80
8. Number of AI journal papers per capita	40
9. Number of AI journal citations*	29
10. Number of AI journal citations per capita	3
11. Number of AI journal references*	88
12. Number of AI journal references per capita	23
Innovation > Patents	
13. Number of AI patents*	84
14. Number of AI patents per capita	9
15. Number of AI patent citations*	5
16. Number of AI patent citations per capita	2
17. Number of AI patent references*	18
18. Number of AI patent references per capita	8
Journal Publications > Deep Learning	
19. Number of Deep Learning papers*	100
20. Number of Deep Learning papers per capita	27
21. Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	32

Economy

Skills	Scaled (0-100)
22. Percentile Rank of AI Skills on Coursera	81
23. AI (% of total enrollment)	65
24. Relative Skill Penetration	76
25. Number of unique AI occupations (job titles)	100
Labor	
26. AI hiring index	65
Investment	
27. Total Amount of Funding*	100
28. Total per capita Funding	37
29. Number of Startups Funded*	100
30. Number of funded startups per capita	14
Robot Installations	
31. Robot Installations (in thousands of units)	26

Inclusion

Gender Diversity	Scaled (0-100)
32. Proportion of female AI authors	53
33. AI Skill Penetration (female)	60
34. Number of unique AI occupations	82

[National_Strategies_AI_Vibrancy_Technical_Appendix]
[Access_Data]

Multilateral and Regional AI Policy

United Nations Activity on Artificial Intelligence is a joint-effort between ITU and 32 UN agencies and bodies, all partners of 2018's AI for Good Global Summit, this report provides information about the diverse and innovative activities related to artificial intelligence (AI) across the UN system.

The WTO foresees that AI will transform the administration of the world trading system. While the world trading system will continue to be tested, they foresee that it will endure and improvements will be made to make it effective with respect to all aspects of global need.

In 2019 presentation Multilateral Trading System and WTO Reform: Making Globalization Serve Society, Joseph Stiglitz argues that as we reform the WTO—to strengthen the rules-based multilateral system—we need to keep paramount that trade is not an end in itself but a means to an end, enhancing the well-being of all citizens of the world.

The High-Level Expert Group on Artificial Intelligence (AI HLEG) has as a general objective to support the implementation of the European Strategy on Artificial Intelligence. HLEG has also released the Ethics Guidelines for Trustworthy AI.

The European AI Alliance constitutes a key forum engaged in a broad and open discussion of all aspects of Artificial Intelligence development and its impacts.

In May 2019, Forty-two countries adopted new OECD Principles on Artificial Intelligence, agreeing to uphold international standards that aim to ensure AI systems are designed to be robust, safe, fair and trustworthy.

OECD Global AI Observatory provides evidence and guidance on AI metrics, policies and practices, facilitating dialogue and sharing best practices on AI policies.

OECD Principles on Artificial Intelligence complements existing OECD standards in areas such as privacy, digital security risk management, and responsible business conduct in the context of AI. The book OECD Artificial Intelligence in Society delineates a plan for implementing the Principles in practice. The OECD Private Equity Investment in Artificial Intelligence shows important increases in investments in AI startups. In 2020, they will release the OECD AI Policy Observatory.

Appendix Preview

Technical Appendix

Papers on Scopus

Source

Elsevier's Scopus database of scholarly publications, which has indexed more than 75 million documents. This data was compiled by Elsevier. In depth methodology on paper indexing, affiliations, geographic coverage, and titles can be found on the Scopus Content Coverage Guide.

Methodology

Scopus tags its papers with keywords, publication dates, country affiliations, and several other bibliographic information.

The Elsevier AI Classifier leveraged the following features that were extracted from the Scopus records that were returned as a result of querying against the provided @ 800 AI search terms. Each record fed into the feature creation also maintained a list of each search term that hit for that particular record:

- hasAbs – Boolean value whether or not the record had an abstract text section in the record (e.g. some records are only title and optional keywords)
- coreCnt – number of core-scored search terms present for the record
- mediumCnt – number of medium-scored search terms present for the record
- lowCnt – number of low-scored search terms present for the record
- totalCnt – total number of search terms present for the record
- pcntCore – coreCnt/totalCnt
- pcntMedium – mediumCnt/totalCnt
- pcntLow – lowCnt/totalCnt
- totalWeight = 5*coreCnt + 3*mediumCnt + 1*lowCnt
- normWeight = if (has Abs) { totalWeight / (title.length + abstract.length) } else { totalWeight/title.length}
- hasASJC – Boolean value – does the record have an associated ASJC list
- isAiASJC – does ASJC list contain 1702
- isCompSciASJC does ASJC list contain a 17XX ASJC code - ("1700", "1701", "1702", "1703", "1704", "1705", "1706", "1707", "1708", "1709", "1710", "1711", "1712")
- isCompSubj – Does the Scopus record have a ComputerScience subject code associated with it. This should track 1:1 to isCompSciASJC, but added in case they didn't.
- pcntCompSciASJC – percentage of ASJC codes for record that are from the CompSci ASJC code lis

Details on Elsevier's dataset defining AI, country affiliations, and AI sub-categories can be found in the 2018 AI Index Report Appendix.

Europe is defined as EU44.

Datasets

FWCI and FWDI sheets

Published Papers: Citation Impact By Region
FWCI and FWDI is Field-Weighted Citation (Download) Impact, a normalized score for citation/download impact - normalized for age of publication, subject area, and type of publication. This is necessary, as number of citations is strongly influenced by these factors - e.g. reviews attract more citations than articles, older publications have more time to accrue citations and so on.

WLD

WLD is global (WORLD) therefore the total number of all publications. Individual regions and/or countries do not add up to WLD as publications can be collaboratively published in the US, China and Europe. This deduplication issue means that country counts generally don't add up to regional ones

Nuance

- The Scopus system is retroactively updated. As a result, the number of papers for a given query may increase over time.
- Members of the Elsevier team commented that data on papers published after 1995 would be most reliable, so we use 1996 as a starting year for Scopus data.

Nuances specific to AI publications by region

- Papers are double counted if they are tagged to multiple regions. This explains why top line numbers in a given year may not match last year's annual paper count.
- "Other" includes all other countries that have published AI paper(s) on Scopus.

Nuances specific to publications by topic

- The 2017 AI Index Report only showed AI papers within the CS category. In 2018 and 2019, all papers tagged as AI were included, regardless of whether they fell into the larger CS category.
- Elsevier has a subject category called 'AI', which is a subset of 'CS' - but this is relevant only for a subject category approach to defining AI papers. The methodology used for the report includes all papers, since increasingly not all AI papers fall into CS.

Nuances specific to methodology

- The entire data collection process was done by Elsevier internally — the AI Index was not involved in the keyword selection process or the counting of relevant papers.
- The boundaries of AI are difficult to establish, in part because of the rapidly increasing applications in many fields, such as speech recognition, computer vision, robotics, cybersecurity, bioinformatics, and healthcare. But limits are also difficult to define because of AI's methodological dependency on many areas such as logic, probability and statistics, optimization, photogrammetry, neuroscience, and game theory — to name a few. Given the community's interest in AI bibliometrics, we believe it would be valuable if groups producing these studies would strive for a level of transparency in their methods which supported reproducibility of results, in particular on different underlying bibliographic databases.

Documentation

Methodological documentation may be downloaded here.

AI Training Set

The training set of ~1,500 publications to define the AI field. The set is only the EID (the Scopus identifier of the underlying publications). Publications can be searched and downloaded either from Scopus directly or via the API.

The Elsevier Developer API (https://dev.elsevier.com) provides more details on the different endpoints available. API keys are available through the developers portal.

Return to Research & Development - Journal Publications: AI Papers in All Publications

Elsevier Appendix Graphs

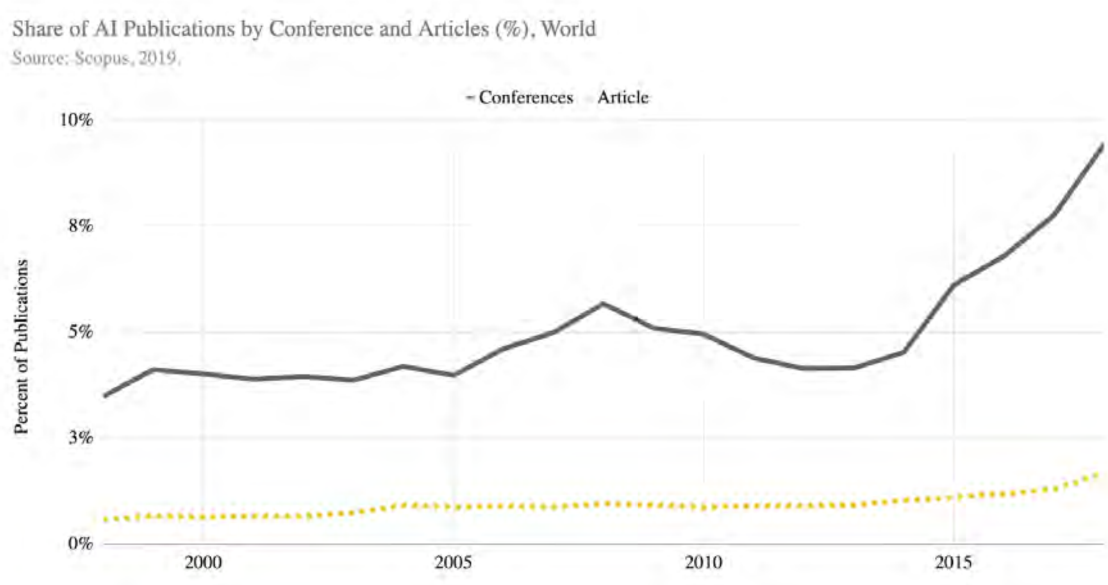

Fig. A1.1a

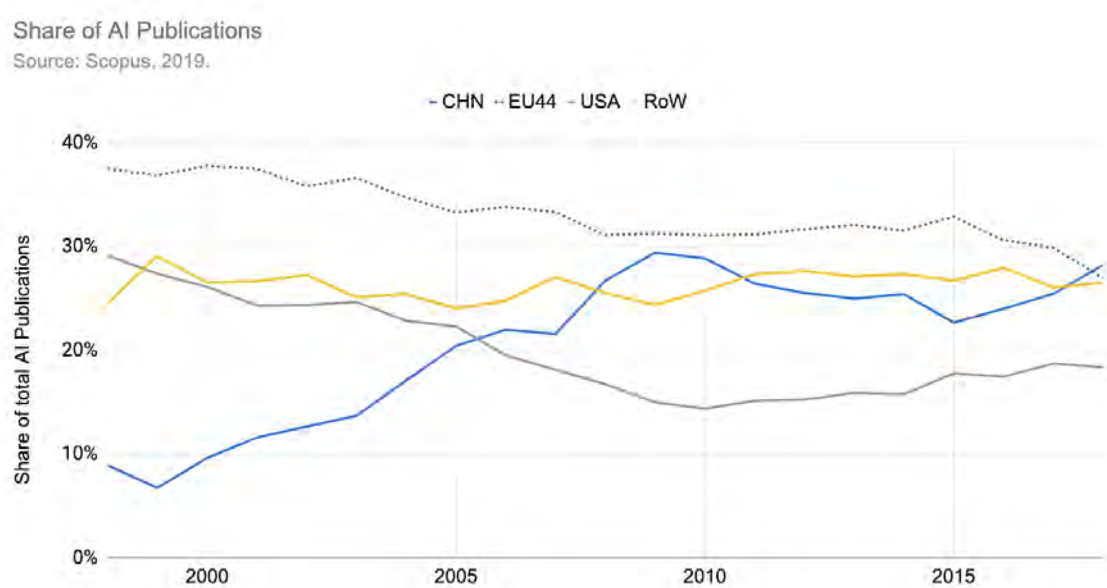

Fig. A1.1b

Return to Research & Development - Journal Publications:
Published Papers: Institutional Affiliation

Institutional Affiliation, Growth

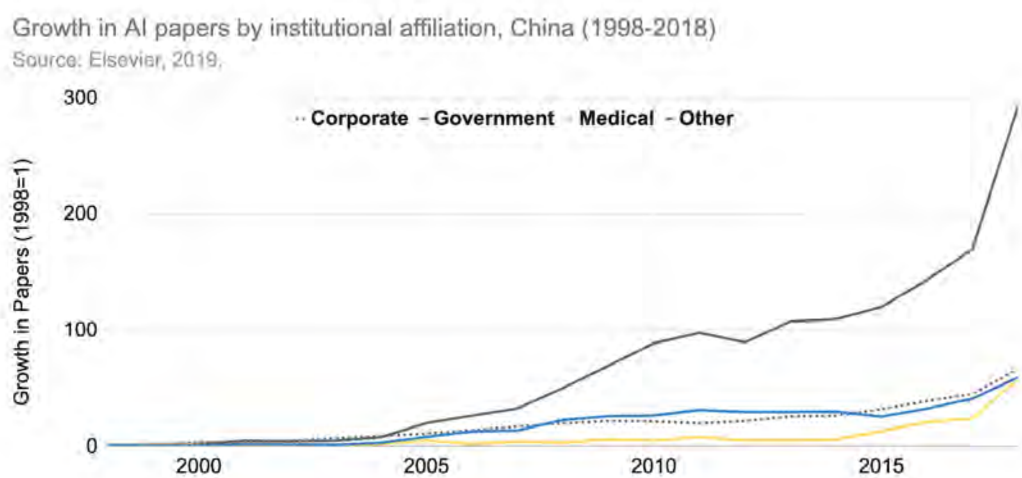

Fig. A1.2a

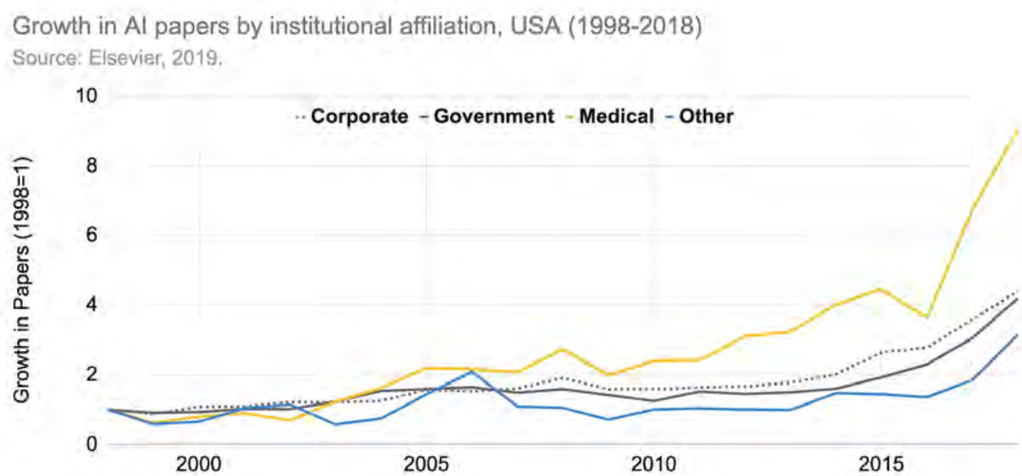

Fig. A1.2b

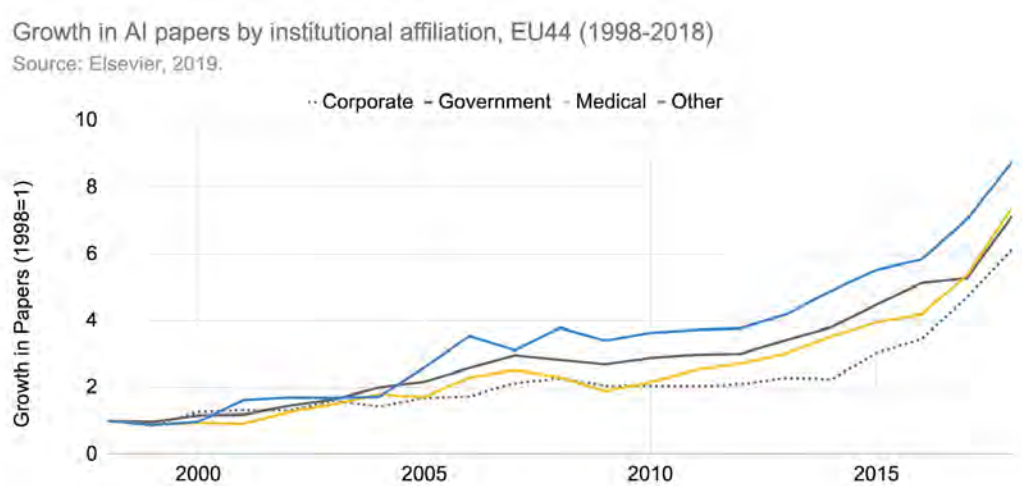

Fig. A1.2b

Return to Research & Development - Journal Publications:
[Published Papers: Institutional Affiliation]

Comparative View of Growth in AI Papers by Regions, Corporate and Government affiliated

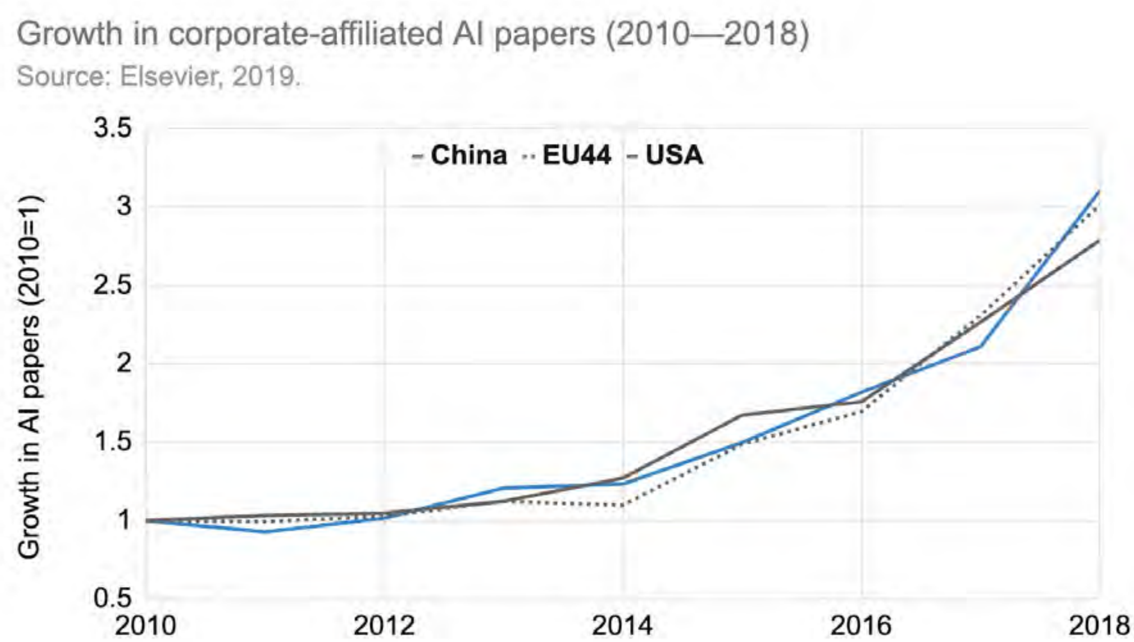

Fig. A1.3a

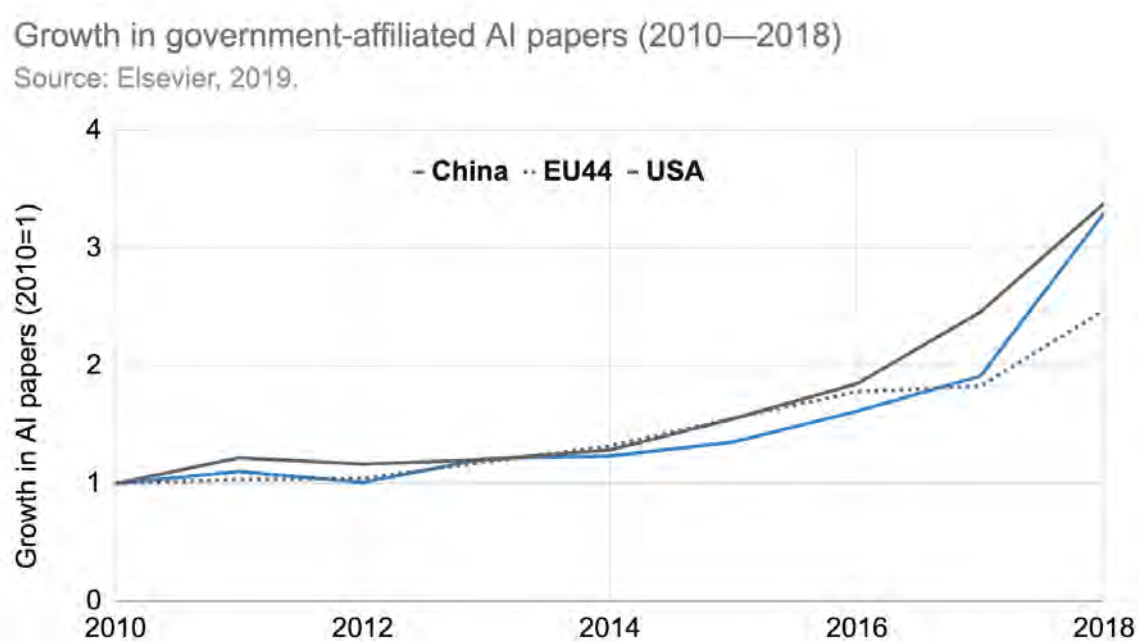

Fig. A1.3b

Return to Research & Development - Journal Publications
[Published Papers: Citation Impact By Region]

Field-Weighted Download Impact

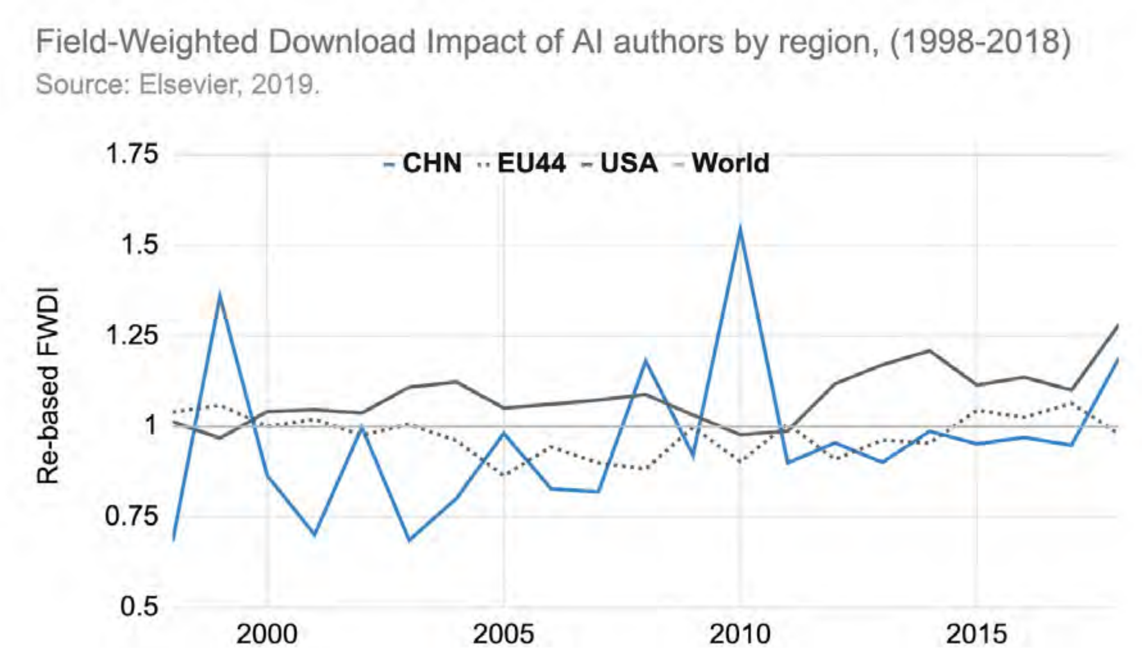

Fig. A1.4

Scholarly AI Output for list of countries, 2014-18

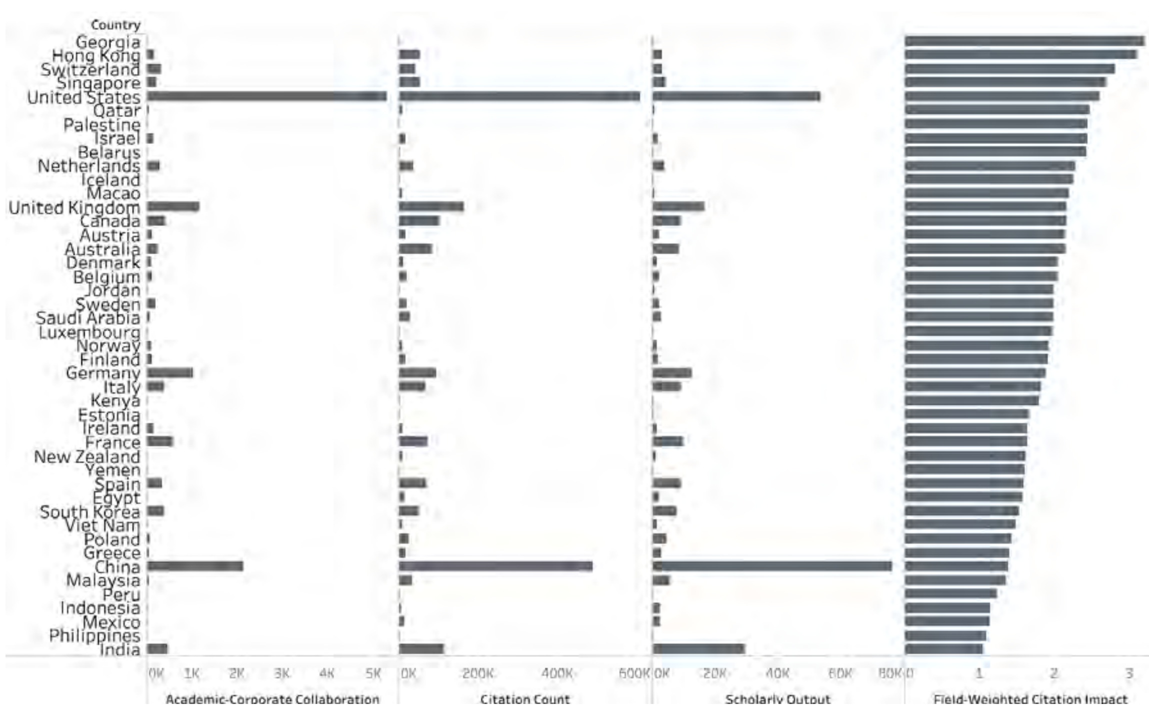

Fig. A1.5

Return to Research & Development - arXiv

Papers on arXiv

Source

arXiv.org is an online archive of research articles in the fields of physics, mathematics, computer science, quantitative biology, quantitative finance, statistics, electrical engineering and systems science, and economics. arXiv is owned and operated by Cornell University. See more information on arXiv.org.

arXiv Methodology

Raw data for our analysis was provided by representatives at arXiv.org. The keywords we selected, and their respective categories, are below:

Artificial intelligence (cs.AI)
Computation and language (cs.CL)
Computer vision and pattern recognition (cs.CV)
Machine learning (cs.LG)
Neural and evolutionary computing (cs.NE)
Robotics (cs.RO)
Machine learning in stats (stats.ML)

For most categories, arXiv provided data years 1999 — 2018. For our analysis, we decided to start at the year 2010 in order to include *Machine Learning in Stats*, which did not exist on arXiv prior.

To see other categories' submission rates on arXiv, see arXiv.org's submission statistics.

Nuance

- Categories are self-identified by authors — those shown are selected as the "primary" category. Thus there is not one automated categorization process. Additionally, the *Artificial intelligence* or *Machine learning* categories may be categorized by other subfields / keywords.
- arXiv team members have shared that participation on arXiv can breed more participation — meaning that an increase in a subcategory on arXiv could drive over-indexed participation by certain communities.
- Growth of papers on arXiv does not reflect actual growth of papers on that topic. Some growth can be attributed to arXiv.org's efforts to increase their paper count, or to the increasing importance of dissemination by AI communities.

AI Index-arXiv full paper search engine

Source

ArXiv papers were filtered by category tags presented in arXiv. Analysis of the arXiv papers was broken into several tasks: Task and sub-task classification, ethics and fairness topic evaluation, and institution/country affiliation.

Methodology

Data Integration

A data pipeline was developed to extract keywords and metrics on Apache Beam. Data is stored in Google Cloud Platform, across an Elasticsearch instance, Google Cloud Storage, and Google BigQuery.

Task and sub-task classification

For task and sub-task classification, we use regex keyword search within the parent class.

Ethics and Fairness Topic Evaluation

A named entity recognition model was trained for ethics and fairness topic evaluation with keywords derived from word2vec. Papers containing ethics entities were identified with a deep bidirectional transformer (BERT) and trained for binary classification.

Institution/country Affiliation

For institution affiliation, a named entity recognition model was trained to filter the output by key terms. After extracting the institution affiliation, a lookup table of global university country codes was used to extract country affiliation. For institutional affiliation outside of academia, regex phrase matching was used on prominent AI technology company contributors.

Return to Research & Development - Microsoft Academic Graph

Microsoft Academic Graph (MAG) Data and Methodology

Source

The Microsoft Academic Graph is a heterogeneous graph containing scientific publication records, citation relationships between those publications, as well as authors, institutions, journals, conferences, and fields of study. This graph is used to power experiences in Bing, Cortana, Word, and in Microsoft Academic. The graph is currently being updated on a weekly basis. Learn more about MAG here.

Methodology

MAG Data Attributions

Each paper is counted exactly once. When a paper has multiple authors/countries, the credit is equally distributed to the unique countries. For example, if a paper has two authors from the US, one from China and one from the UK, then the US, UK, and China get 1/3 each.

Metric

Total number of published papers

Datasets

Combined dataset: OutAiPaperCountByYearDocType

Qualitymetrics Data

Metric

Number of citation counts

Datasets

Combined dataset: OutAiPaperCitationCountryPairByYearDocType

Definition

The citation and reference count represents the number of respective metrics for AI papers collected from ALL papers. For example, in "OutAiPaperCitationCountryPairByYearConf.csv", a row "China, United States, 2016, 14955" means that the China's conference AI papers published in 2016 received 14955 citations from (all) US papers indexed by MAG.

Curating the MAG Dataset and References

Generally speaking, the robots sit on top of a Bing crawler to read everything from the web and have access to the entire web index. As a result, MAG is able to program the robots to conduct more web searches than a typical human can do. This is really helpful in disambiguating entities with same names. For example, for authors, MAG get to additionally use all the CVs and institutional homepages on the web as signals to recognize and verify claims (see [1] and [2] for some details). MAG has found this approach much better than the results of the best of the KDD Cup 2013 competition [3] that use only data from within all publication records and ORCIDs.

[1] https://www.microsoft.com/en-us/research/project/academic/articles/microsoft-academic-uses-knowledge-address-problem-conflation-disambiguation/
[2] https://www.microsoft.com/en-us/research/project/academic/articles/machine-verification-paper-author-claims/
[3] https://www.kaggle.com/c/kdd-cup-2013-author-paper-identification-challenge

The statistics of all CS papers can be found at https://academic.microsoft.com/publications/41008148 and one can further navigate the fields of study hierarchy to see the historical publication volume for any subfields.

Similarly, the page https://preview.academic.microsoft.com/institutions puts all institutions on a map. However, it appears a bug at the website is preventing the map from showing up consistently. Nevertheless, MAG can provide you with the raw data for you to create your own visualization art work. MAG tracks author locations through their affiliations.

MAG can define what AI means for specific case. MAG's current ontology, generated party by machine and defined by humans in the top 3 layers, treats major AI applications such as computer vision, speech recognition and natural language processing as "sibling" fields rather than subfields of AI. It is possible to further develop customized scripts to include or exclude publications/patents in areas that are appropriate for different use cases.

Readers can refer to the "A Century of Science" paper for extracting data on citation and reference between countries.

Dong, Y., Ma, H., Shen, Z., & Wang, K. (2017). A Century of Science: Globalization of Scientific Collaborations, Citations, and Innovations. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1437–1446). https://arxiv.org/pdf/1704.05150.pdf

Return to Research & Development - [Microsoft Academic Graph]

Volumetrics: Total Count of AI Papers (Journal, Conference, Patents) by region, 1990-2019

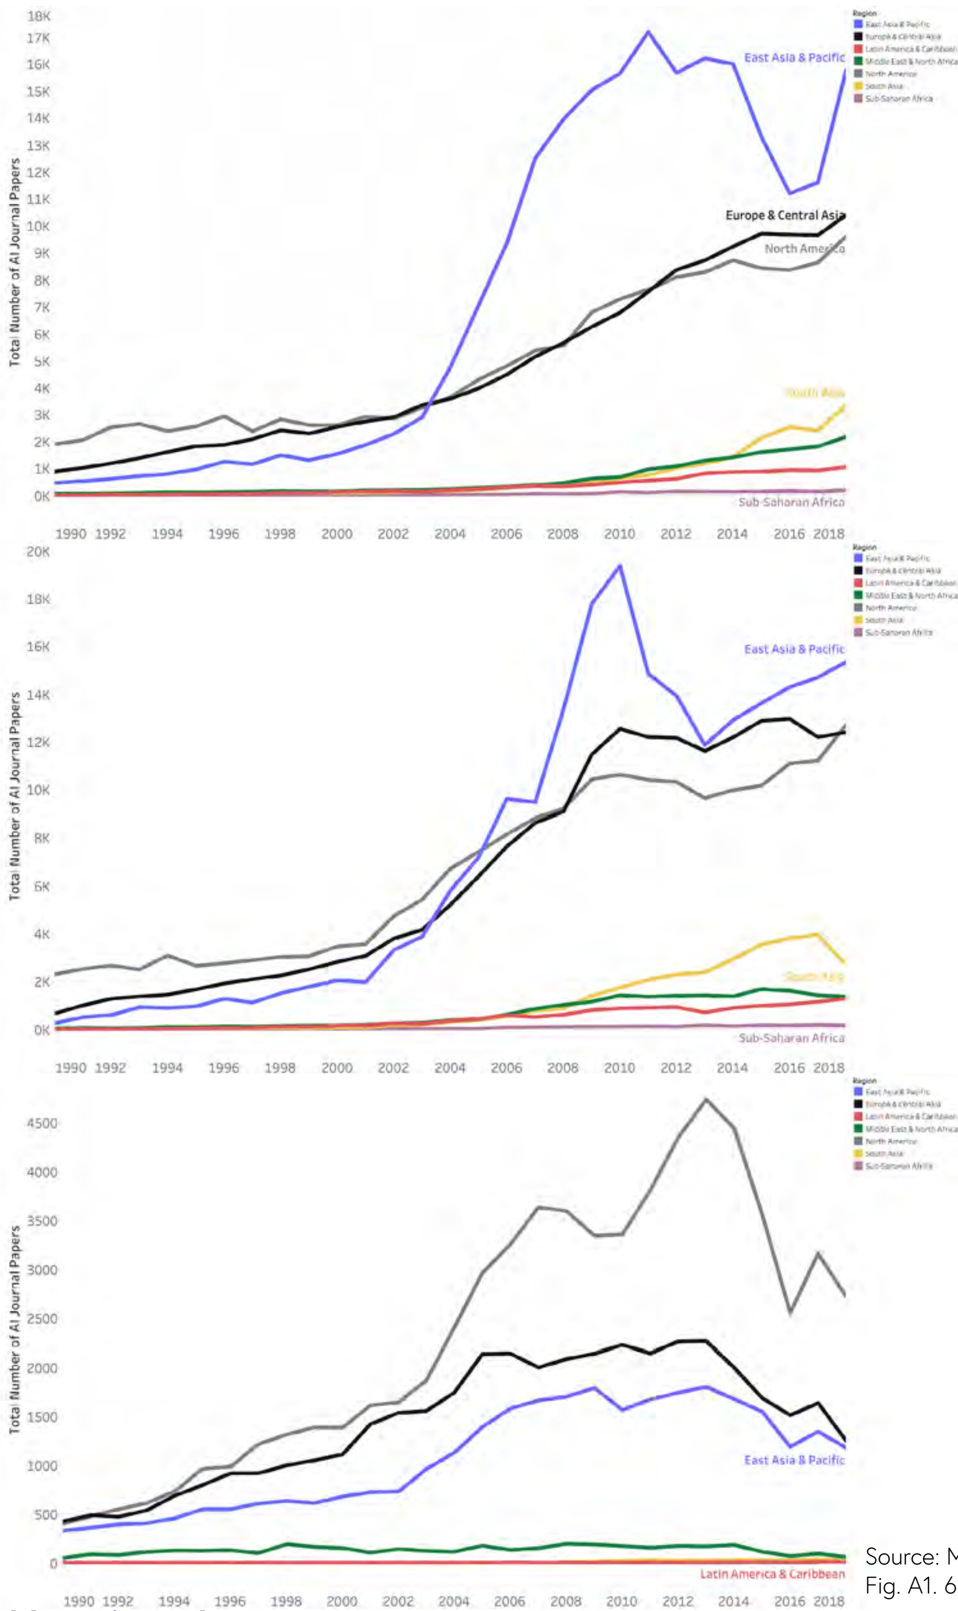

Source: MAG, 2019.
Fig. A1. 6a, b, & c.

Return to Research & Development - [Microsoft Academic Graph]

Total Publications against Per capita Publication, 2015-18

Journals
Source: MAG 2019

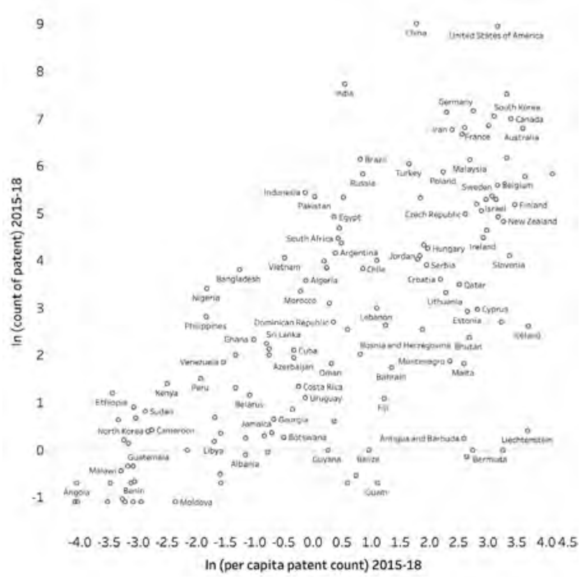

Conferences
Source: MAG 2019

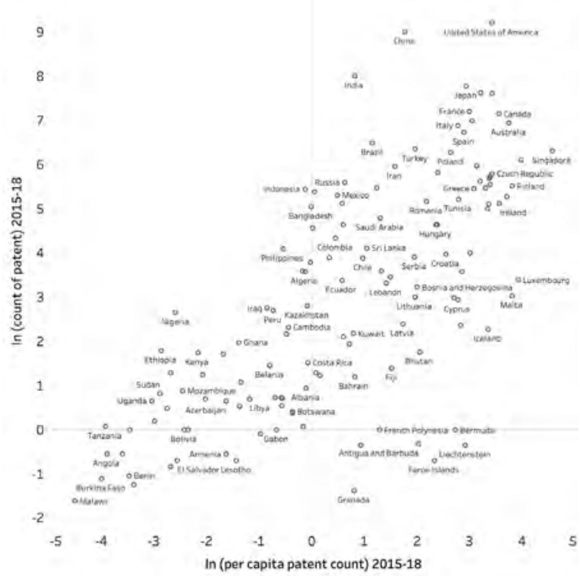

Patents
Source: MAG 2019

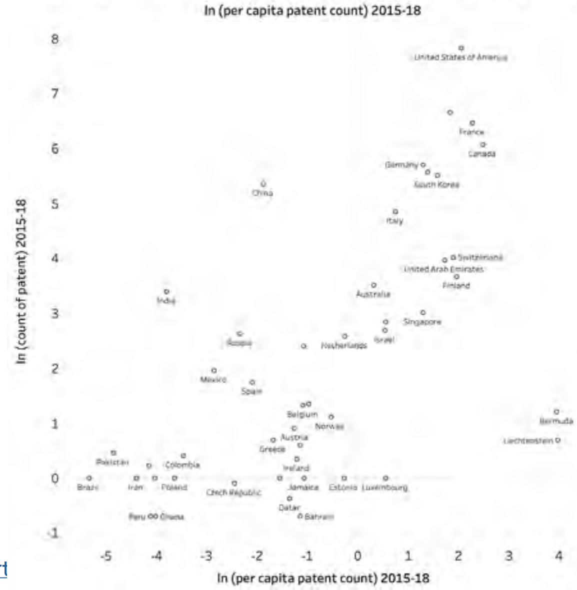

Fig. A1. 7a, b, & c.

Return to Research & Development - [Microsoft Academic Graph (Patents)]

Regional Trends in AI Patents

Global Trends in AI Patents, 1990-2018

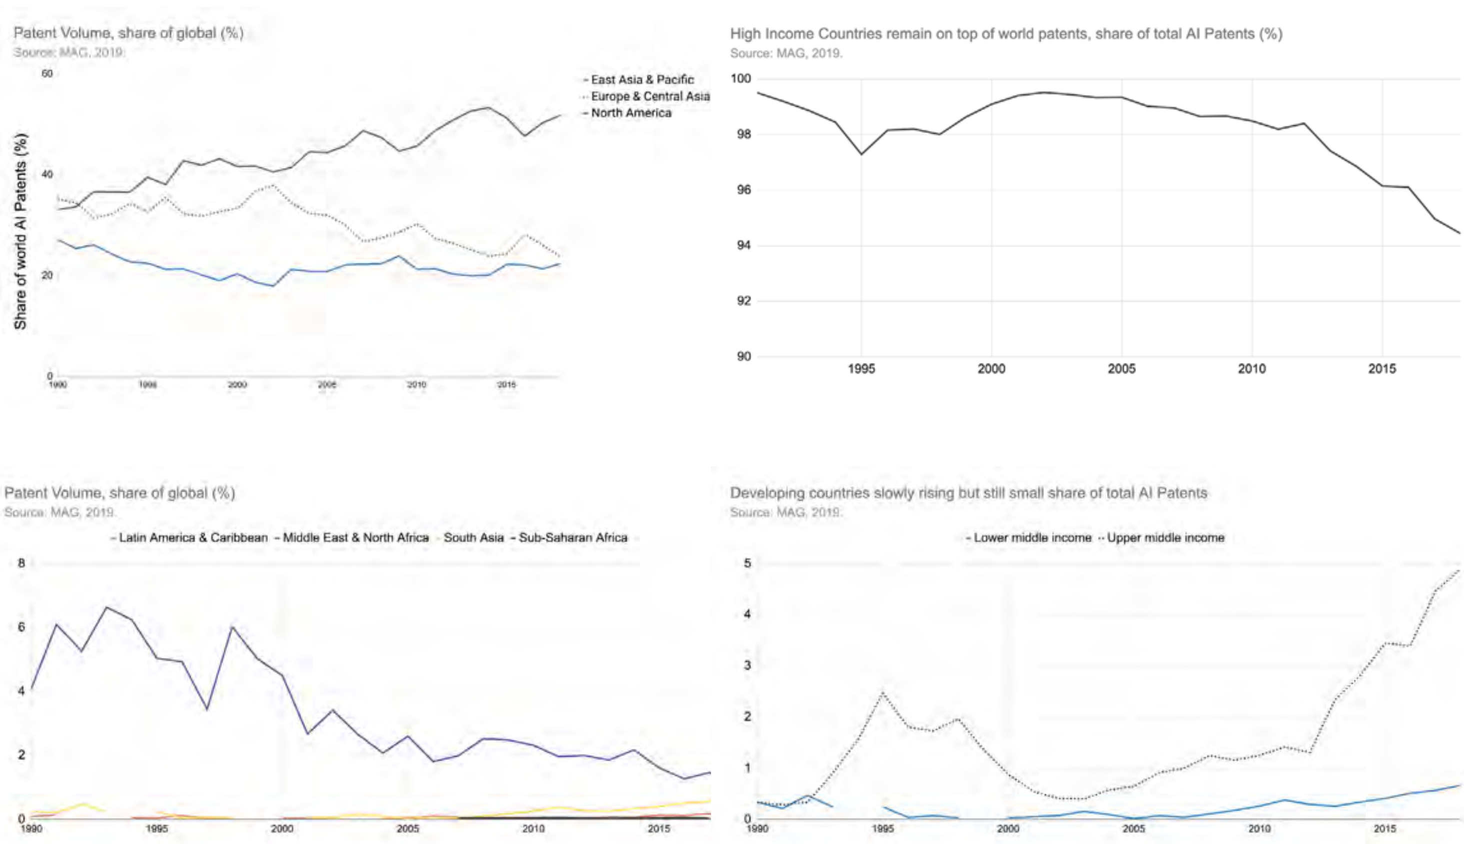

Fig. A1. 8a, b, c, & d.

Published Papers: AI Conference Sectors

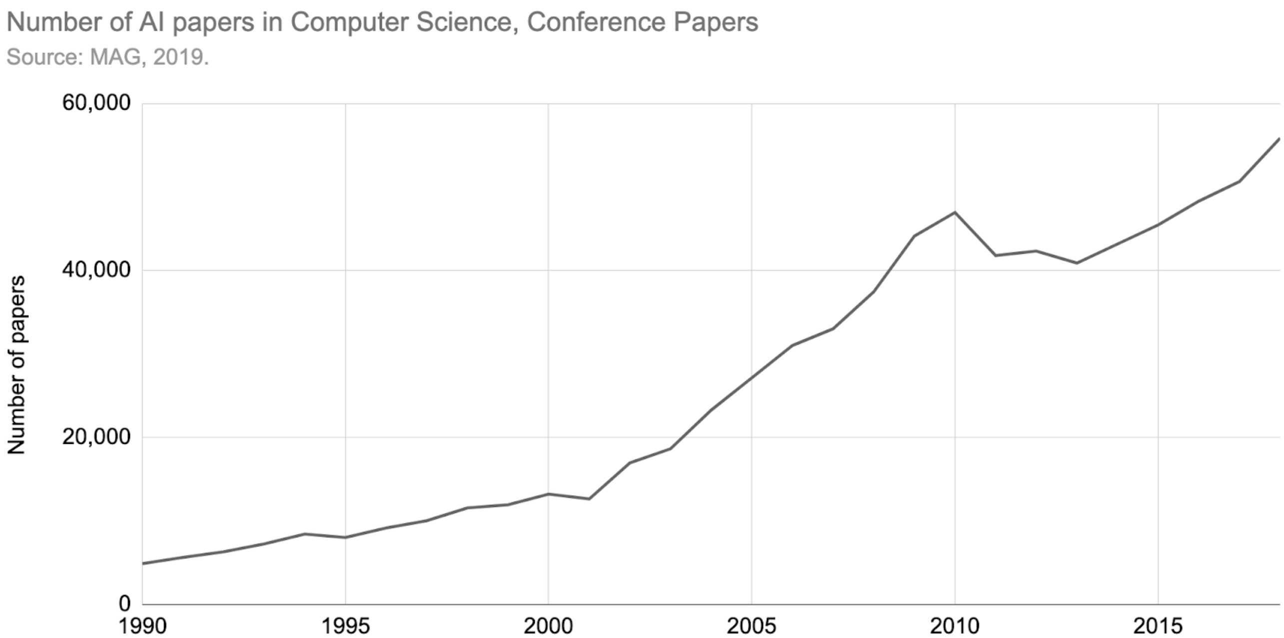

Fig. A1. 9a.

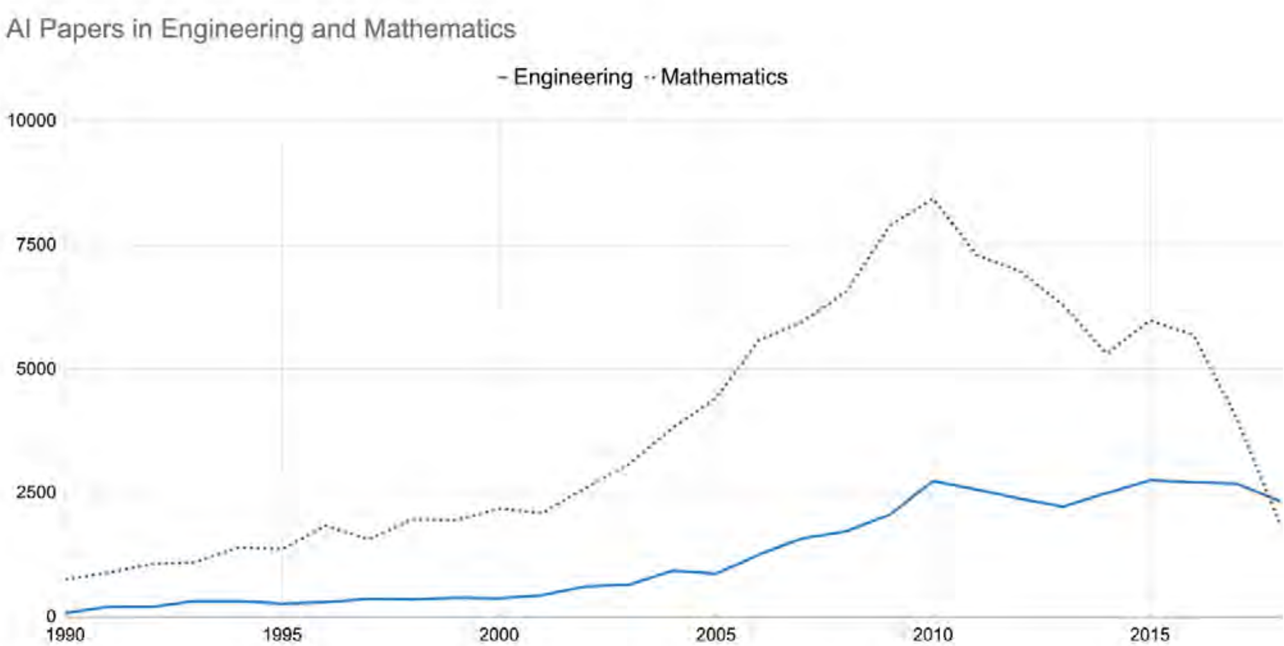

Fig. A1. 9b.

Return to Research & Development - arXiv Deep Learning

NESTA Data and Methodology

Source

Details can be found in the following publication:
Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology.

Methodology

Deep Learning papers were identified through a topic modelling analysis of the abstracts of arXiv papers in the Computer Science and Statistics: Machine Learning category.

Access Data

GitHub repo with the code and data for the regional / national analysis using arXiv data.
https://github.com/Juan-Mateos/ai_index_data

The output_data contains tables with DL paper counts and revealed comparative advantage indices by year and split by pre-post 2012. Analysts can change the parameters at the top to generate similar tables by country or modifying the citation thresholds and watershed years.

Total number of Deep Learning Papers and Per capita measures for all countries, 2018

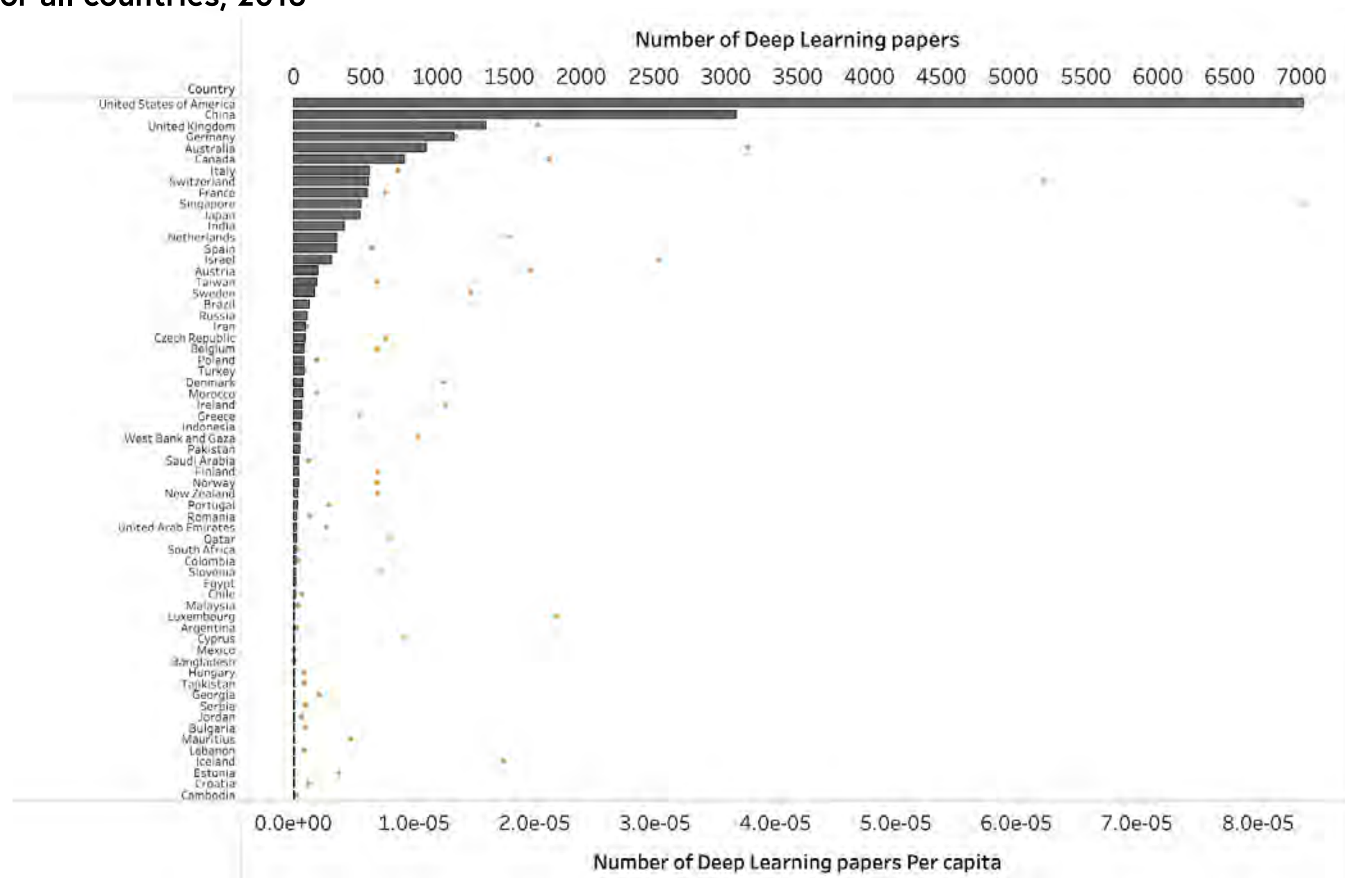

Fig. A1. 10.

Return to Research & Development - [arXiv Deep Learning]

Scatter Plot of Total Number of Deep Learning Papers and Per capita Deep Learning papers on arXiv, 2015-18

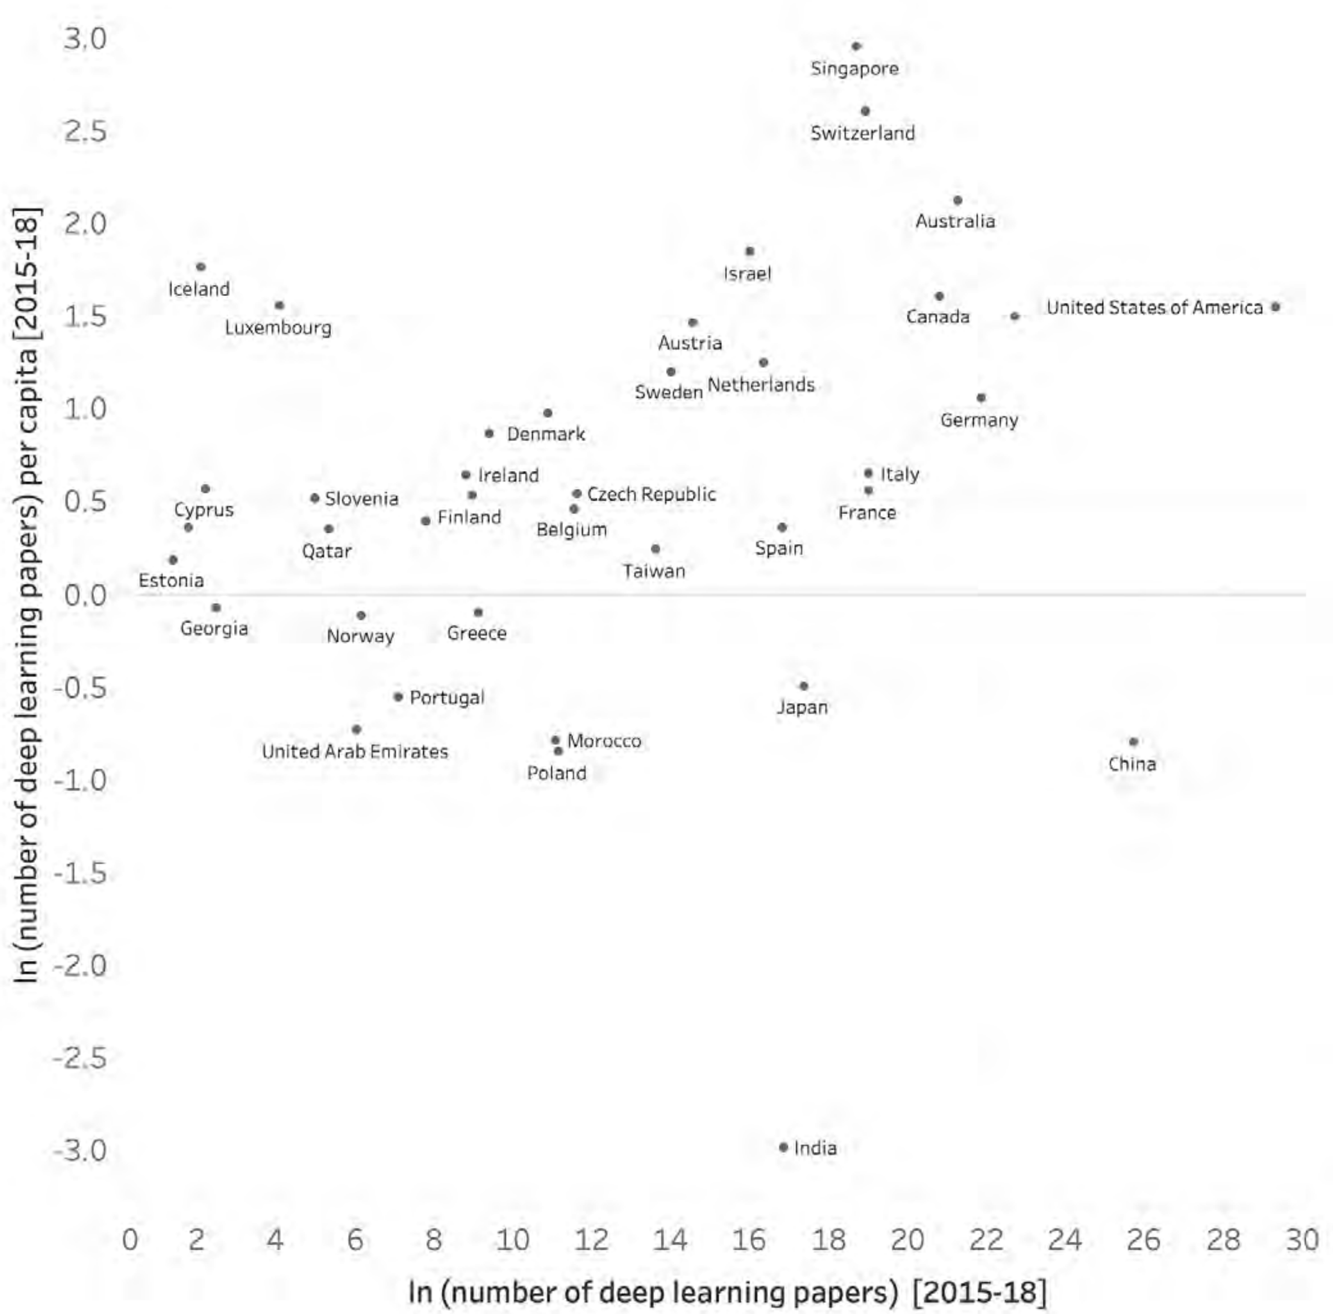

Fig. A1. 11.

Github Stars

Source

We used the GitHub archive stored on Google BigQuery.

Methodology

The visual in the report shows the number of stars for various GitHub repositories over time. The repositories include:

apache/incubator-mxnet, BVLC/cafe, cafe2/cafe2, dmlc/mxnet, fchollet/keras, Microsoft/CNTK, pytorch/pytorch, scikit-learn/scikit-learn, tensorflow/tensorflow, Theano/Theano, Torch/Torch7

GitHub archive data is stored on Google BigQuery. We interfaced with Google BigQuery to count the number of "WatchEvents" for each repository of interest. A sample of code for collecting the data over the course of 2016 is to the right:

```
SELECT
    project,
    YEAR(star_date) as yearly,
    MONTH(star_date) as monthly,
    SUM(daily_stars) as monthly_stars
FROM (
SELECT
    repo.name as project,
    DATE(created_at) as star_date,
    COUNT(*) as daily_stars
FROM
    TABLE_DATE_RANGE(
        [githubarchive:day.],
        TIMESTAMP("20160101"),
        TIMESTAMP("20161231"))
WHERE
    repo.name IN (
        "tensorflow/tensorflow",
         "fchollet/keras",
        "apache/incubator-mxnet",
         "scikit-learn/scikit-learn",
        "cafe2/cafe2", "pytorch/pytorch",
         "Microsoft/CNTK", "Theano/Theano",
        "dmlc/mxnet", "BVLC/cafe")
    AND type = 'WatchEvent'
GROUP BY project, star_date
)
GROUP BY project, yearly, monthly
ORDER BY project, yearly, monthly
```

Nuance

The GitHub Archive currently does not provide a way to count when users remove a Star from a repository. Therefore, the data reported slightly overestimates the count of Stars. Comparison with the actual number of Stars for the repositories on GitHub shows that the numbers are fairly close and the trends remain unchanged.

There are other ways to retrieve GitHub Star data. The star-history tool was used to spot-check our results.

While Forks of GitHub project are also interesting to investigate, we found that the trends of repository Stars and Forks were almost identical.

Women in AI

Source

The data is based on NESTA paper titled Gender Diversity in AI Research.

Methodology

The analysis relies on several data collection and processing steps that are described below and can be inspected on GitHub. All papers are extracted from arXiv using the API, yielding 1,372,350 papers (after cleaning) which we used in the analysis. Based on strategy described by Klinger, et al. (2018), information from the arXiv corpus was matched with MAG .87 per cent of the arXiv preprints were matched with MAG. Authors' geolocation was determined by looking up their institution in Google Places API, a commercial cloud service that provides names, addresses, and other information for over 150 million places. 93 per cent of the 8,351 affiliations were successfully geocoded.

Gender API, the biggest platform on the internet to determine gender by a first name, a full name or an email address was used for inferring gender from names. This database contains 1,877,874 validated names from 178 different countries. The inference of the gender from author names in corpus follows this approach:

- Query the Gender API with full names. The last name is used to improve results on gender-neutral names.
- Exclude results where the first name field contained only an initial
- Remove results with less than 80 per cent accuracy
- Remove any papers where gender cannot be determined for more than 50 per cent of the authors

Following this procedure, about 480K of the roughly 772K author names in arXiv were labelled. As with all other inference systems, Gender API has limitations. It may underestimate the number of female names and its performance degrades with Asian and especially South-East Asian names. Moreover, it assumes that gender identity is both a fixed and binary concept. We acknowledge that this limitation restricts the scope of our analysis to binary genders, and will account for identities beyond binary in future analyses.

The approach was implemented in the following way: first, text was lowercased and tokenised, stop words, punctuation and numeric characters were removed from all of the abstracts. Then bigrams and trigrams were created. Then, two models were applied to the data:

- Word2Vec with the Continuous Bag-of-Words (CBOW) architecture
- Term frequency, Inverse document frequency (TF-IDF)

Lastly, the pretrained word2vec is queried for AI related terms to extract a list of similar tokens, the most common and rare ones are filtered using their inverse document frequency (IDF) and the paper abstracts are searched search for the rest. More details can be found in the Methodology Paper.

Proportion of Female AI authors, Netherlands,US, and Japan, 1995-2018

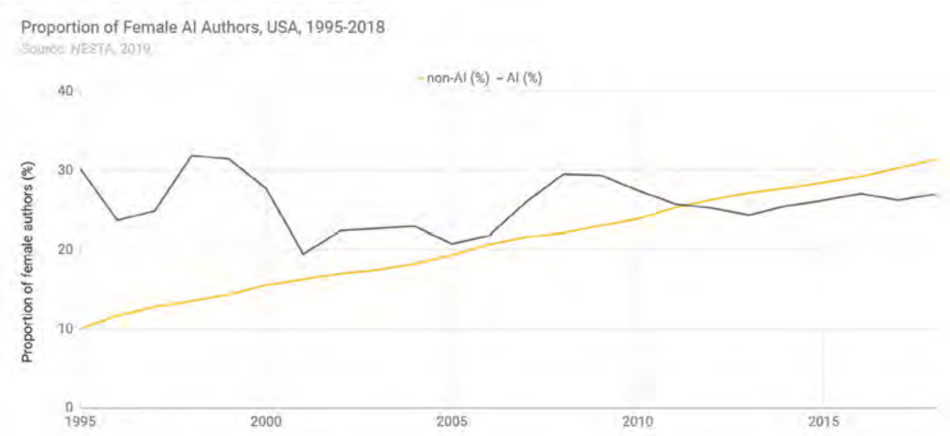

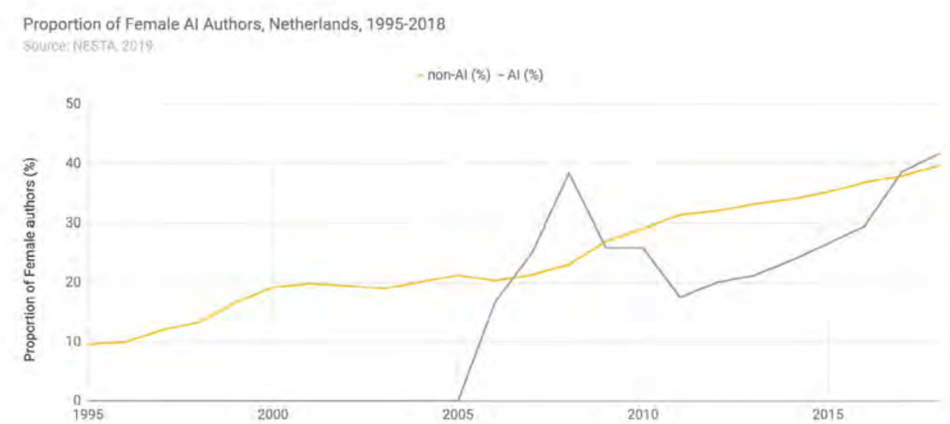

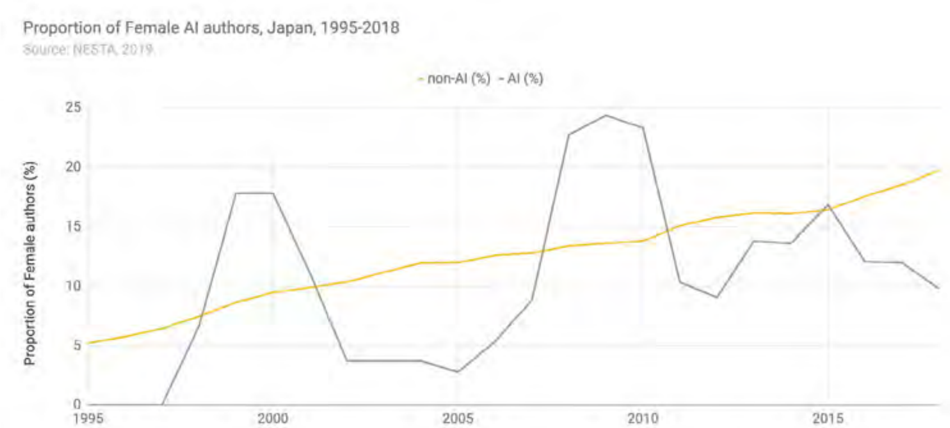

Fig. A1. 12a, b, & c.

Conference Participation

Source

Conference attendance data was collected directly from conference / organization representatives. Data was collected from the following conferences:

AAAI — Association for the Advancement of Artificial Intelligence
AAMAS — International Conference on Autonomous Agents and Multiagent Systems
AI4ALL
ACL — Association for Computational Linguistics
CVPR — Conference on Computer Vision and Pattern Recognition
ICAPS — International Conference on Automated Planning and Scheduling
ICLR — International Conference on Learning Representations
ICML — International Conference on Machine Learning
ICRA — International Conference on Robotics and Automation
IJCAI — International Joint Conferences on Artificial Intelligence
KR — International Conference on Principles of Knowledge Representation and Reasoning NeurIPS — Conference on Neural Information Processing Systems
UAI — Conference on Uncertainty in Artificial Intelligence
WiML — Women in Machine Learning workshop

Methodology

We defined large conferences as those with 2,000 or more attendees in 2018, and small conferences as those with fewer than 2,000 attendees in 2018. Conferences selected are those that lead in AI research and were also able to supply yearly attendance data.

AI4ALL and WiML were selected for their progress on AI inclusion and their availability of data. We look forward to adding more organizations / conferences that cater to underrepresented groups in future reports.

AI4ALL Open Learning was launched with 8 educational partners across the US who are using the curriculum in their classrooms and clubs, including Girl Scouts of Northeast Texas, National Society of Black Engineers Bay Area, and the Stockton Unified School District, among others. The program is slated to reach over 750 high school students through our education partners and other students using the platform by the end of 2019.

Nuance

Nuances specific to conferences

- Some conference organizers were only able to provide estimates of attendance — we have accepted estimates as accurate.
- Some conferences do not run annually, and some have skipped years.
- Several conference organizers have let us know that because conference venues are determined over a year in advance, the supply of spots are often limited. Therefore, the number of conference attendees doesn't necessarily reflect demand.

Nuances specific to AI4ALL / WiML

- It is important to note that several other formal and informal programs exist to support inclusivity in AI.
- Participation does not necessarily indicate progress in increasing the number of women and underrepresented groups in the field.

AAAI papers Statistics

Source

The Association for the Advancement of Artificial Intelligence (AAAI) hosts conferences every year, including the annual "AAAI conference". Raw data on 2019 AAAI paper submissions / acceptances by country was provided by AAAI representatives. Learn more about the AAAI conferences.

Methodology

We collected data on AAAI submissions / acceptance by country from the AAAI team. AAAI was the only conference where we were able to obtain this level of detail. The AI Index hopes to include equivalent data for other conferences in future reports.

Nuance

- Countries included in this analysis are those that submitted 10 or more papers to the AAAI conference.
- This data is from the 2019 conference. The landscape of submitted / accepted papers may look different for other years.
- Acceptance is largely limited due to space constraints.

Ethics at AI Conference

Source

Prates, Marcel, Pedro Avelar, Luis C. Lamb. 2018. On Quantifying and Understanding the Role of Ethics in AI Research: A Historical Account of Flagship Conferences and Journals. 21 Sep 2018.

Methodology

The percent of keywords has a straightforward interpretation: for each category (classical / trending / ethics) the number of papers for which the title (or abstract, in the case of the AAAI and NIPS figures) contains at least one keyword match. The percentages do not necessarily add up to 100% (i.e. classical / trending / ethics are not mutually exclusive). One can have a paper with matches on all three categories.

To achieve a measure of how much Ethics in AI is discussed, ethics-related terms are searched for in the titles of papers in flagship AI, machine learning and robotics conferences and journals.

The **ethics keywords** used were the following:
Accountability, Accountable, Employment, Ethic, Ethical, Ethics, Fool, Fooled, Fooling, Humane, Humanity, Law, Machine bias, Moral, Morality, Privacy, Racism, Racist, Responsibility, Rights, Secure, Security, Sentience, Sentient, Society, Sustainability, Unemployment and Workforce.

The classical and trending keyword sets were compiled from the areas in the most cited book on AI by Russell and Norvig [2012] and from curating terms from the keywords that appeared most frequently in paper titles over time in the venues.

The keywords chosen for the **classical keywords** category were:
Cognition, Cognitive, Constraint satisfaction, Game theoretic, Game theory, Heuristic search, Knowledge representation, Learning, Logic, Logical, Multiagent, Natural language, Optimization, Perception, Planning, Problem solving, Reasoning, Robot, Robotics, Robots, Scheduling, Uncertainty and Vision.

The curated **trending keywords** were:
Autonomous, Boltzmann machine, Convolutional networks, Deep learning, Deep networks, Long short term memory, Machine learning, Mapping, Navigation, Neural, Neural network, Reinforcement learning, Representation learning, Robotics, Self driving, Self-driving, Sensing, Slam, Supervised/Unsupervised learning and Unmanned.

The terms searched for were based on the issues exposed and identified in papers below, and also on the topics called for discussion in the First AAAI/ACM Conference on AI, Ethics, and Society.

J. Bossmann. Top 9 ethical issues in artificial intelligence. 2016. World Economic Forum - https://www.weforum.org/agenda/2016/10/top-10-ethical-issues-in-artificialintelligence/[Online; 21-Oct-2016].

Emanuelle Burton, Judy Goldsmith, Sven Koenig, Benjamin Kuipers, Nicholas Mattei, and Toby Walsh. Ethical considerations in artificial intelligence courses. AI Magazine, 38(2):22–34, 2017.

The Royal Society Working Group, P. Donnelly, R. Browsword, Z. Gharamani, N. Griffiths, D. Hassabis, S. Hauert, H. Hauser, N. Jennings, N. Lawrence, S. Olhede, M. du Sautoy, Y.W. Teh, J. Thornton, C. Craig, N. McCarthy, J. Montgomery, T. Hughes, F. Fourniol, S. Odell, W. Kay, T. McBride, N. Green, B. Gordon, A. Berditchevskaia, A. Dearman, C. Dyer, F. McLaughlin, M. Lynch, G. Richardson, C. Williams, and T. Simpson. Machine learning: the power and promise of computers that learn by example. The Royal Society, 2017.

Conference and Public Venue - Sample

The AI group contains papers from the main Artificial Intelligence and Machine Learning conferences such as AAAI, IJCAI, ICML, NIPS and also from both the Artificial Intelligence Journal and the Journal of Artificial Intelligence Research (JAIR).

The Robotics group contain papers published in the IEEE Transactions on Robotics and Automation (now known as IEEE Transactions on Robotics), ICRA and IROS.

The CS group contains papers published in the mainstream Computer Science venues such as the Communications of the ACM, IEEE Computer, ACM Computing Surveys and the ACM and IEEE Transactions.

Codebase

The code and data are hosted in this Github repository https://github.com/marceloprates/Ethics-AI-Data

The "correlation-matrix" analysis refers to titles only. It measures the correlation between the number of papers matching for ethics keywords and the number of papers matching for trending keywords (for example). Although the correlation coefficients are close to zero, both classical and trending matches are negatively correlated with ethics. This could suggest both that traditional (classical and trending) papers in leading conferences fail to mention ethics and that ethics papers are perhaps too immersed in their own subjects to mention hot topics in other areas.

Correlation Matrix for Classical, Trending, and Ethics keywords

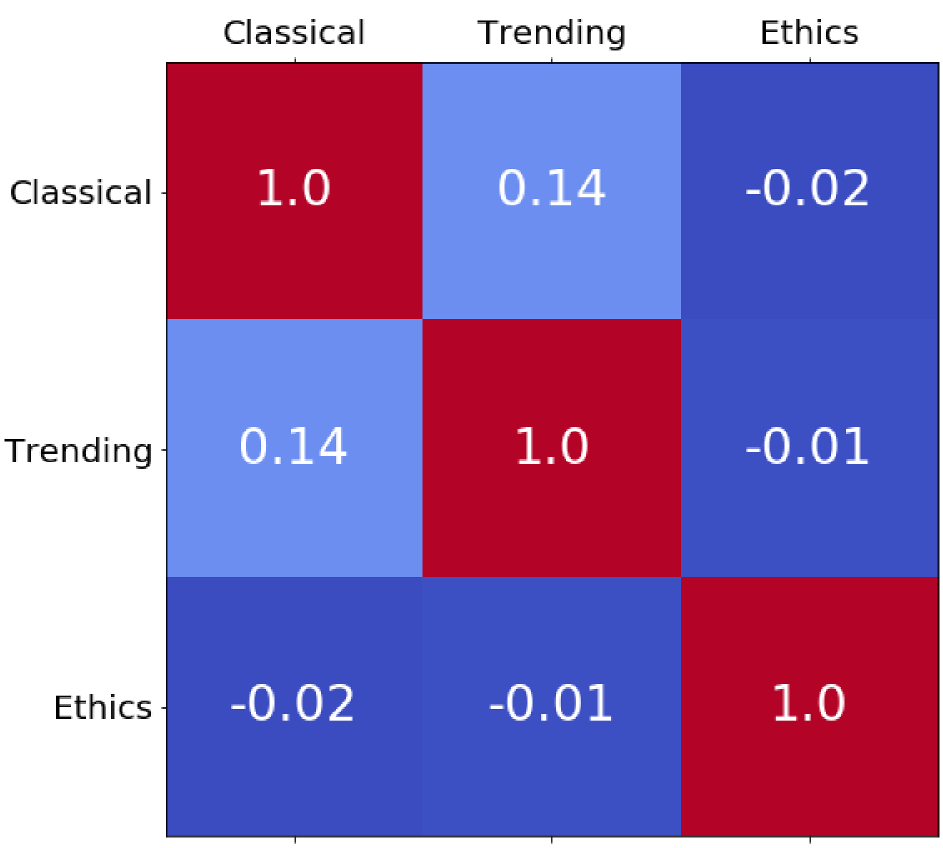

Fig. A2. 1.

Return to Technical Performance - Computer Vision: ImageNet

ImageNet

Source

Data on ImageNet accuracy was retrieved through an arXiv literature review. All results reported were tested on the LSRVC 2012 validation set - their ordering may differ from the results reported on the LSRVC website, since those results were obtained on the test set. Dates we report correspond to the day when a paper was first published to arXiv, and top-1 accuracy corresponds to the result reported in the most recent version of each paper. We selected a top result at any given time point from 2012 to November 17, 2019. Some of the results we mention were submitted to LSRVC competitions over the years. Image classification was part of LSRVC through 2014; in 2015 it was replaced with an object localization task, where results for classification were still reported, but were no longer a part of the competition, and were instead replaced with more difficult tasks.

For papers published in 2014 and later, we report the best result obtained using a single model (we did not include ensembles) and using single-crop testing. For the three earliest models (AlexNet, ZFNet, Five Base) we reported the results for ensembles of models.

While we report the results as described above, due to the diversity in models, evaluation methods and accuracy metrics, there are many other ways to report ImageNet performance. We list the possible choices below:

- Evaluation set: validation set (available publicly) or test set (available only LSRVC organizers)
- Performance Metric: Top-1 Accuracy (whether the correct label was the same as the first predicted label for each image) or Top-5 Accuracy (whether the correct label was present among the top 5 predicted labels for each image)
- Evaluation method: single-crop or multi-crop

Return to Technical Performance - Computer Vision: ImageNet

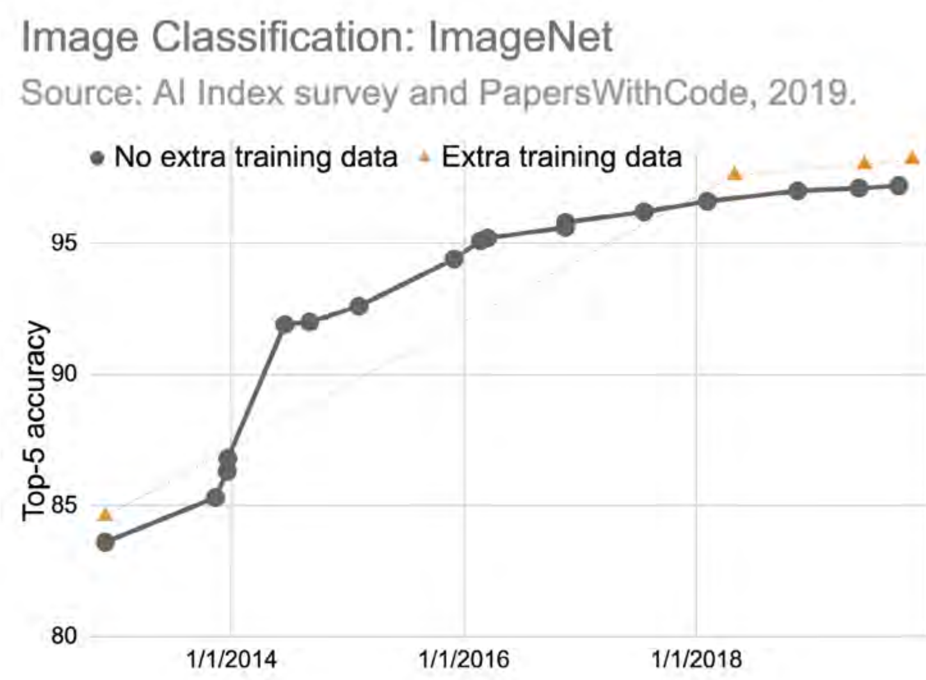

Fig. A3.1. Top-5 Accuracy on ImageNet

To highlight progress here, we have taken scores from the following papers:

- ImageNet Classification with Deep Convolutional Neural Networks
- Visualizing and Understanding Convolutional Networks
- Some Improvements on Deep Convolutional Neural Network Based Image Classification
- OverFeat: Integrated Recognition, Localization and Detection using Convolutional Networks
- Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition
- Very Deep Convolutional Networks for Large-Scale Image Recognition
- Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification
- Rethinking the Inception Architecture for Computer Vision
- Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning
- Identity Mappings in Deep Residual Networks
- Aggregated Residual Transformations for Deep Neural Networks
- PolyNet: A Pursuit of Structural Diversity in Very Deep Networks
- Learning Transferable Architectures for Scalable Image Recognition
- Regularized Evolution for Image Classifier Architecture Search
- GPipe: Efficient Training of Giant Neural Networks using Pipeline Parallelism
- EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks
- RandAugment: Practical data augmentation with no separate search
- Self-training with Noisy Student improves ImageNet classification
- Fixing the train-test resolution discrepancy
- Exploring the Limits of Weakly Supervised Pretraining
- Revisiting Unreasonable Effectiveness of Data in Deep Learning Era

The estimate of human-level performance is from Russakovsky et al, 2015. Learn more about the LSVRC ImageNet competition and the ImageNet data set.

Return to Technical Performance - Computer Vision: [ImageNet]

Training Time on Private Infrastructure

Trends can also be observed by studying research papers that discuss the time it takes to train ImageNet on any infrastructure. This gives us a sense of the difference between public cloud and private cloud infrastructure, and also provides another view of progress in this domain. To gather this data, research papers published over the last few years were analyzed, which seek to train systems to competitive wallclock times while achieving top-1 accuracy on ImageNet. This maps to the contemporary state of the art. The reference to papers and details on hardware can be found here.

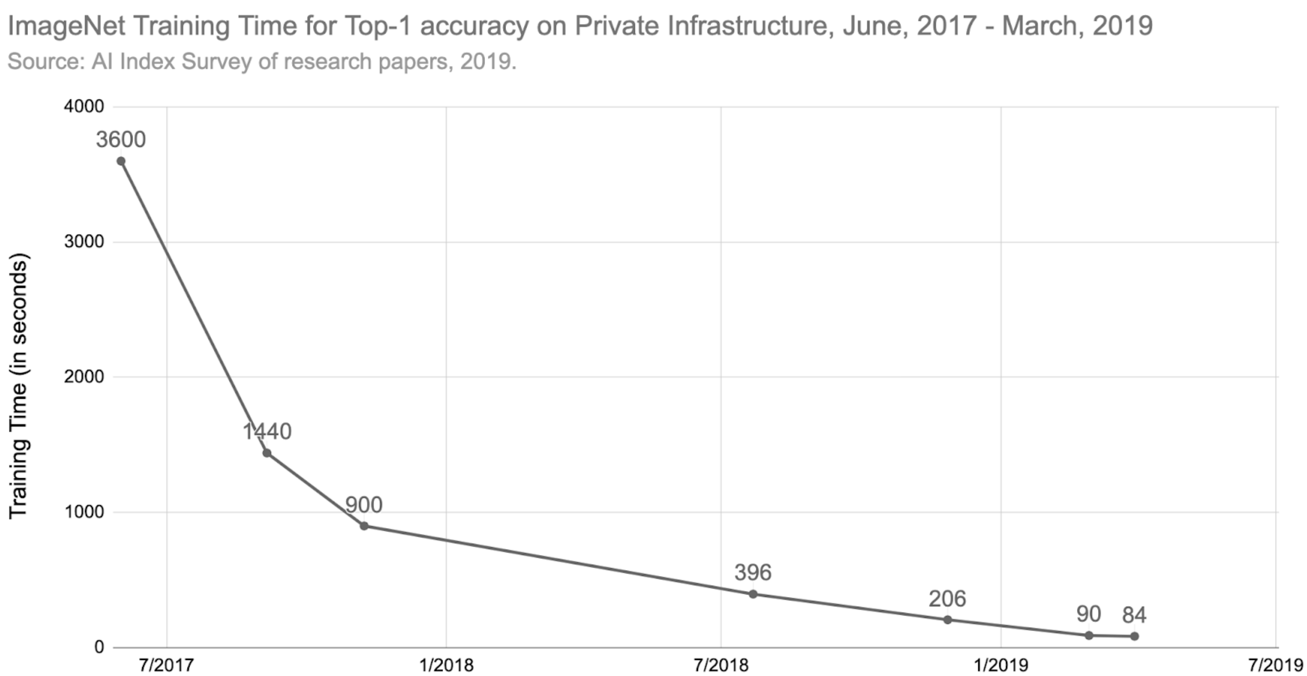

Fig. A3.2.

DAWN Benchmark

Source

DAWNBench is a benchmark suite for end-to-end deep learning training and inference. Computation time and cost are critical resources in building deep models, yet many existing benchmarks focus solely on model accuracy. DAWNBench provides a reference set of common deep learning workloads for quantifying training time, training cost, inference latency, and inference cost across different optimization strategies, model architectures, software frameworks, clouds, and hardware.

More details available:
https://dawn.cs.stanford.edu

Methodology and Definition

The following metrics are introduced to compute training time and cost.

ImageNet Compute Economic Metrics

Metric	Definition	Units
Training Time	Time taken to train an image classification model to a top-5 validation accuracy of 93% or greater on ImageNet.	Time to 93% Accuracy
Training Cost	Total cost of public cloud instances to train an image classification model to a top-5 validation accuracy of 93% or greater on ImageNet.	Cost USD
Inference Latency	Latency required to classify one ImageNet image using a model with a top-5 validation accuracy of 93% or greater.	1-example Latency (milliseconds)
Inference Cost	Average cost on public cloud instances to classify 10,000 validation images from ImageNet using an image classification model with a top-5 validation accuracy of 93% or greater.	Cost USD

Return to Technical Performance - [ImageNet,](#) [ImageNet Training,](#) [SQuAD](#)

ImageNet Inference Latency and Inference Cost

[Access Data](#)

The inference latency i.e. the 1-example latency in milliseconds (to classify one ImageNet image using a model with a top-5 validation accuracy of 93% or greater) the inference cost i.e. the USD cost (on public cloud instances to classify 10,000 validation images from ImageNet using an image classification model with a top-5 validation accuracy of 93% or greater) results are presented. The inference time has reduced from 22ms in November, 2018 to 0.82 ms in July, 2019. The inference cost has become almost zilch.

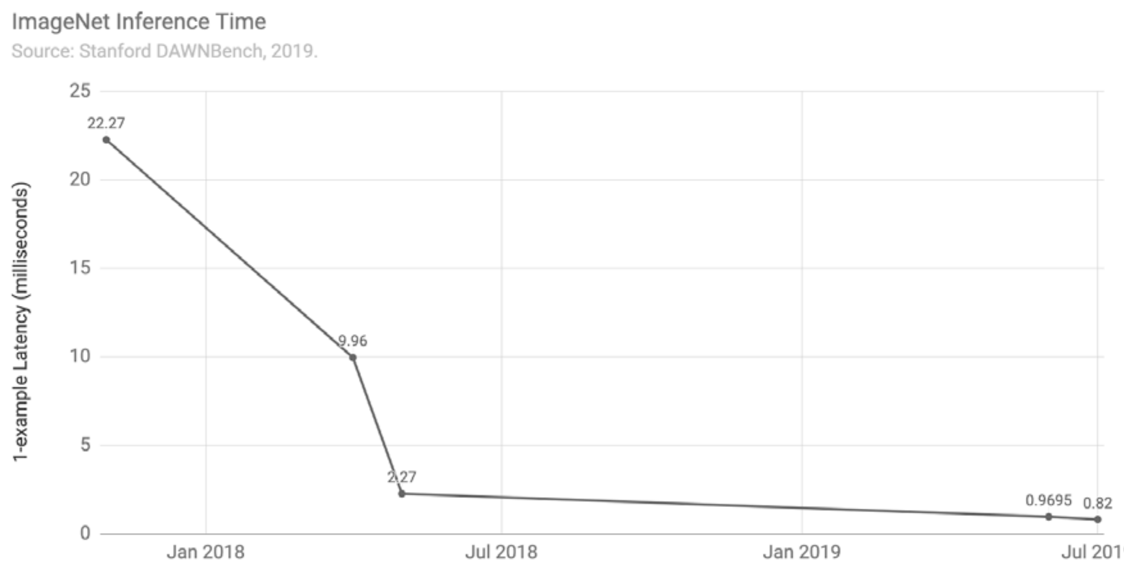

Fig. A3.3a.

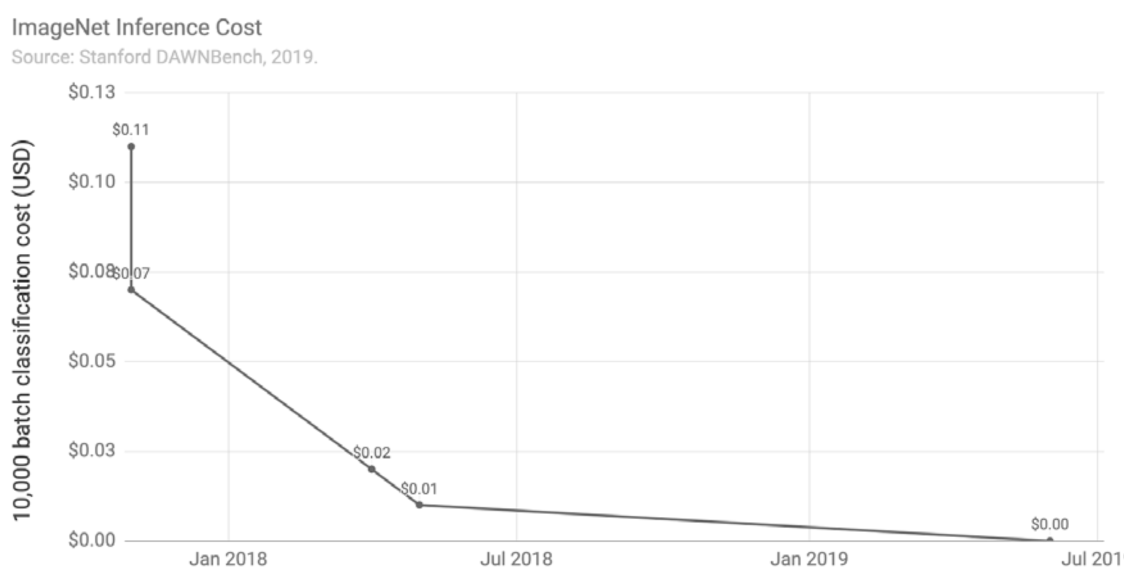

Fig. A3.3b.

Image Classification

Return to Technical Performance - Computer Vision: Image Classification

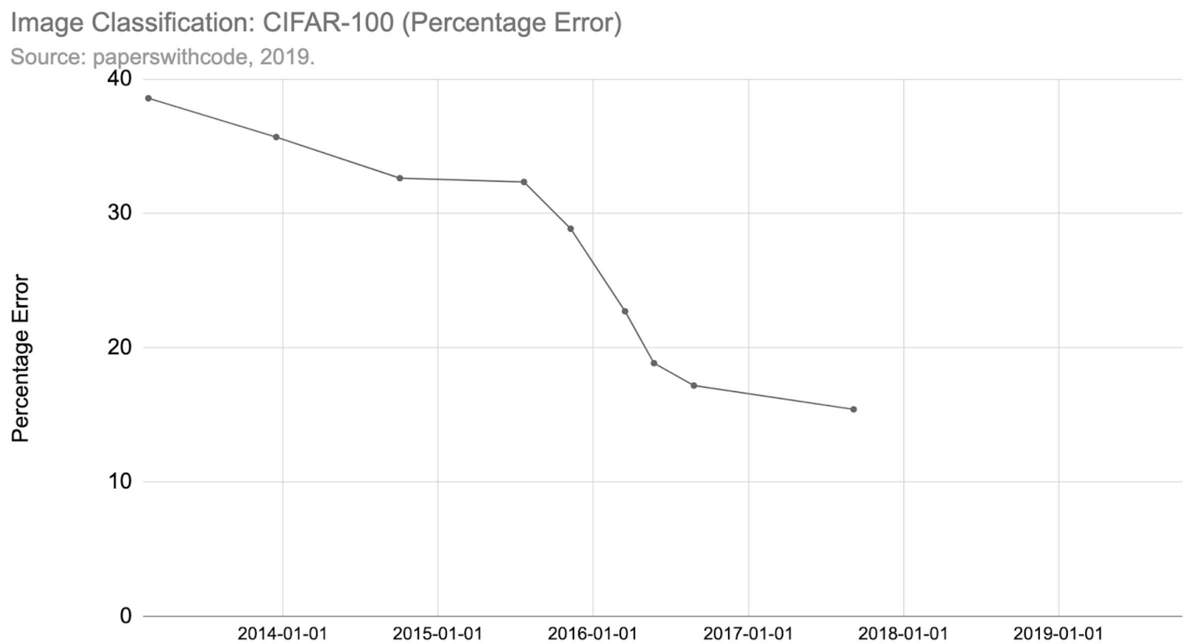

Fig. A3.4.

Image Generation

Return to Technical Performance - Computer Vision: Image Generation

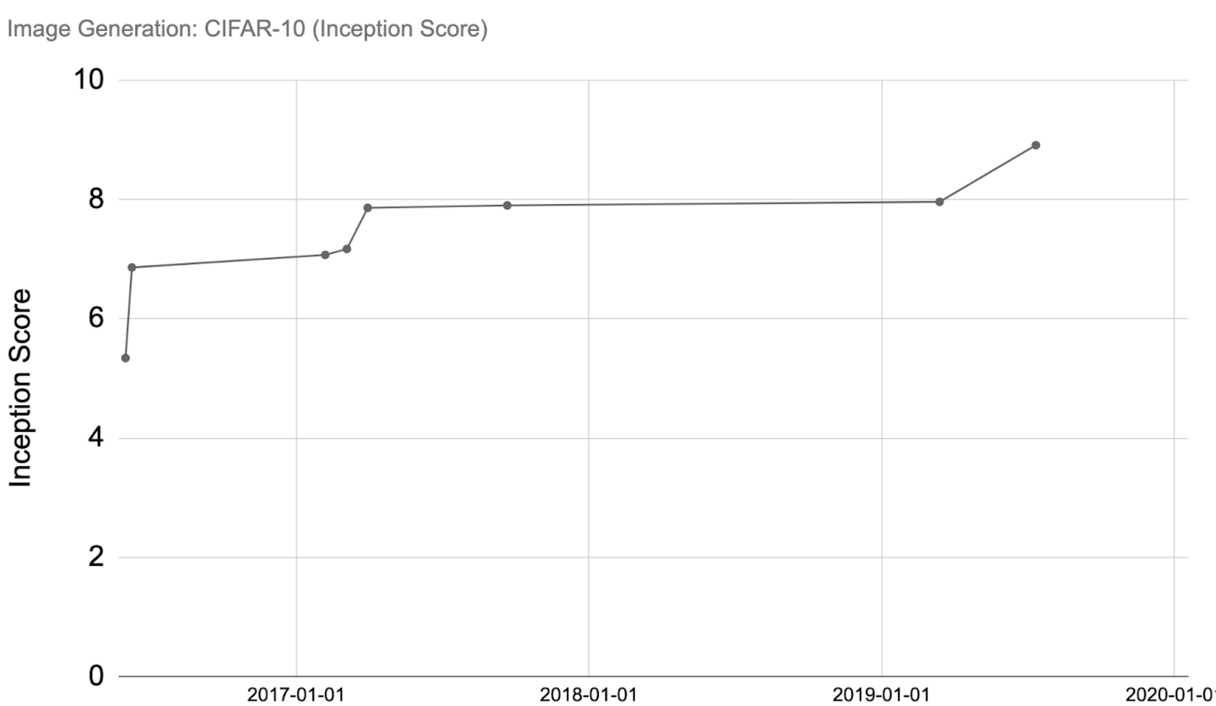

Fig. A3.5a.

Semantic Segmentation

Datasets and Challenges

The Cityscapes dataset contains a diverse set of stereo video sequences recorded in street scenes from 50 different cities, with high quality pixel-level annotations of 5 000 frames in addition to a larger set of 20,000 weakly annotated frames.

PASCAL Context dataset additional annotations for PASCAL VOC 2010. It goes beyond the original PASCAL semantic segmentation task by providing annotations for the whole scene. The statistics section has a full list of 400+ labels.

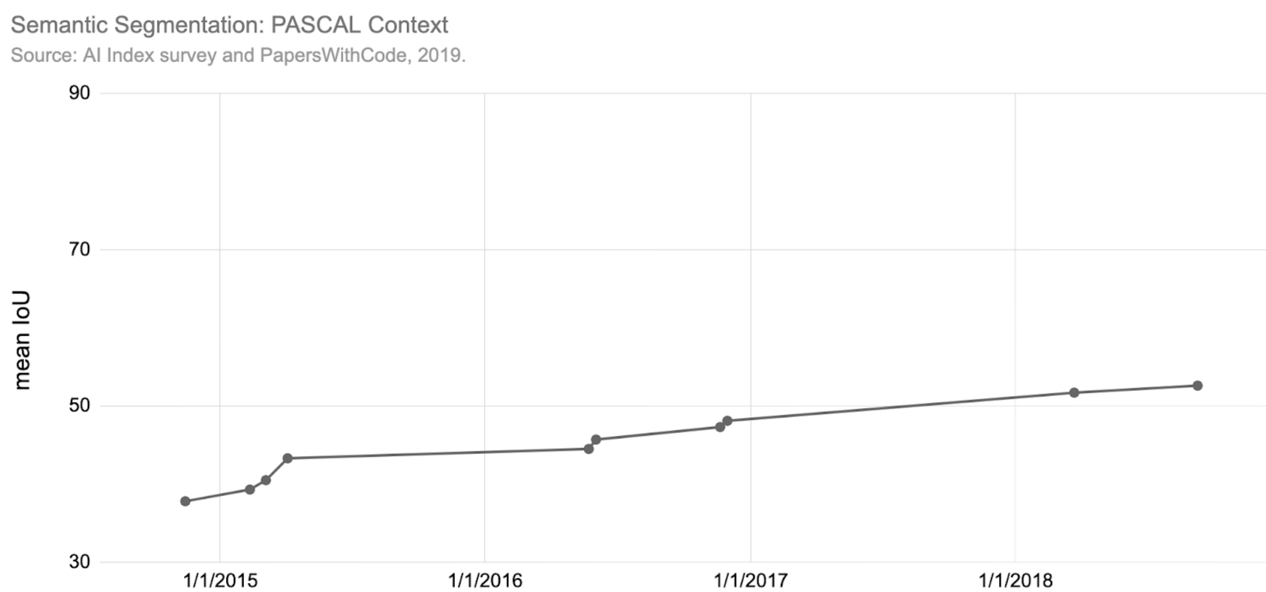

Fig. A3.5b.

Visual Question Answering (VQA)

Source

VQA accuracy data was provided by the VQA team. Learn more about VQA here. More details on VQA 2019 are available here.

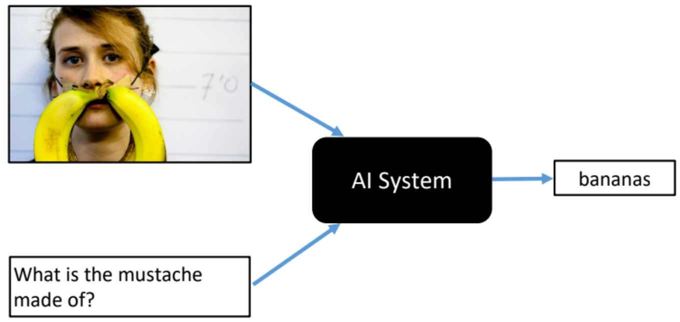

Fig. A3.6.

Methodology

Given an image and a natural language question about the image, the task is to provide an accurate natural language answer. The challenge is hosted on EvalAI. Challenge link: https://evalai.cloudcv.org/web/challenges/challenge-page/163/overview

The VQA v2.0 train, validation and test sets, containing more than 250K images and 1.1M questions, are available on the download page. All questions are annotated with 10 concise, open-ended answers each. Annotations on the training and validation sets are publicly available.

VQA Challenge 2019 is the fourth edition of the VQA Challenge. Previous three versions of the VQA Challenge were organized in past three years, and the results were announced at VQA Challenge Workshop in CVPR 2018, CVPR 2017 and CVPR 2016. More details about past challenges can be found here: VQA Challenge 2018, VQA Challenge 2017 and VQA Challenge 2016.

VQA had 10 humans answer each question. More details about the VQA evaluation metric and human accuracy can be found here (see Evaluation Code section) and in sections 3 (the subsection on Answers) and 4 (the subsection on Inter-human Agreement) of the paper.

Paper and Code Linking

Source

ImageNet accuracy and model complexity, Semantic Segmentation, Image Generation, and CIFAR-100 data was pulled from paperswithcode. Learn more about here.

Methodology

Paper and code linking. For papers we follow specific ML-related categories on arxiv (see [1] below for the full list) and the major ML conferences (NeurIPS, ICML, ICLR, etc..). For code we follow github repositories mentioning papers.We have a good coverage of core ML topics, but are missing some applications, e.g. applications of ML in medicine or bioinformatics, which are usually in journals behind paywalls . For code the dataset is pretty unbiased (as long as the paper is freely available).

For tasks (e.g. "Image classification"), the dataset has annotated those on 1600 SOTA papers from the database, published in 2018 Q3.

For SOTA tables (i.e. "Image classification on ImageNet") - the data has been scraped from a couple of different sources (full list here: https://github.com/paperswithcode/sota-extractor) and hand-annotated a large number focusing on CV and NLP.

A significant proportion of our data was contributed by users, and they've added data based on their own preferences and interests.

[1] Arxiv categories we follow:

```
ARXIV_CATEGORIES = {
        "cs.CV",
        "cs.AI",
        "cs.LG",
        "cs.CL",
        "cs.NE",
        "stat.ML",
        "cs.IR",
}
```

The public link is the following
https://paperswithcode.com/sota

Process of Extracting Dataset at Scale

1) Follow various paper sources (as described above) for new papers
2) Do a number of pre-defined searches on github (e.g. for READMEs containing links to arxiv)
3) Extract github links from papers
4) Extract paper links from github
5) Run validation tests to decide if links from 3) and 4) are bona-fide links or false positives
6) Let the community fix any errors, and/or add any missing values

Sample of Task Areas represented on paperswithcode

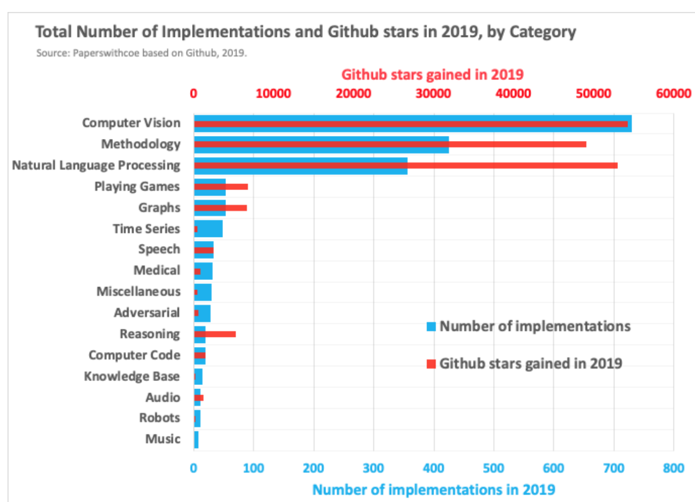

Fig. A3.7.

Note: Number of implementations is the number of _independent_ implementations.

Return to Technical Performance - Language: GLUE

GLUE

Source

GLUE benchmark data was pulled from the GLUE leaderboard.
Learn more about the GLUE benchmark here.

Methodology

Participants download the GLUE tasks and submit to the leaderboard through the GLUE website. Scores are calculated for each task based on the task's individual metrics. All metrics are scaled by 100x (i.e., as percentages). These scores are then averaged to get the final score. For tasks with multiple metrics (including MNLI), the metrics are averaged.

On the leaderboard, only the top scoring submission of a user is shown or ranked by default. Other submissions can be viewed under the expanded view for each user. Competitors may submit privately, preventing their results from appearing. The AI Index visual does not include any private submissions. MNLI matched and mismatched are considered one task for purposes of scoring.

The AI Index has only collected scores that beat scores from previous submissions. If a submission is lower than any of the previous submissions, it is not included in our visual.

Read more about the rules and submission guidelines here.

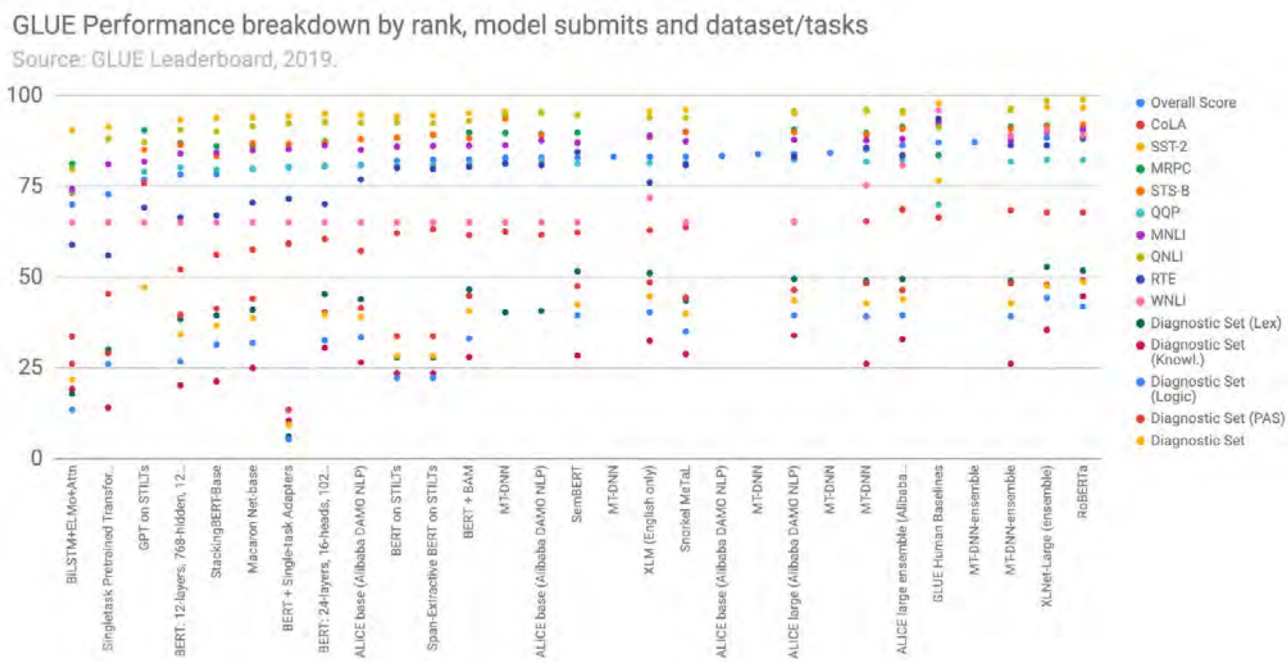

Fig. A3.8.

Return to Technical Performance - Language: GLUE

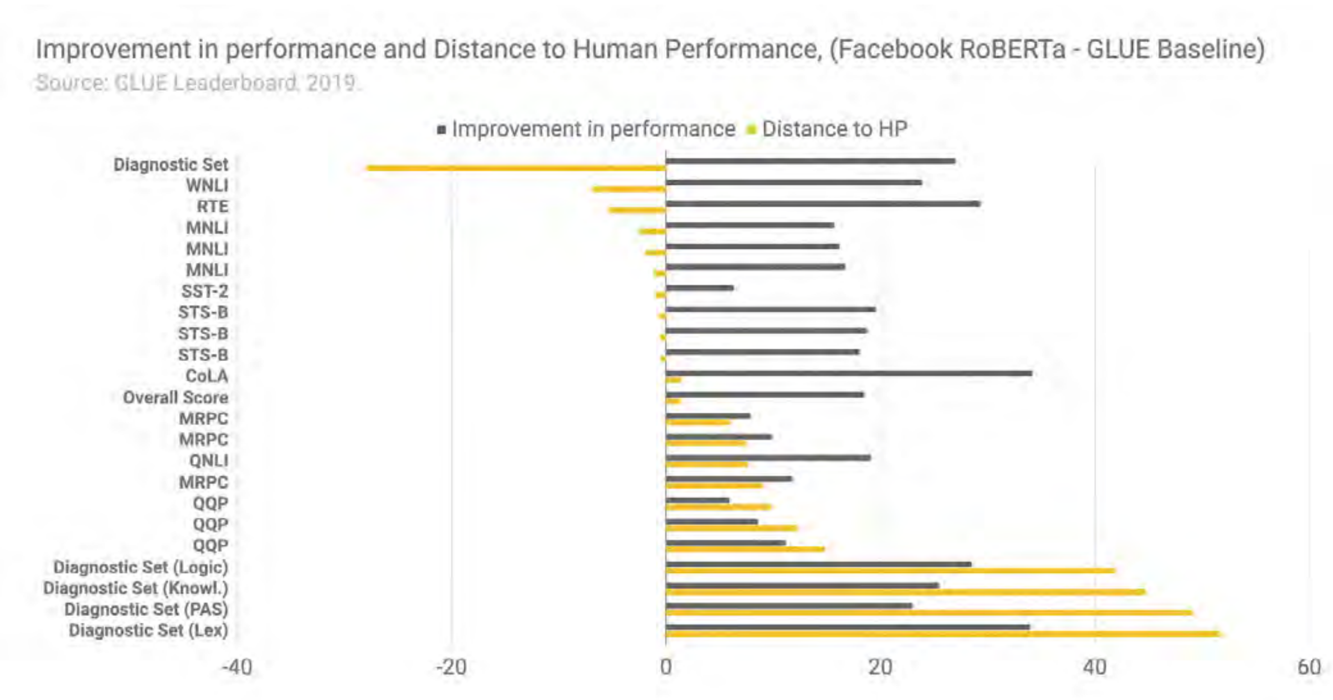

Fig. A3.9.

Return to Technical Performance - Language: SuperGLUE

SuperGLUE

Source

The SuperGLUE benchmark data was pulled from the SuperGLUE Leaderboard. Learn more about the SuperGLUE benchmark. Refer to the SuperGLUE paper and SuperGLUE Software Toolkit for more details.

The different tasks and evaluation metrics for SuperGLUE are presented below.

Name	Identifier	Metric
Broadcoverage Diagnostics	AX-b	Matthew's Corr
CommitmentBank	CB	Avg. F1 / Accuracy
Choice of Plausible Alternatives	COPA	Accuracy
Multi-Sentence Reading Comprehension	MultiRC	F1a / EM
Recognizing Textual Entailment	RTE	Accuracy
Words in Context	WiC	Accuracy
The Winograd Schema Challenge	WSC	Accuracy
BoolQ	BoolQ	Accuracy
Reading Comprehension with Commonsense Reasoning	ReCoRD	F1 / Accuracy
Winogender Schema Diagnostics	AX-g	Gender Parity / Accuracy

Return to Technical Performance - Language: [Reasoning: AI2]() and [ARC]()

Reasoning: AI2 Leaderboards

ARC

Source

AI2 Reasoning Challenge (ARC) is hosted by the [Allen Institute for Artificial Intelligence](). ARC performance data was retrieved from the ARC leaderboards. Find leaderboards for the [easy set]() and the [challenge set]() in the corresponding links.

Methodology

Participants download the ARC data set and submit to the leaderboard through the Allen Institute website.

Examples of questions from the Easy development corpus:

Which technology was developed most recently? (A) cellular telephone (B) television (C) refrigerator (D) airplane [Grade 4]

A student hypothesizes that algae are producers. Which question will best help the student determine if this is correct? (A) Do algae consume other organisms? (B) Which organisms consume algae? (C) Do algae use sunlight to make food? (D) Could an ecosystem survive without algae? [Grade 8]

Examples from the Challenge development corpus:

Juan and LaKeisha roll a few objects down a ramp. They want to see which object rolls the farthest. What should they do so they can repeat their investigation? (A) Put the objects in groups. (B) Change the height of the ramp. (C) Choose different objects to roll. (D) Record the details of the investigation. [Grade 4]

High-pressure systems stop air from rising into the colder regions of the atmosphere where water can condense. What will most likely result if a high-pressure system remains in an area for a long period of time? (A) fog (B) rain (C) drought (D) tornado [Grade 8]

Each question is worth one point. Models are allowed to give multiple answers, in which case a model that gives N answers gets 1/N points if one of its N answers is correct, and 0 otherwise. The overall score is the average of the scores of the individual questions.

The AI Index has only collected scores that beat scores from previous submissions. If a submission is lower than any of the previous submissions, it is not included in our visual.

Read more about the rules and submission guidelines [here]().

Return to Technical Performance - Language: SQuAD

SQuAD

SQuAD 1.1, the previous version of the SQuAD dataset, contains 100,000+ question-answer pairs on 500+ articles. here are a few NLP competitions on CodaLab Worksheets

https://codalab-worksheets.readthedocs.io/en/latest/Competitions/#list-of-competitions

Training Time and Cost on SQuAD

Total cost for public cloud instances to train a question answering model to a F1 score of 0.75 or greater on the SQuAD development dataset

Source: Stanford DAWNBench, 2019.

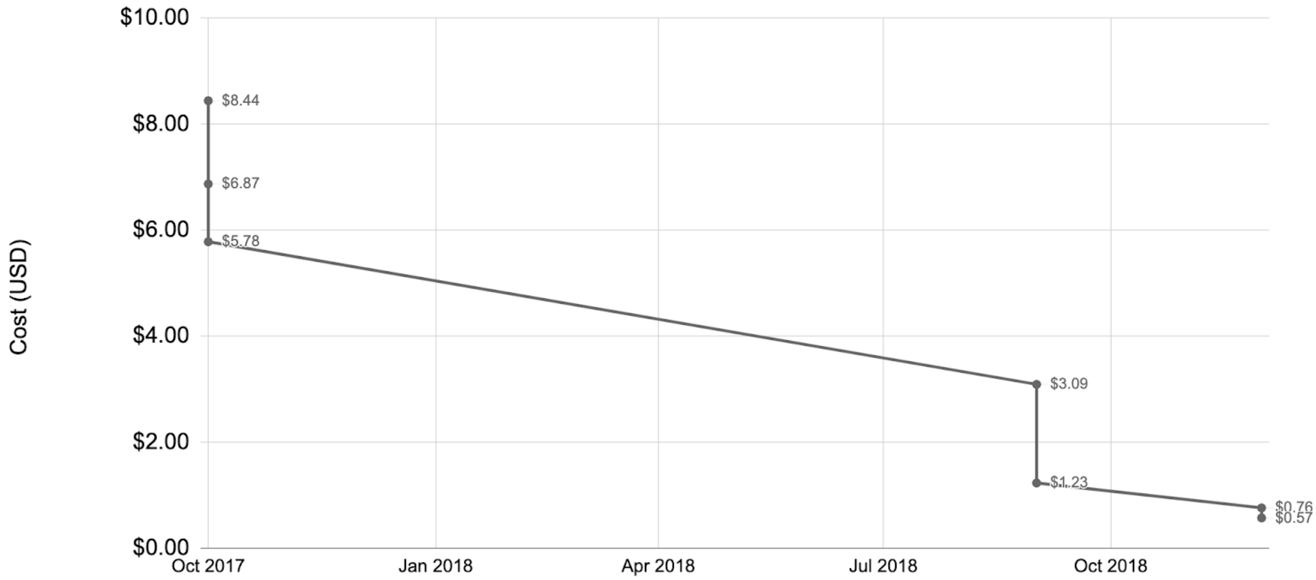

Fig. A3.10a.

Latency required to answer one SQuAD question using a model with a F1 score of at least 0.75 on the development dataset.

Source: Stanford DAWNBench, 2019.

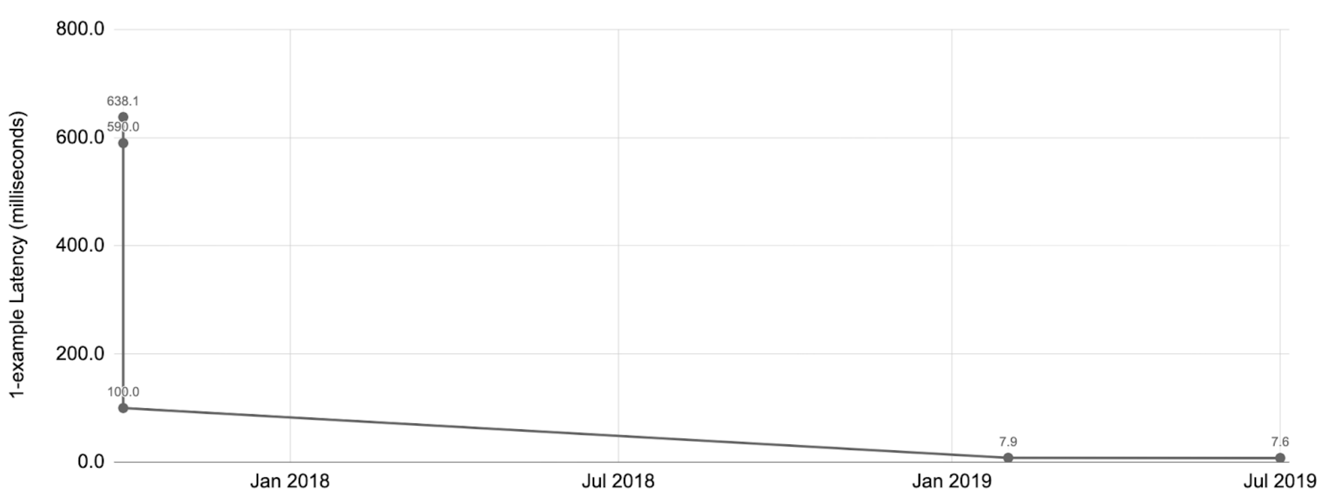

Fig. A3.10b.

Machine Translation (Commercial System)

Source

Intento provides evaluation for third party AI model to help use the right models in a vendor-agnostic fashion.

Learn more about Intento here.

Evaluation Scope

There are more than 20 commercial MT systems with pre-trained models, provided by Alibaba, Amazon, Baidu, CloudTranslate, DeepL, Google, GTCom, IBM, Microsoft, ModernMT, Naver, Niutrans, PROMT, SAP, SDL, Sogou, Systran, Tilde, Tencent, Yandex, and Youdao.

The accuracy of a given MT system for a specific translation project depends on a number of factors: linguistic performance on the language pair, amount of in-domain data in the training set, available means of domain adaptation, learning curve, and data quality tolerance to name a few.

The benchmark studies used in this report evaluate only translation accuracy over different language pairs. Other factors are controlled: the most general-purpose domain (News) was used and only pre-trained models considered. Also, only commercial systems were considered assuming them using all commercially reasonable efforts to acquire training data and improve performance.

It should be noted that the translation accuracy is understood in a very narrow meaning -- a distance from reference human translation, calculated using a specialized metric (hLEPOR).

Language pair selection

Combined, all systems studied support 14136 language pairs (as of June 2019). Ideally, performance for every one of them would be evaluated, even if it's supported by a single MT system.

Several factors limit the scope of our study. Very few datasets are publicly available (general purpose, with a license to use in evaluation and relatively low amount of noise) and studies must be performed under limited time and budget.

To prioritize language pairs, we referred to the Web Content Language Popularity index (https://w3techs.com/technologies/overview/content_language/all). We split all languages into four groups based on the percentage of websites in this language: ≥ 2.0%, 0.5%-2%, 0.1-0.3%, <0.1%. The first group contains English, Russian, Japanese, German, Spanish, French, Portuguese, Italian and Chinese. We focused our effort on 16 language pairs between English and this first group of languages. Later, some language pairs were added between those languages (without English) and between English and some languages of the second group, as shown in the picture below.

Language pairs without English were selected based on dataset availability., To avoid this selection bias, they are not included in this overview report. detailed information on them can be found in this report.

Dataset selection

We have to make several choices around the dataset selection.

Public datasets are good at keeping the evaluation transparent and reproducible. The potential downside is that they may be (and probably are) used by every MT provider to train their models. Private datasets provide a cleaner experiment, but the study is impossible to reproduce. We have made several experiments and found no signs of NMT overfitting in the scores of sentences from the public datasets. Hence, we decided to follow the public dataset path.

Here's the full list of datasets used in the last study (June 2019):

- WMT-2013 (translation task, news domain) - en-es, es-en
- WMT-2015 (translation task, news domain) - fr-en, en-fr
- WMT-2016 (translation task, news domain) - ro-en, en-ro
- WMT-2018 (translation task, news domain) - tr-en, en-tr, cs-en
- WMT-2019 (translation task, news domain) - zh-en, en-zh, en-cs, de-en, en-de, ru-en, en-ru, fi-en, en-fi, de-fr, fr-de
- NewsCommentary-2011 - en-ja, ja-en, en-pt, pt-en, en-it, it-en, ru-de, de-ru, ru-es, ru-fr, ru-pt, ja-fr, de-ja, es-zh, fr-ru, fr-es, it-pt, zh-it, en-ar, ar-en, en-nl, nl-en, de-it, it-de, ja-zh, zh-ja
- Tatoeba, JHE - en-ko, ko-en

Return to Technical Performance - Language: Commercial Machine Translation

Dataset selection

Virtually every dataset we selected contains some amount of noise. We decided not to invest in the dataset cleaning, considering that dealing with the source noise (grammatical issues and typos) is one of the MT success factors and a small number of mistranslations won't skew the relative MT quality picture. We have to decide how many sentences include in the test set. We tried a different size of random samples and analyzed how the average score changes with adding more sentences. For most of the language pairs, we found that average score converges after 1,500 sentences, hence we randomly sampled 2,000 sentences for each language pair.

Historical consistency

Another choice we had to make was if we should keep the same datasets we used in the initial May 2017 benchmark or update the datasets to the latest available. We preferred relevance over historical consistency. We observed that updating the dataset may change quality scores up to 10% in either direction, correlated across all MT providers.

Evaluation metric

We use LEPOR metric: automatic machine translation evaluation metric considering the enhanced Length Penalty, n-gram Position difference Penalty and Recall. We found it more reliable than BLEU, because it combines both precision and recall, and also because it may be reliably used both on corpus and sentence levels.

https://www.slideshare.net/AaronHanLiFeng/lepor-an-augmented-machine-translation-evaluation-metric-thesis-ppt

https://github.com/aaronlifenghan/aaron-project-lepor

In our evaluation, we used hLEPORA v.3.1A (best metric at the ACL-WMT 2013 contest). The score for a test set is calculated as an average of the sentence scores.

For hieroglyphic languages, we performed the tokenization similar to used by WMT (https://www.statmt.org/wmt17/tokenizeChinese.py)

Return to Technical Performance - Language: [Commercial Machine Translation](#)

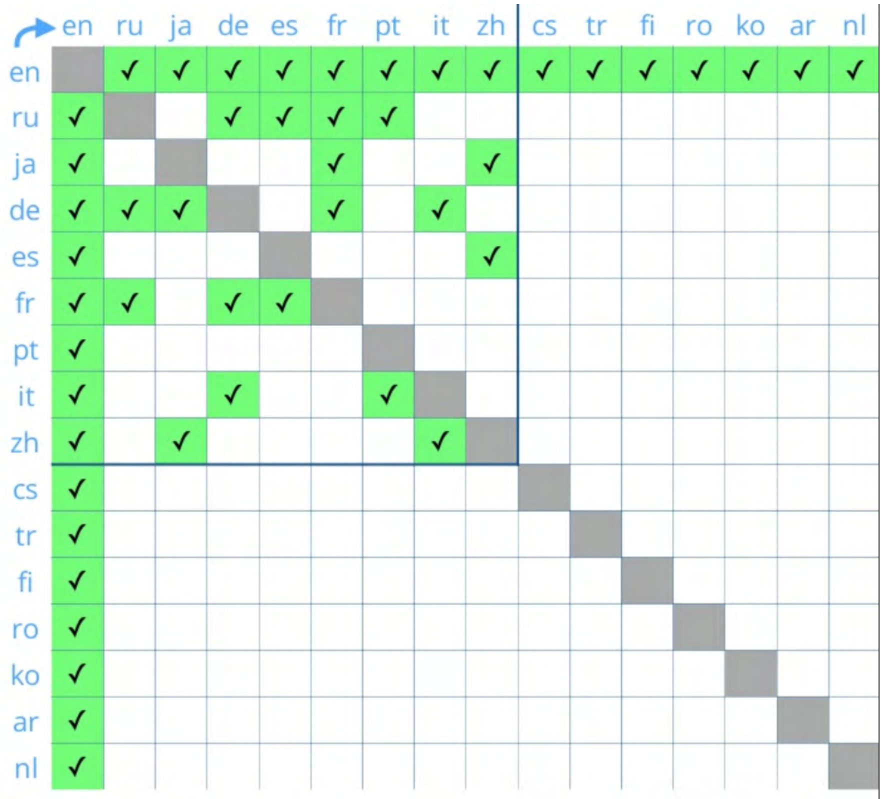

Fig. A3.11a.

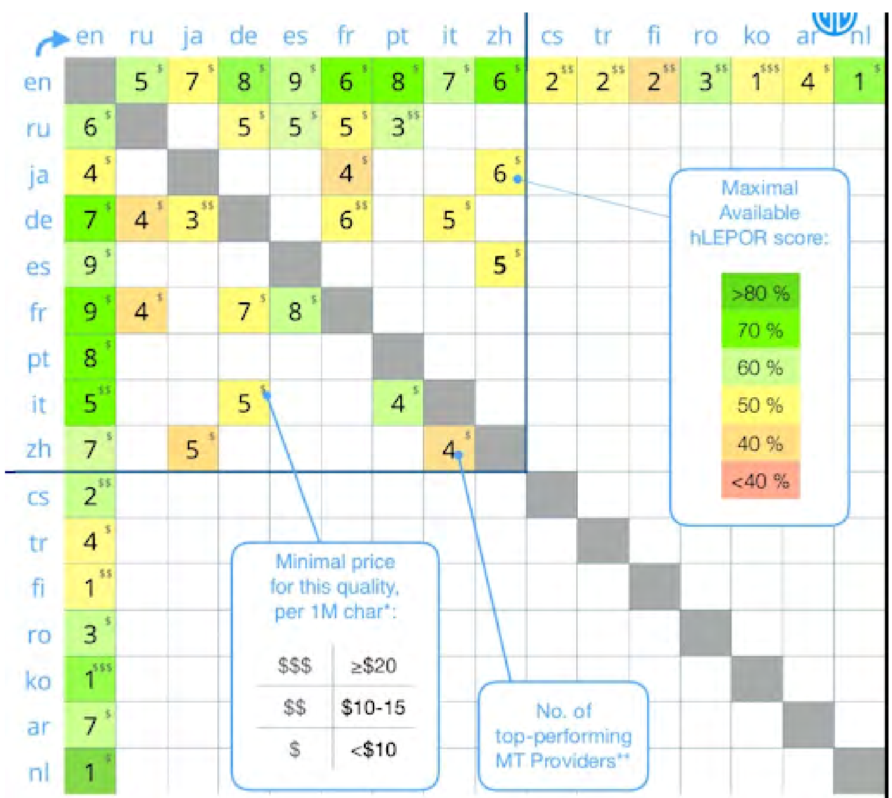

Fig. A3.11b.

Return to Technical Performance - Omniglot

Omniglot Challenge

The Omniglot challenge is to build a single model that can perform five-concept learning tasks at a human level (see Figure A3.12a). The authors measured the quality of the samples using visual Turing tests (Figure 6 in Human-level concept learning through probabilistic program induction). Unfortunately subsequent work hasn't adopted this metric, and thus there is no metric that makes these samples directly comparable beyond their visual form. The Omniglot Challenge: a 3-year progress report.

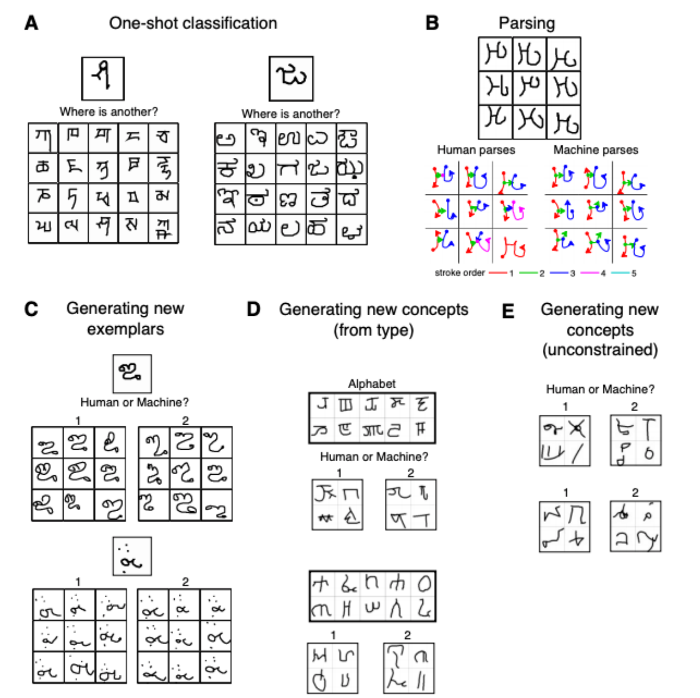

Fig. A3.12a.

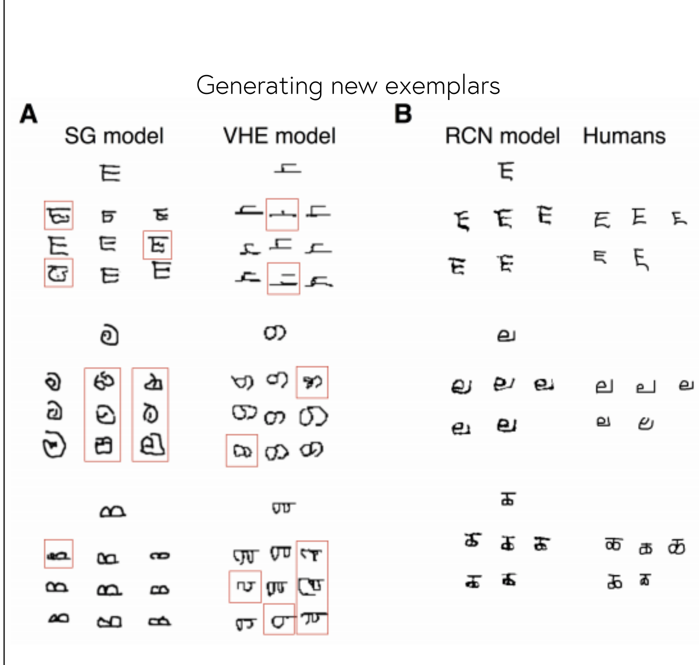

Fig. A3.12b.

Fig. A3.12a. The Omniglot challenge of performing five concept learning tasks at a human level. A) Two trials of one-shot classification, where a single image of a new character is presented (top) and the goal is to select another example of that character amongst other characters from the same alphabet (in the grid below). In panels B-E, human participants and Bayesian Program Learning (BPL) are compared on four tasks. B) Nine human drawings (top) are shown with the ground truth parses (human) and the best model parses (machine). C) Humans and BPL were given an image of a new character (top) and asked to produce new examples. D) Humans and BPL were given a novel alphabet and asked to produce new characters for that alphabet. E) Humans and BPL produced new characters from scratch. The grids generated by BPL are C (by row): 1, 2; D: 2, 2; E: 2, 2. Reprinted from The Omniglot Challenge: a 3-year progress report.

Fig. A3.12b. Generating new exemplars with deep neural architectures. The task is to generate new examples (shown in grid) given an image of a new character (above each grid). A) The sequential generative model and variational homoencoder produce compelling examples in some cases, while showing too much variation in others (highlighted in red). B) The recursive cortical network (RCN) produces reasonable new examples but has too little variation relative to human examples from Omniglot, suggesting the model is not capturing all the degrees of freedom that people grasp in these concepts. Reprinted from The Omniglot Challenge: a 3-year progress report.
.

4.1 Jobs

Diverse datasets are introduced for the first time with a deeper focus on cross country, sub-national, sectoral, and gender related labor market metrics. The goal of AI labor market metrics should be not just to provide the evolution of volume to represent proxies for job growth but also quality, sophistication, and complexity of AI related labor supply and demand. These diverse metrics help to provide a more complete picture of AI and its impact on the labor market than before. The comprehensive list of metrics is provided below in the Appendix Table.

So far metrics have provided online job posting measures that provided a perspective on labor demand. The various new metrics include (a) AI job posting per million from Indeed; (b) AI jobs posted across jobs sites presented as a share of total jobs and as a share of IT jobs online by Burning Glass. This metric is available for 5 countries and regional data for almost 400 hundred metropolitan areas in the US; (c) AI Hiring index for almost 30 countries which measures the relative growth in AI hiring, (d) AI Skill penetration rate, skill penetration relative to the global AI skill penetration available for countries and regions within the United States by LinkedIn.

Table A4.1. Summary of Job Metrics

Metric	Definition	Source	Years	Freq	County coverage
AI hiring index	AI hiring rate is the percentage of LinkedIn members who had any AI skills (see appendix 2 for the AI skill grouping) on their profile and added a new employer to their profile in the same month the new job began, divided by the total number of LinkedIn members in the country. This rate is then indexed to the average month in 2015-2016; for example, an index of 1.05 indicates a hiring rate that is 5% higher than the average month in 2015-2016.	LinkedIn Economic Graph	2015-19	M	28
AI jobs posted (per million jobs)	Number of AI jobs posted per million jobs posted	Indeed	2015-19	M	5
AI jobs posted (% of total jobs/% of IT jobs)	AI job postings as a percent of total jobs posted, or as a percent of IT job	BurningGlass	2010-2019	A	5
AI Skill Penetration Index	Relative skill penetration rate (this is a method to compare how prevalent AI skills are for each country against a global average/benchmark based on the same set of occupations)	LinkedIn Economic Graph	2018	A	15
LinkedIn members with AI skills	Total number of LinkedInmembers with AI skills on their profile	LinkedIn Economic Graph	2015-2019	M	28
Count of AI hire	Total number of AI hires on LinkedIn	LinkedIn Economic Graph	2015-2019	M	28

Return to Economy - Jobs: Global Hiring, US Metropolitan Areas and Cities

LinkedIn

Source

This is the first data collaboration effort between LinkedIn Economic Graph team and Stanford AI Index team. The goal is to jointly publish metrics that measure AI technology adoption and AI talent characteristics based on LinkedIn data in the 2019 annual report from Stanford AI Index. We hope this will be the starting point for more extensive research collaboration around the AI theme between the two teams.

Methodology for AI Hiring Index:

AI hiring rate is the percentage of LinkedIn members who had any AI skills on their profile and added a new employer to their profile in the same month the new job began, divided by the total number of LinkedIn members in the country. By only analyzing the timeliest data, we can make month-to-month comparisons and account for any potential lags in members updating their profiles. This rate is then indexed to the average month in 2015-2016; for example, an index of 1.05 indicates a hiring rate that is 5% higher than the average month in 2015-2016.

Sample: Countries were included if they met the following conditions: 1) sufficient labor force coverage (roughly >40%); and 2) at least 10 AI talents in any given month. Countries meeting these conditions are: United States, Netherlands, Ireland, Denmark, Australia, United Kingdom, Luxembourg, Canada, Singapore, Belgium, New Zealand, Norway, Sweden, United Arab Emirates, France, Portugal, Switzerland, Chile, Spain, Italy, Hong Kong (SAR), Finland, Israel, Costa Rica, Brazil.

China and India were included in the sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly.

AI Growth and Economic Development

AI also offers opportunities for labor reallocation and job creation and to (?) address the growing polarization of labor markets. The demand for jobs for the future global workforce would be led by the technology infrastructure that powers the system of AI applications. As automation of routine tasks and codification of job tasks becomes more prevalent and macro-critical, national economies are already starting to reallocate labor based on such consumer preferences and forces of global demand.

Trends with economic stages of development i.e. natural log of GDP per capita and economic growth i.e. GDP per capita growth are plotted below. A normalized version of the metric is the percent of LinkedIn members with AI skills plotted against the stages of development and economic growth. The relationship is statistically significant and positive. In particular, it is noted that Israel, Singapore, Finland, India, and Greece are positive outliers, indicating higher relative AI specialization on LinkedIn than it would be predicted by the stage of development or level of economic growth.

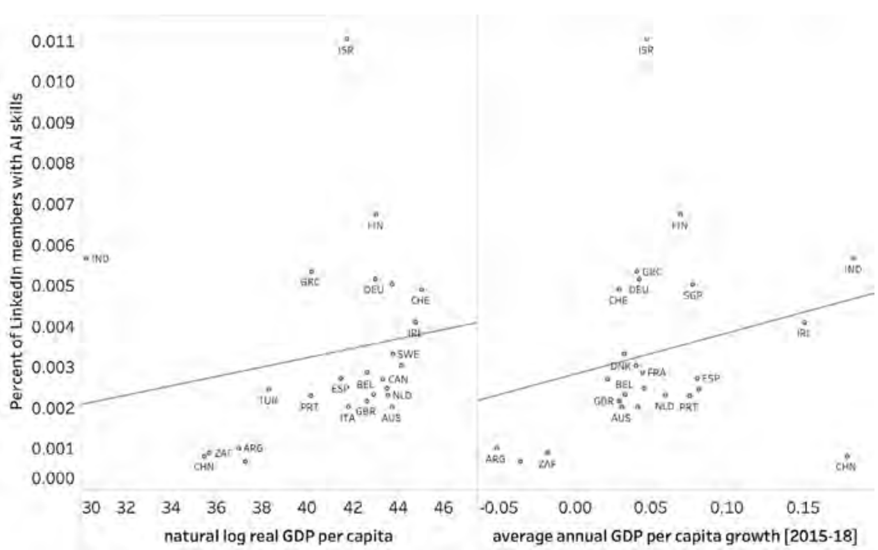

Fig. A4.1.

[Jobs]_[Appendix_Start]
[Access_Data]

AI skill penetration by sectors for countries over time

Methodology

Relative skill penetration rate (this is a method to compare how prevalent AI skills are for each country against a global average/benchmark based on the same set of occupations)

Sample Specifications

- Sample sectors: Software & IT Services, Hardware and Networking, Education, Finance, Manufacturing. These are the top 5 sectors with the highest AI skill penetration globally. Data is pooled for these 5 sectors.
- Sample countries: All countries with at least 40% labor force coverage and sufficient occupation and skill data. China and India also were included in this sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly.
- Sample timeframe: Pooled skill adds during 2015 to 2018.

3 steps to calculate relative skill penetration

1. Identify top 50 skills in each occupation in each country: use TF-IDF approach to give higher weights to skills that are added by more members and are more unique for each occupation.

2. Calculate penetration rates by dividing the number of AI skills (using LinkedIn taxonomy of AI skill groups - Box "Artificial Intelligence Top Skill Names from LinkedIn Economic Graph") over the total number of skills (50) for each occupation and each country

3. Calculate relative penetration rates by taking the ratio between the average penetration rates across all occupations in a given country, and the global average penetration rate of AI skills across the countries for the same set of occupations.

Metric interpretation: "For a given country, the relative skill group penetration is the ratio between the penetration rate of a given skill group in each country, and the global average penetration rate."

For example: "Ranking first among selected countries, the average penetration of AI skills in India in selected sectors is 2.5 times the global average across the same set of occupations."

Methodology published in the World Bank publication "The Future of Work in Africa : Harnessing the Potential of Digital Technologies for All"

Return to Economy - Jobs: Skill Penetration

Occupations with High Skill Similarity between Genders (0.95 or higher cosine similarity)

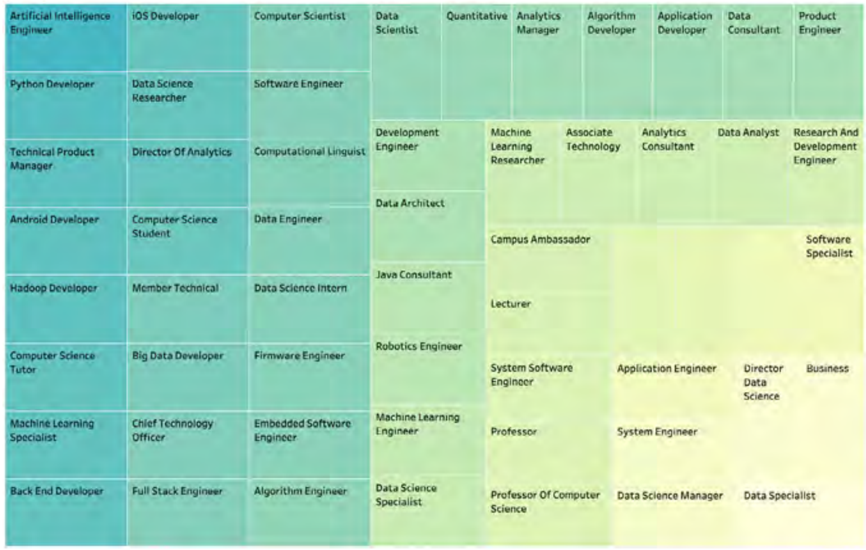

Occupations with Medium Skill Similarity between Genders (0.90-0.95 cosine similarity)

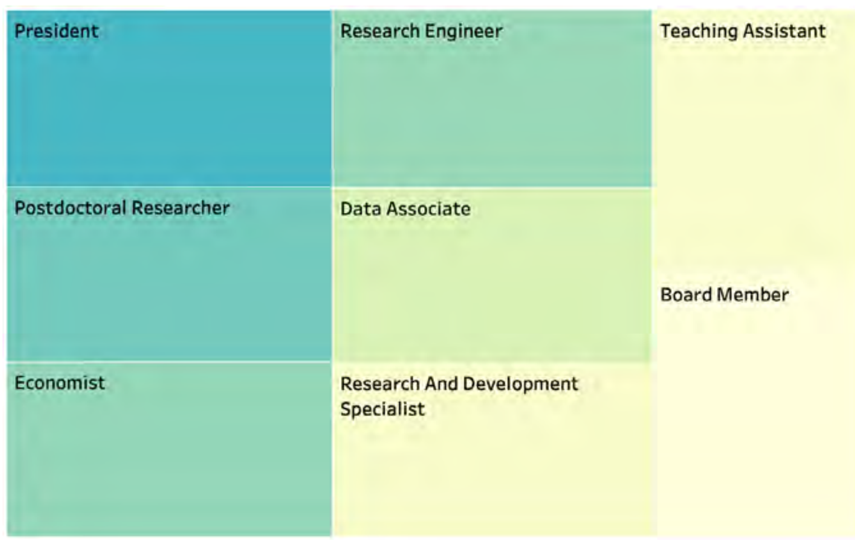

Occupations with Low Skill Similarity between Genders (0.95 or lower cosine similarity)

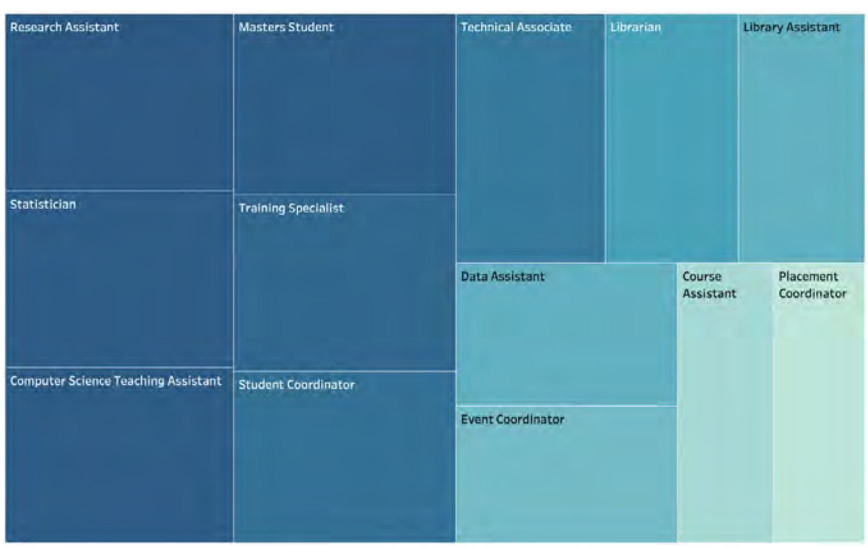

Fig. A4.2a, b, & c.

Return to Economy - Jobs: Skill Penetration

Country and Occupational Group: AI Skill Similarity

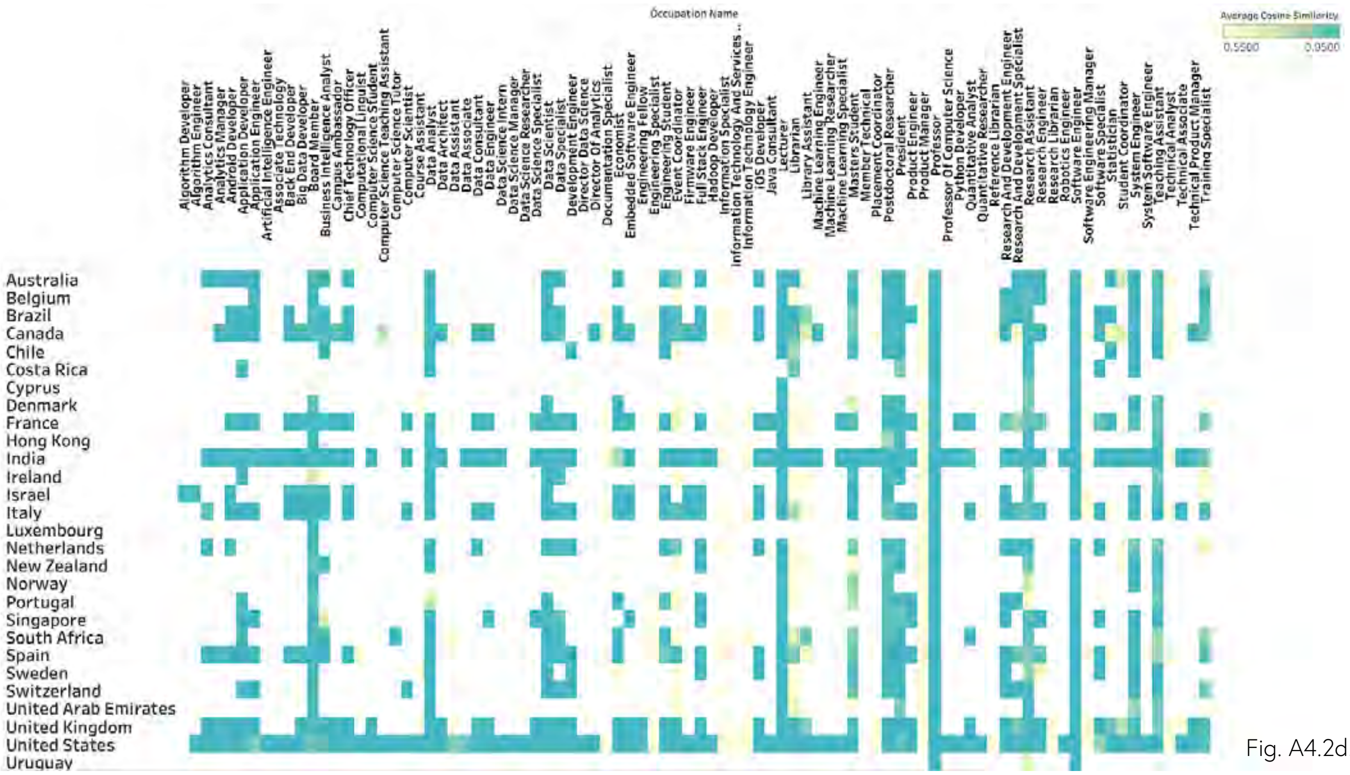

Fig. A4.2d.

Relative Penetration of AI Skills and Number of AI Occupations

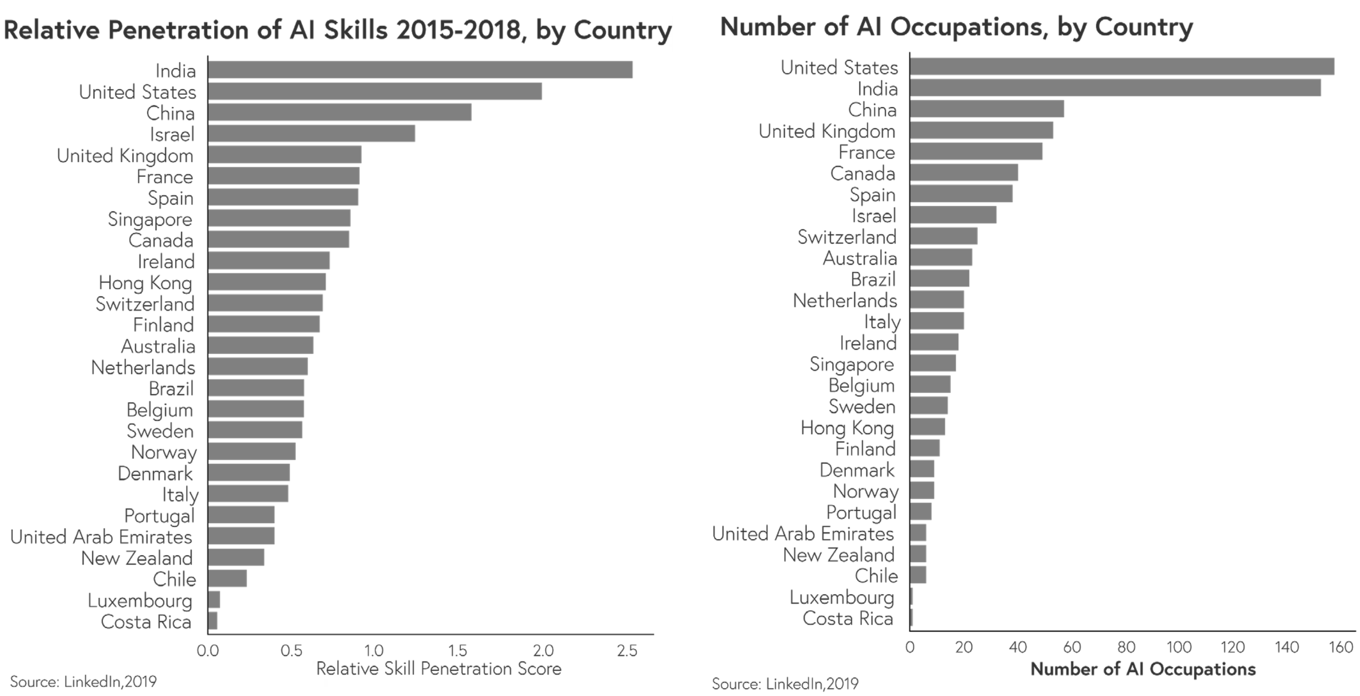

Fig. A4.3.

Notes: China and India were included in this sample due to their increasing importance in the global economy, but LinkedIn coverage in these countries does not reach the 40% of the workforce. Insights for these countries may not provide as full a picture as other countries, and should be interpreted accordingly.

AI skill penetration by US regions

Methodology

City sample: US cities with at least 500 skills that were added between 2015 and 2018 (to ensure comparability and adequate skill information).

AI skill penetration measures the concentration of AI skills among top skills in each city that are added by LI members. The metric measures the number of AI skills (defined by the AI skill group - See Box "Artificial Intelligence Top Skill Names from LinkedIn Economic Graph") among the top 500 skills for each city and top 500 number of skills in each city.

- Sample: Top 10 cities with highest AI skill penetration 2018.
- Results: AI skills have been increasing across many cities in the US. The leading ones include tech hubs such as San Francisco and Seattle, but also university towns such as Bryan-College Station (TX), Lafayette (IN), Binghamton (NY) and Urbana-Champaign (IL), suggesting that R&D talents and training programs in AI are rising to catch up with the industry trend.

Table A4.2. Ranking of US Regions based on AI Skill Penetration, 2018 (Rank: 20-50)

Source: LinkedIn Economic Graph, 2019.

City	Rank	City	Rank
Provo, UT	**20**	Chicago, IL	**36**
Salt Lake City, UT	**21**	Iowa City, IA	**37**
New York City, NY	**22**	Tallahassee, FL	**38**
Detroit, MI	**23**	Tucson, AZ	**39**
Syracuse, NY	**24**	Albuquerque, NM	**40**
Roanoke, VA	**25**	Denver, CO	**41**
Sacramento, CA	**26**	Baltimore, MD	**42**
Washington, D.C.	**27**	Lubbock, TX	**43**
Athens, GA	**28**	Burlington, VT	**44**
Los Angeles, CA	**29**	Tuscaloosa, AL	**45**
Fort Collins, CO	**30**	Lincoln, NE	**46**
Dayton, OH	**31**	Knoxville, TN	**47**
Charlottesville, VA	**32**	Charlotte, NC	**48**
Phoenix, AZ	**33**	Allentown, PA	**49**
Albany, NY	**34**	Dallas-Fort Worth, TX	**50**
Orange County, CA	**35**		

AI skill penetration by US regions

Table A4.3. Ranking of US Regions based on AI Skill Penetration, 2018 (Rank: 50-80)
Source: LinkedIn Economic Graph, 2019.

City	Rank	City	Rank
Providence, RI	**51**	Miami-Fort Lauderdale, FL	**67**
Hartford, CT	**52**	Nashville, TN	**68**
St. Louis, MO	**53**	Lansing, MI	**69**
Columbus, OH	**54**	Columbia, SC	**70**
Tampa-St. Petersburg, FL	**55**	Minneapolis-St. Paul, MN	**71**
Atlanta, GA	**56**	Philadelphia, PA	**72**
Pocatello, ID	**57**	Buffalo-Niagara, NY	**73**
Melbourne, FL	**58**	Norfolk, VA	**74**
Fayetteville, AR	**59**	Peoria, IL	**75**
Houston, TX	**60**	Hickory-Lenoir, NC	**76**
El Paso, TX	**61**	Fargo, ND	**77**
Columbia, MO	**62**	Hawaii	**78**
Spokane, WA	**63**	Davenport, IA	**79**
Portland, OR	**64**	Bloomington-Normal, IL	**80**
Cleveland-Akron, OH	**65**		
Huntsville, AL	**66**		

Return to Economy - Jobs: Skill Penetration

AI skill penetration by US regions

Table A4.4. Ranking of US Regions based on AI Skill Penetration, 2018 (Rank: 80-100)

Source: LinkedIn Economic Graph, 2019.

City	Rank	City	Rank
Augusta, GA	**81**	Indianapolis, IN	**97**
Topeka, KS	**82**	Greenville, SC	**98**
Asheville, NC	**83**	Kalamazoo, MI	**99**
Colorado Springs, CO	**84**	Greensboro-Winston-Salem, NC	**100**
Orlando, FL	**85**		
Lexington, KY	**86**		
Killeen-Temple, TX	**87**		
Clarksville, TN	**88**		
Lafayette, LA	**89**		
Bakersfield, CA	**90**		
Oklahoma City, OK	**91**		
Eugene, OR	**92**		
Boise, ID	**93**		
Wichita, KS	**94**		
San Antonio, TX	**95**		
Cincinnati, OH	**96**		

AI skill penetration by US regions

Table A4.5. Ranking of US Regions based on AI Skill Penetration, 2018 (Rank: 100-118)
Source: LinkedIn Economic Graph, 2019.

City	Rank	City	Rank
Sioux Falls, SD	**101**	Birmingham, AL	**117**
Reno, NV	**102**	Little Rock, AR	**118**
Kansas City, MO	**103**		
West Palm Beach, FL	**104**		
Memphis, TN	**105**		
Grand Rapids, MI	**106**		
Toledo, OH	**107**		
Lancaster, PA	**108**		
Scranton, PA	**109**		
Omaha, NE	**110**		
Richmond, VA	**111**		
Milwaukee, WI	**112**		
Harrisburg, PA	**113**		
Louisville, KY	**114**		
Baton Rouge, LA	**115**		
Tulsa, OK	**116**		

Box. Artificial Intelligence Top Skill Names from LinkedIn Economic Graph
LinkedIn members self-report their skills on their LinkedIn profiles. Currently, there are more than 35,000 distinct, standardized skills classified by LinkedIn. These have been coded and classified by taxonomists at LinkedIn into 249 skill groupings, which are the skill groups represented in the dataset. This analysis focuses on AI and NLP skill groups, including the following top individual skills: Machine Learning, Data Structures, Artificial Intelligence, Computer Vision, Apache Spark, Deep Learning, Pattern Recognition, OpenCV, Artificial Neural Networks, Neural Networks, NumPy, Weka, Information Extraction, Scikit-Learn, Lisp, Recommender Systems, Classification, Graph Theory, SciPy, Support Vector Machine (SVM), Reinforcement Learning, Statistical Inference, Web Mining, Computational Intelligence, among others.

Indeed

Source

Indeed is an employment-related search engine for job listings. Learn more about Indeed here.

Methodology

For indeed, job postings where the title contains one or more of the following terms: "artificial intelligence," "ai engineer," "ai research," "ai researcher," "ai scientist," "ai developer," "ai technical," "ai programmer," ai architect," "machine learning," "ml engineer," "ml research," "ml researcher," "ml scientist," "ml developer," "ml technical," "ml programmer," "ml architect," "natural language processing," "nlp," "deep learning," "computer vision," "robotics engineer," "robotics research," "robotics researcher," "robotics scientist," "robotics developer," "robotics technical," "robotics programmer," "robotics architect."

English speaking countries and coverage details

The following countries are available to use the English terms in: United States, Canada, Great Britain, Australia, Ireland, New Zealand, United Arab Emirates, Bahrain, India, Kuwait, Malaysia, Oman, Philippines, Pakistan, Qatar, Singapore, South Africa. As they cannot all be pulled individually due sample size concerns, we suggest the approach be bucketing them geographically as this should alleviate most of the sample size issues (Australia+New Zealand, Singapore+Malaysia+Philippines, etc.). We have indexed the shares per million so that countries are more directly comparable.

Definition

For the breakout of the components, if the title of a postings was "artificial intelligence engineer / machine learning engineer," it would be bucketed in AI as well as in ML since it contains the respective terms for both classifications. For the other metrics, where the various components were not broken out, it was counted once via an OR statements filter.

The publicly available dataset provides the following metrics:

- The share of postings per million by type of AI job posting. The time range is from March 2014 to March 2019, by month and is for US job postings. The definition of AI is detailed below in the methodology note; for example, both "natural language processing" or "nlp" were used to identify natural language processing postings, "machine learning" or "ml" to identify machine learning postings, etc.
- The share of postings per million, by country over time. The time range is from March 2014 to March 2019. The definition of AI job postings was consistent across countries and is detailed below.
- The top ten metropolitan statistical areas (MSA) and their respective percentage of AI jobs out of all AI jobs located within a US MSA. Note there are more than ten listed due to a tie between a few of the MSAs.
- he top ten states and their respective percentage of AI jobs out of all AI jobs located in a US state. Note there are more than ten listed due to a tie between a few of the states.

Methodology: For this data, the definition of AI jobs were job postings whose title contained the terms "artificial intelligence," "machine learning," "deep learning," "natural language processing" or "npl." We also included "ai" and "ml," though with some caveats to ensure against false positives. Please note this definition of AI job postings is slightly different from previous definitions used by Indeed.

Return to Economy - Jobs: [US Labor Demand By Job Cluster]

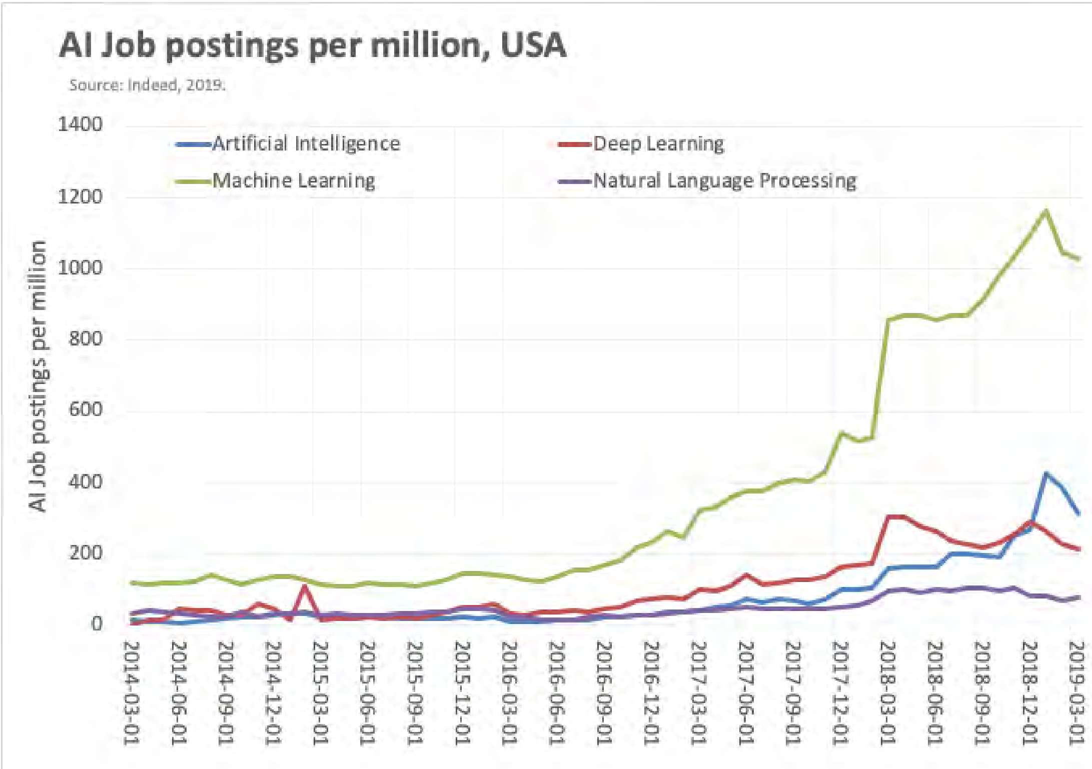

Fig. A4.4.

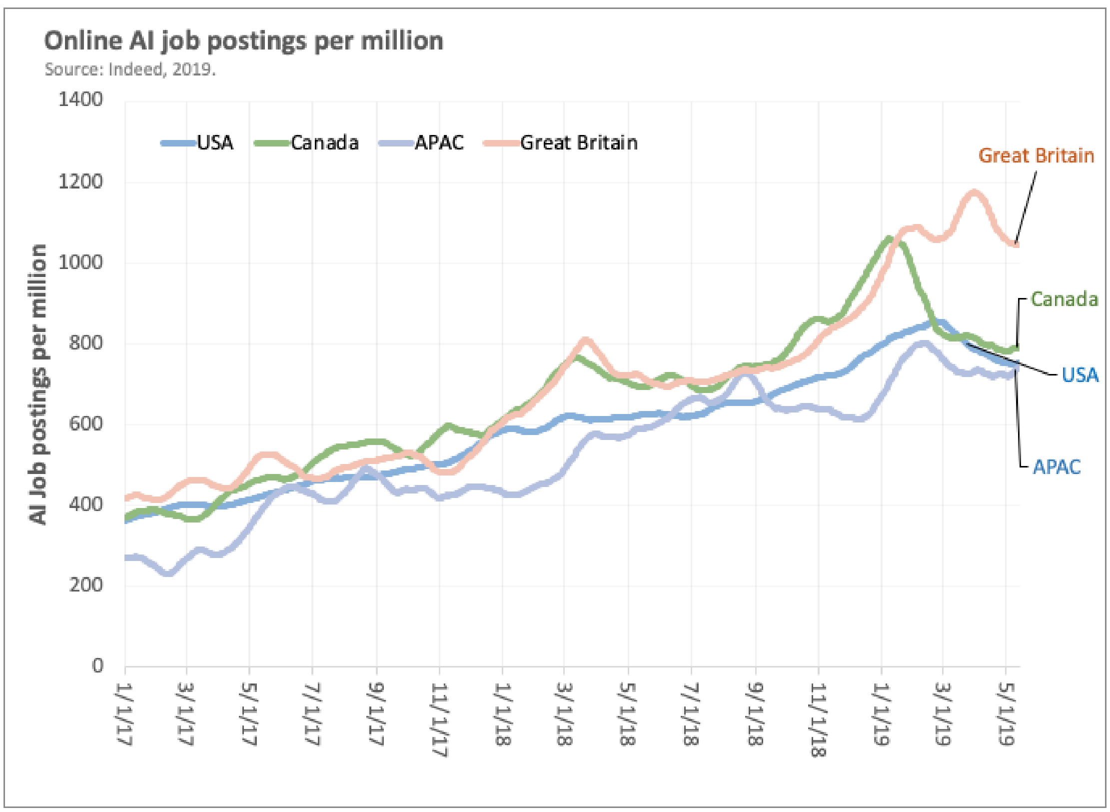

Fig. A4.5.
Notes: The countries the pull was possible for was the US, Canada, Great Britain & Ireland and APAC.
APAC includes: Australia, New Zealand, Malaysia, the Philippines and Singapore.

Return to Economy - Jobs: US Labor Demand By Job Cluster

AI Labor Demand Growth by Clusters on Indeed

Data from Indeed is presented below for the US, where the left axis presents the AI jobs posted per million jobs on Indeed. Job posts mentioning Machine Learning captured the largest proportion of AI jobs posted (58% of AI jobs and 0.003% of the total jobs posted), followed by Artificial Intelligence (24% of AI jobs and 0.001% of the total jobs), Deep Learning (9% of AI jobs and 0.0007% of total jobs), and NLP (8% of AI jobs and 0.0002% of the total jobs).

Between 2015 and 2018, Deep Learning grew the fastest by over 12x, followed by Artificial Intelligence (almost 5x), Machine Learning (4x), and NLP (2x). It is noted that the share of AI jobs as a percent of the total jobs posted remains smaller than 1% on Indeed.

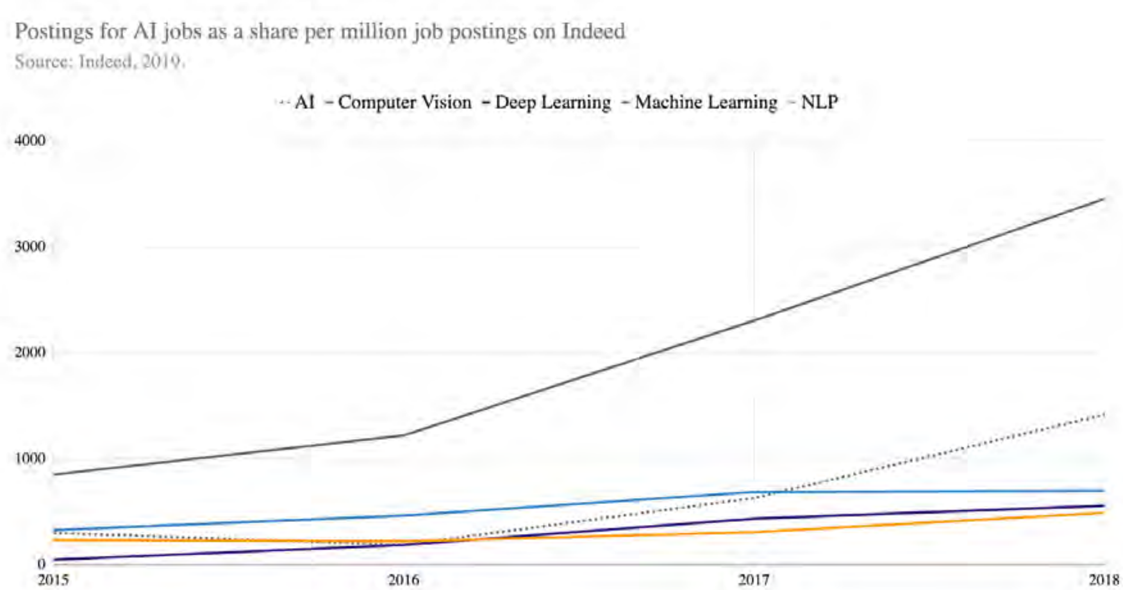

Fig. A4.6a.

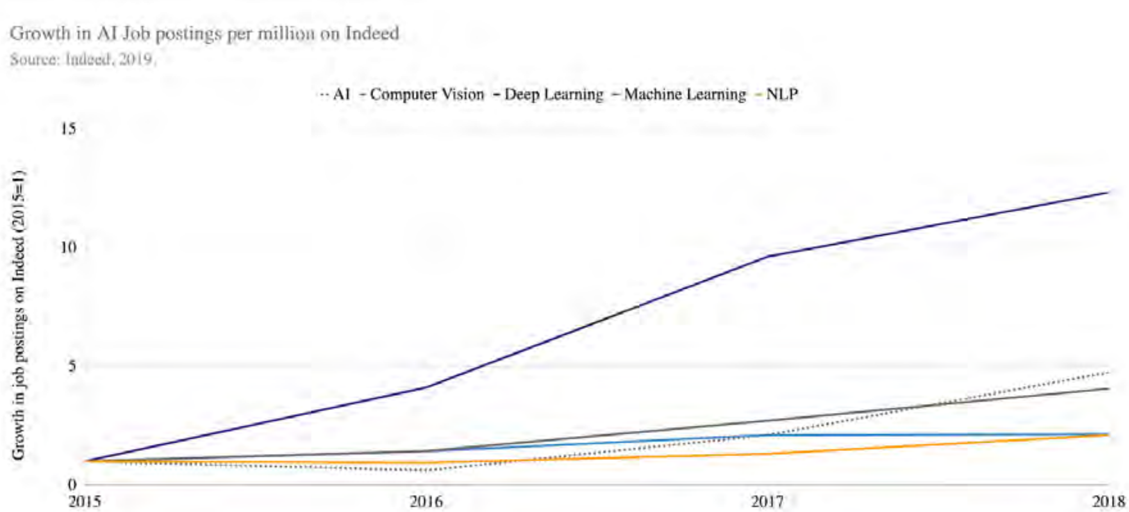

Fig. A4.6b.

Return to Economy - Jobs: Sectoral Diffusion

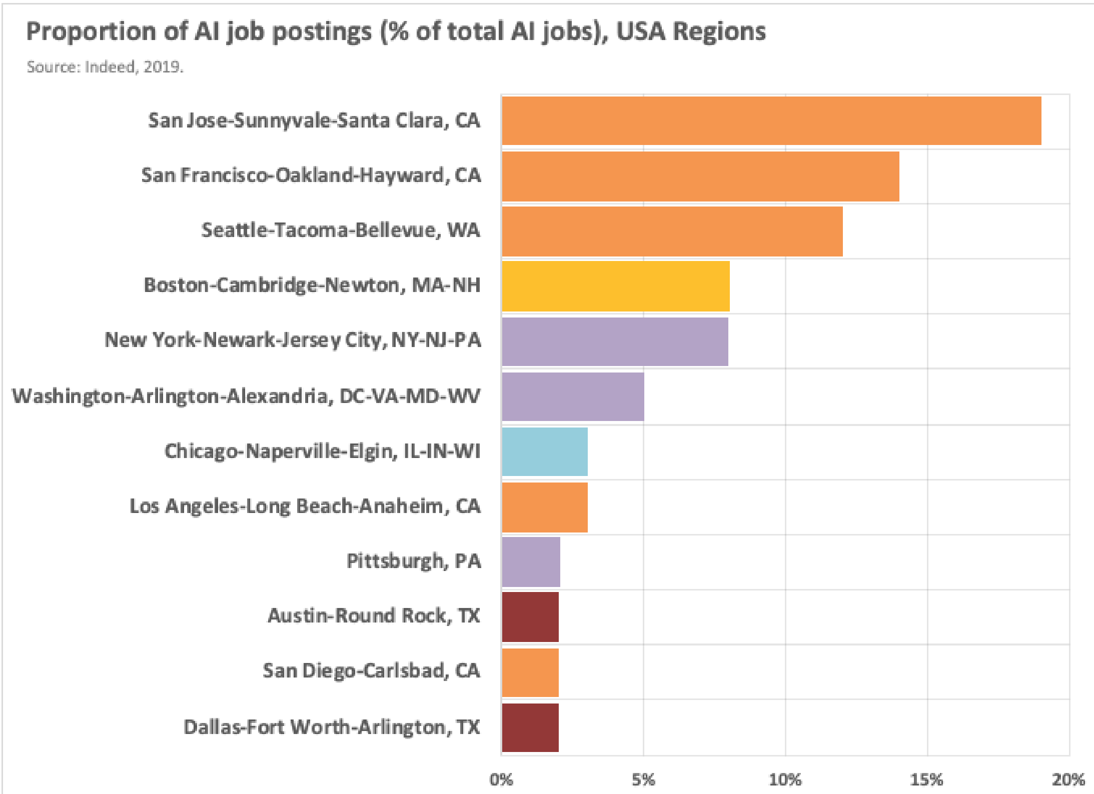

	West Coast	50%
	New England	8%
	East Coast	15%
	IL-IN-WI	3%
	Texas	4%
	Others	20%

Fig. A4.7.

Burning Glass Technologies

About Burning Glass Technologies

Burning Glass Technologies delivers job market analytics that empower employers, workers, and educators to make data-driven decisions. The company's artificial intelligence technology analyses hundreds of millions of job postings and real-life career transitions to provide insight into labor market patterns. This real-time strategic intelligence offers crucial insights, such as which jobs are most in demand, the specific skills employers need, and the career directions that offer the highest potential for workers. For more information, visit burning-glass.com.

Job Postings Data

To support these analyses, Burning Glass mined its dataset of millions of job postings collected since 2010. Burning Glass collects postings from over 45,000 online job sites to develop a comprehensive, real-time portrait of labor market demand. We aggregate job postings, remove duplicates, and extract data from job posting text. This includes information on job title, employer, industry, region, and required experience, education, and skills. .

Job postings are useful for understanding trends in the labor market because they allow for a detailed, real-time look at the skills employers seek. In order to assess the representativeness of job postings data, Burning Glass conducts a number of analyses to compare the distribution of job postings to the distribution of official government and other third-party sources in the US. The primary source of government data on job postings in the US is the Job Openings and LAbor Turnover Survey (JOLTS) program conducted by the Bureau of Labor Statistics.

To understand the share of job openings captured by Burning Glass data, it is important to first note that Burning Glass and JOLTS collect data on job postings differently. Burning Glass data captures new postings: a posting appears in the data only on the first month it is found, and is considered a duplicate and removed in subsequent months. JOLTS data captures active postings: a posting appears in the data in all months that it is still actively posted, meaning the same posting can be counted in two or more consecutive months if it has not been filled.   T allow for apples-to-apples volume comparison in postings, the Burning Glass data needs to be inflated to account for active postings, not only new postings. The number of postings from Burning Glass can be inflated using the new jobs to active jobs ratio in Help Wanted OnLine™ (HWOL), a method used in Carnevale, Jayasundera and Repnikov (2014). Based on this calculation, the share of jobs online as captured by Burning Glass is roughly 85% of the jobs captured in JOLTS in 2016.

The labor market demand captured by Burning Glass data represents over 85% of the total labor demand. Jobs not posted online are usually in small businesses (the classic example being the "help wanted" sign in the restaurant window) and union hiring halls.

Measuring the demand for AI

In order to measure the demand by employers of AI skills, Burning Glass used its skills taxonomy of over 17,000 skills. The list of AI skills from Burning Glass data are shown in the table below, with associated skill clusters. While we considered some skills to be in the AI cluster specifically, for the purposes of this report all skills in the table below were considered AI skills A job postings was considered an AI job if it requested one or more of these skills.

Return to Economy - Jobs: US Labor Demand By Job Cluster, Skill Penetration, Regional Dynamics

Table A4.6. AI Skill Cluster

In addition, Burning Glass' taxonomy assigns skills to Skill Clusters. The following Skill

Skill	Skill Cluster	Skill	Skill Cluster	Skill	Skill Cluster
Artificial Intelligence	AI	Boosting (Machine Learning)	ML	Blue Prism	Robotics
Expert System	AI	Chi Square Automatic Interaction Detection (CHAID)	ML	Electromechanical Systems	Robotics
IBM Watson	AI	Classification Algorithms	ML	Motion Planning	Robotics
IPSoft Amelia	AI	Clustering Algorithms	ML	Motoman Robot Programming	Robotics
Ithink	AI	Decision Trees	ML	Robot Framework	Robotics
Virtual Agents	AI	Dimensionality Reduction	ML	Robotic Systems	Robotics
Autonomous Systems	AD	Google Cloud Machine Learning Platform	ML	Robot Operating System (ROS)	Robotics
Lidar	AD	Gradient boosting	ML	Robot Programming	Robotics
OpenCV	AD	H2O (software)	ML	Servo Drives / Motors	Robotics
Path Planning	AD	Libsvm	ML	Simultaneous Localization and Mapping (SLAM)	Robotics
Remote Sensing	AD	Machine Learning	ML	Computer Vision	Vision
ANTLR	NLP	Madlib	ML	Image Processing	Vision
Automatic Speech Recognition (ASR)	NLP	Mahout	ML	Image Recognition	Vision
Chatbot	NLP	Microsoft Cognitive Toolkit	ML	Machine Vision	Vision
Computational Linguistics	NLP	MLPACK (C++ library)	ML	Object Recognition	Vision
Distinguo	NLP	Mlpy	ML	Caffe Deep Learning Framework	NN
Latent Dirichlet Allocation	NLP	Random Forests	ML	Convolutional Neural Network (CNN)	NN
Latent Semantic Analysis	NLP	Recommender Systems	ML	Deep Learning	NN
Lexalytics	NLP	Scikit-learn	ML	Deeplearning4j	NN
Lexical Acquisition	NLP	Semi-Supervised Learning	ML	Keras	NN
Lexical Semantics	NLP	Supervised Learning (Machine Learning)	ML	Long Short-Term Memory (LSTM)	NN
Machine Translation (MT)	NLP	Support Vector Machines (SVM)	ML	MXNet	NN
Modular Audio Recognition Framework (MARF)	NLP	Semantic Driven Subtractive Clustering Method (SDSCM)	ML	Neural Networks	NN
MoSes	NLP	Torch (Machine Learning)	ML	Recurrent Neural Network (RNN)	NN
Natural Language Processing	NLP	Unsupervised Learning	ML	Pybrain	NN
Natural Language Toolkit (NLTK)	NLP	Vowpal	ML	TensorFlow	NN
Nearest Neighbor Algorithm	NLP	Xgboost	ML		
OpenNLP	NLP				
Sentiment Analysis / Opinion Mining	NLP				
Speech Recognition	NLP				
Text Mining	NLP				
Text to Speech (TTS)	NLP				
Tokenization	NLP				
Word2Vec	NLP				

Note: AD is Autonomous Driving, NLP is Natural Language Programming, AI is Artificial Intelligence, ML is Machine Learning, NN is Neural Networks, Vision is Visual Image Recognition.

Return to Economy - Jobs: [US Labor Demand By Job Cluster], [Skill Penetration], [Regional Dynamics]

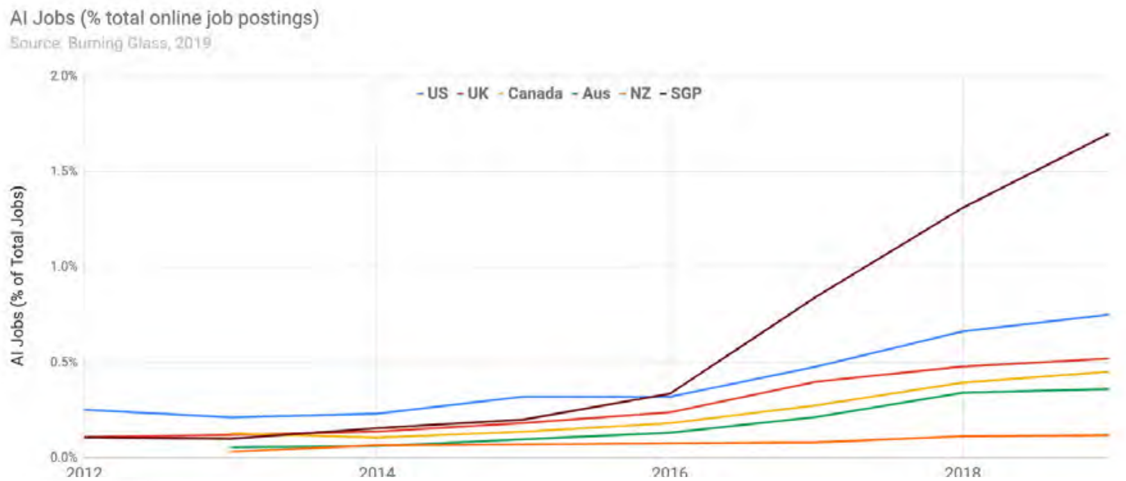

Fig. A4.8a.

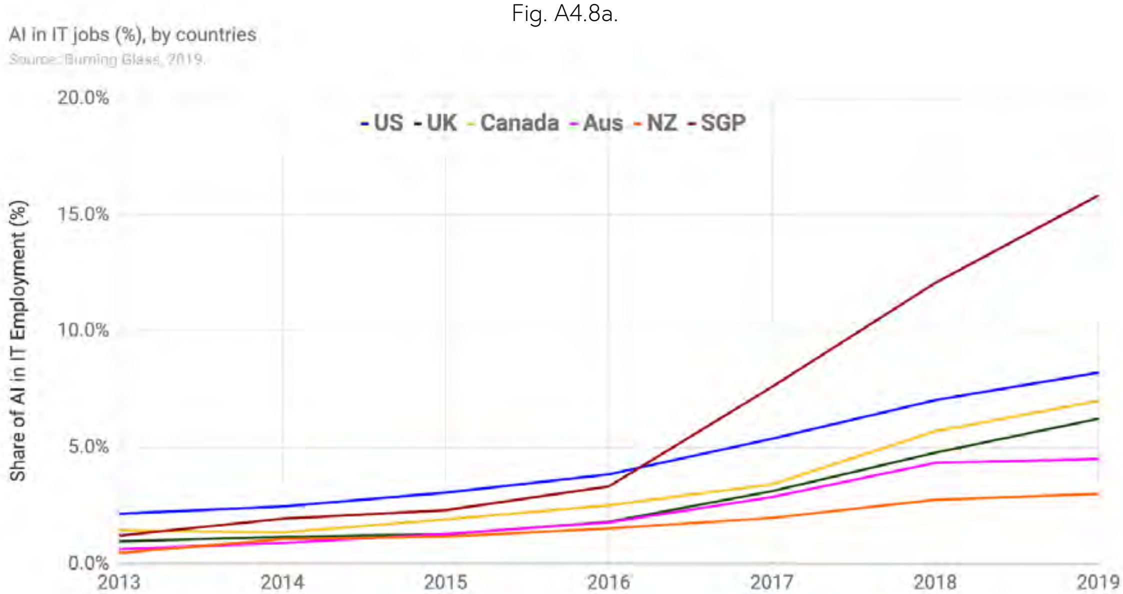

Fig. A4.8b.

Volume of AI Jobs by Skill Clusters

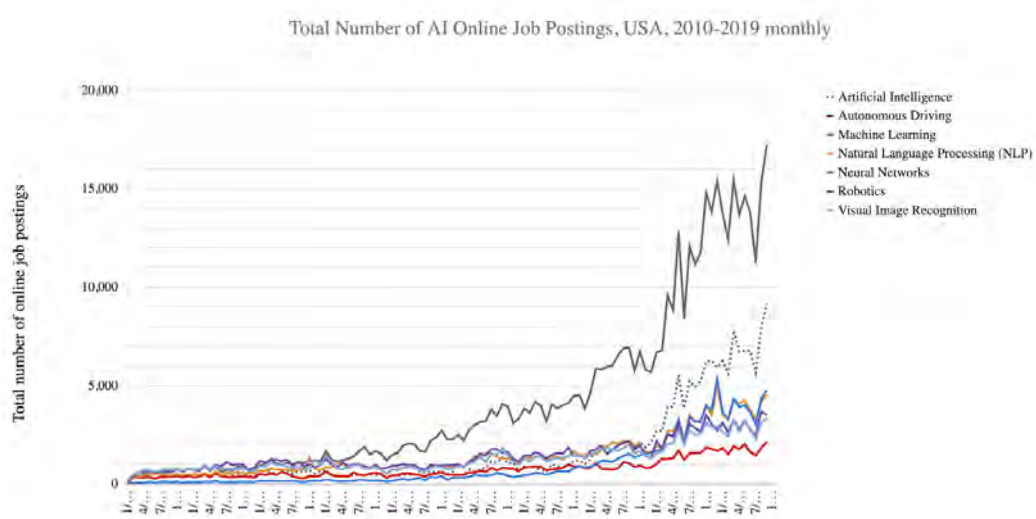

Fig. A4.9.

Return to Economy - Jobs: US Labor Demand By Job Cluster, Skill Penetration, Regional Dynamics

Unconditional Convergence in Jobs posted, absolute

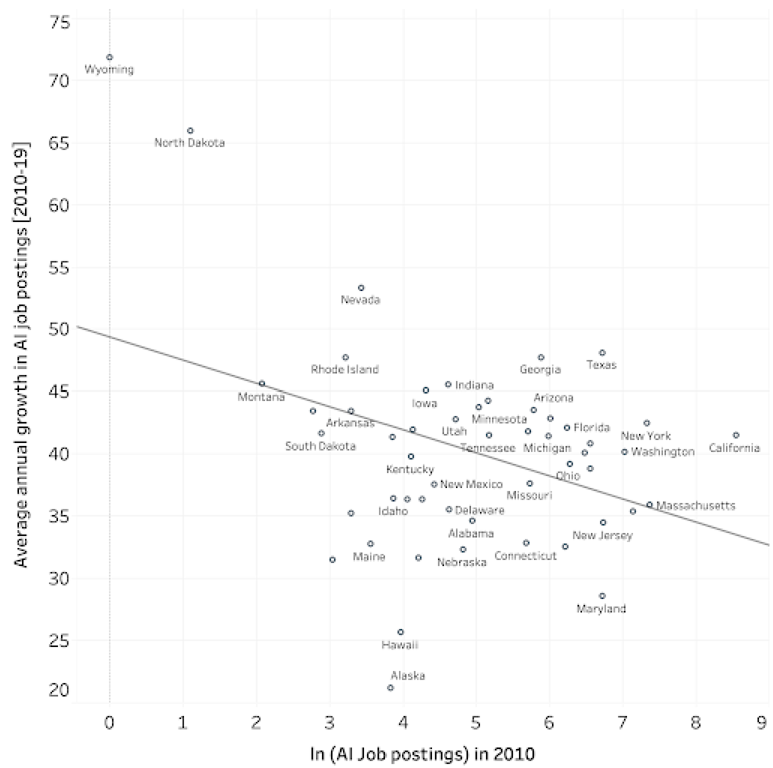

Fig. A4.10a.

Unconditional Divergence in Jobs posted, relative (as a % of total jobs posted)

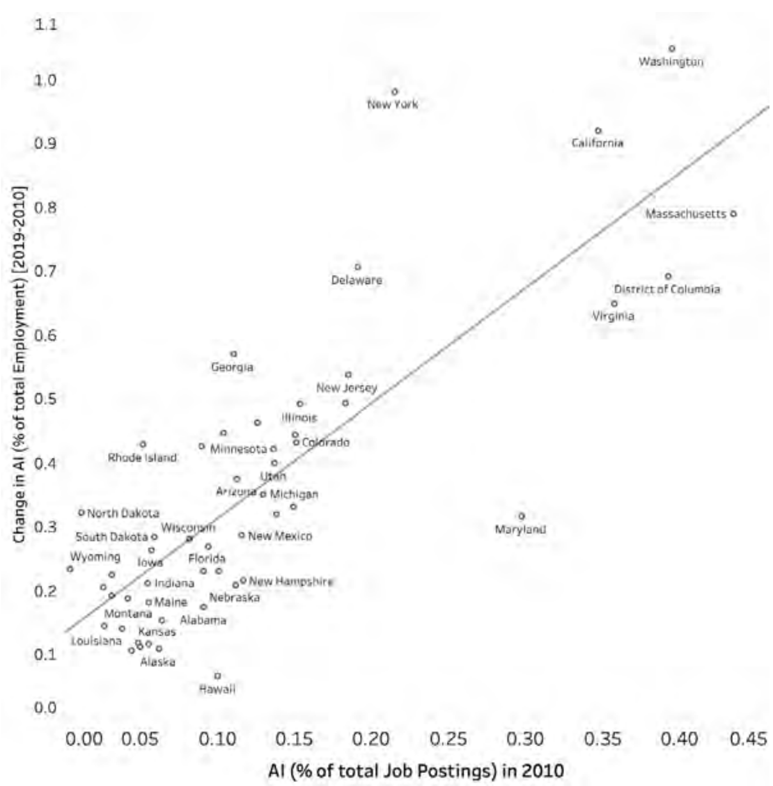

Fig. A4.10b.

Return to Economy - Jobs: [US Labor Demand By Job Cluster], [Skill Penetration], [Regional Dynamics]

AI Job Postings, State Analysis

Ranking of US States based on relative AI labor demand, 2018-19

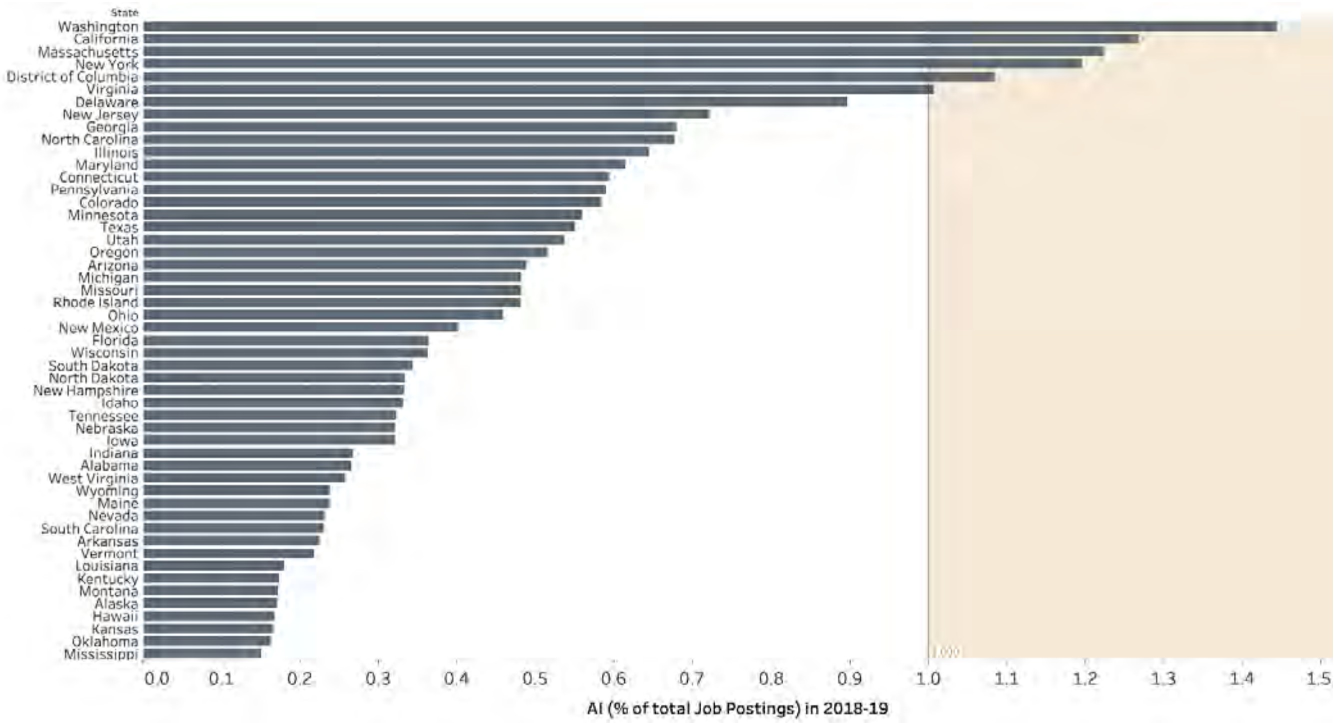

Fig. A4.11a.

Unconditional Divergence in Jobs posted, relative (as a % of total jobs posted)

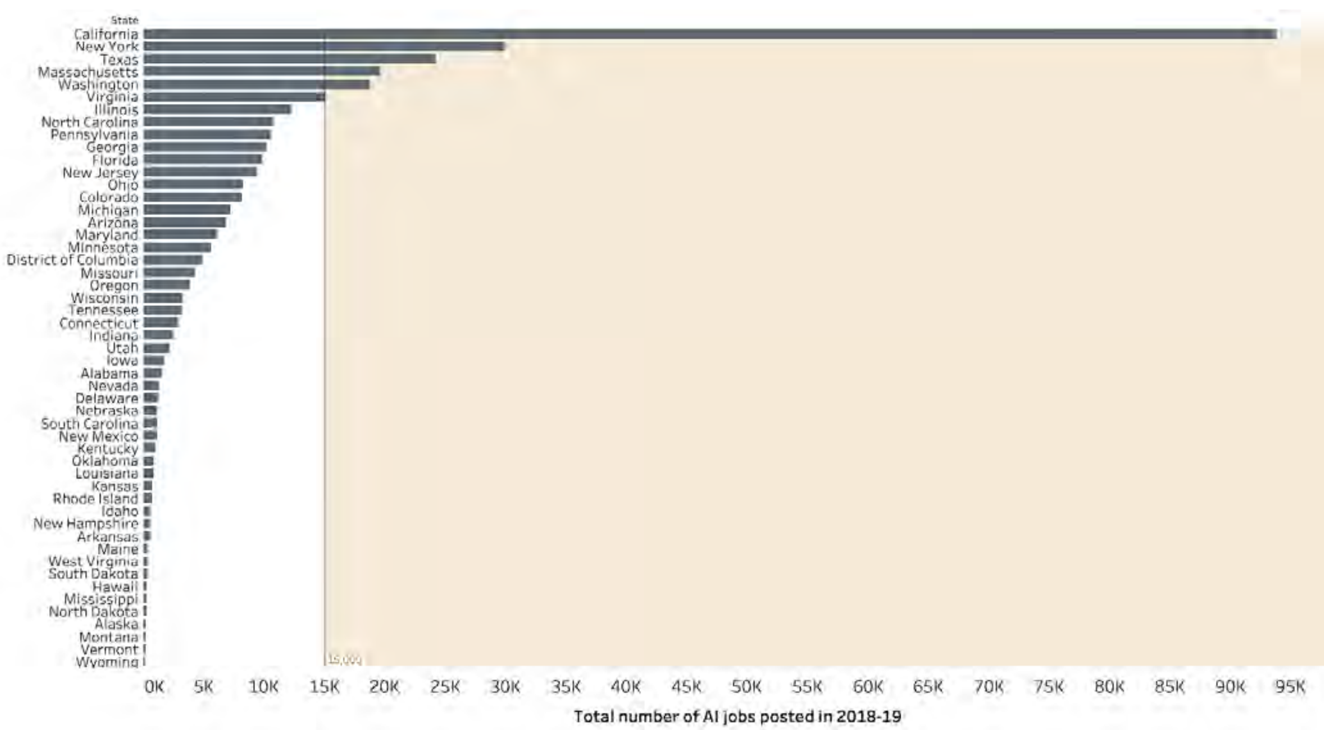

Fig. A4.11b.

Return to Economy - Jobs: US Labor Demand By Job Cluster, Skill Penetration, Regional Dynamics

Benchmarking states in absolute and relative growth in AI labor demand, 2010-19

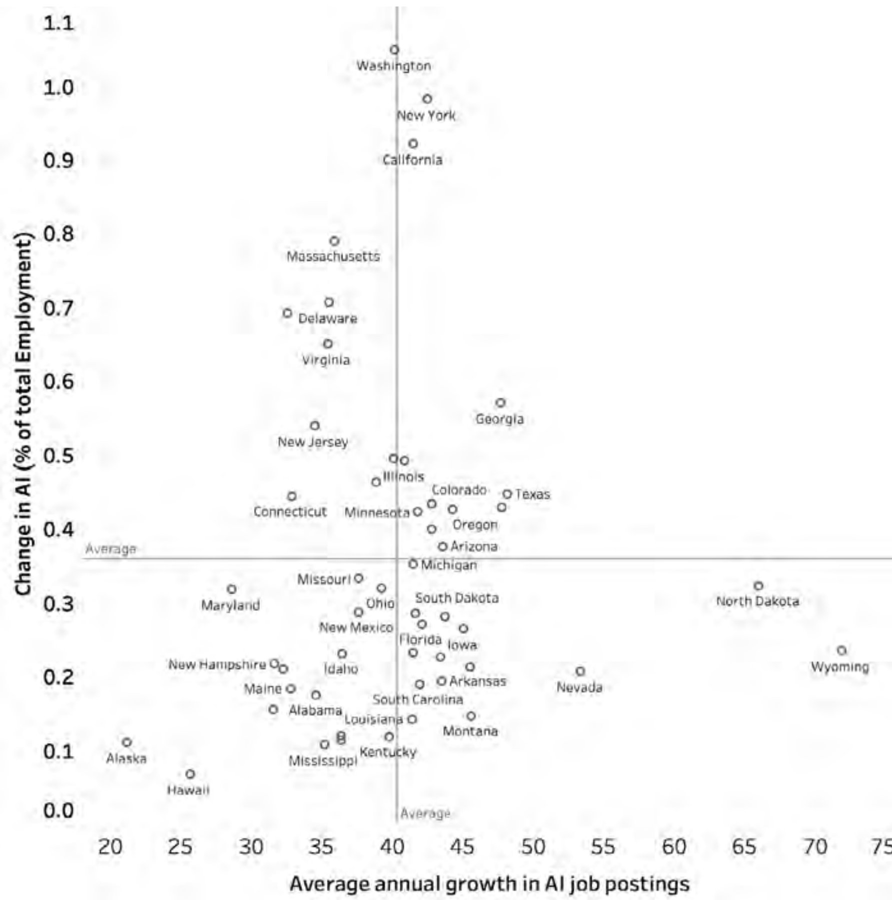

		Average annual growth in AI jobs postings, absolute	
		Lower than average	Higher than Average
Changein AI jobs postings, relative as a % of total jobs posted	Higher than Average	AI jobs growing in as part of the labor market but terms but did not not as much in absolute size	AI jobs growing both in absolute and relative importance of overall labor market
	Lower than average	AI jobs are not growing in either absolute size or relative importance	AI jobs growing in absolute terms but did not increase as much in relative importance of labor market

Fig. A4.12.

Return to Economy - Jobs: [US Labor Demand By Job Cluster], [Skill Penetration], [Regional Dynamics]

No convergence in AI labor demand across MSA's, absolute (2010-19)

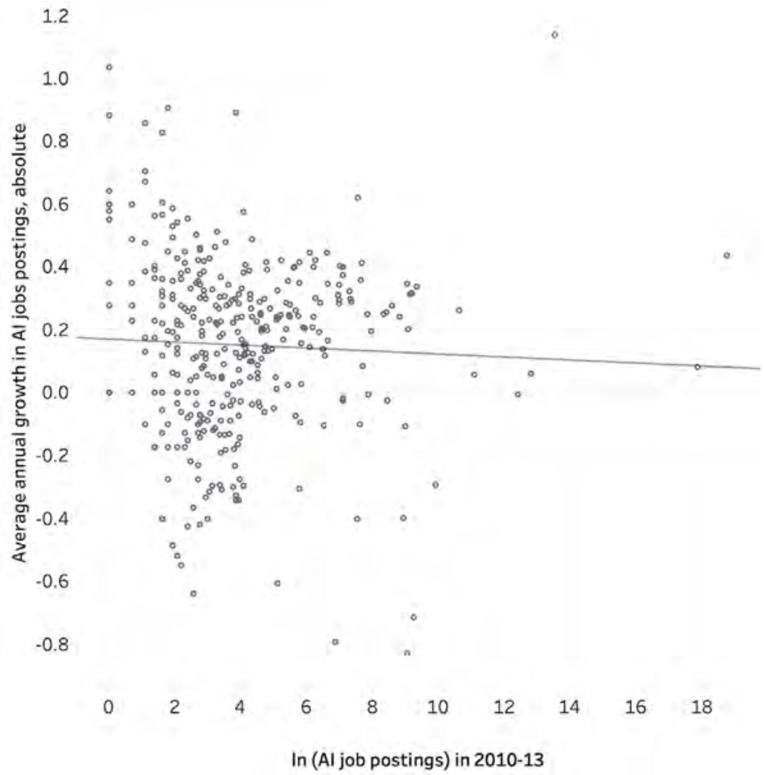

Fig. A4.13a.

No convergence in AI labor demand across MSA's, absolute (2010-19)

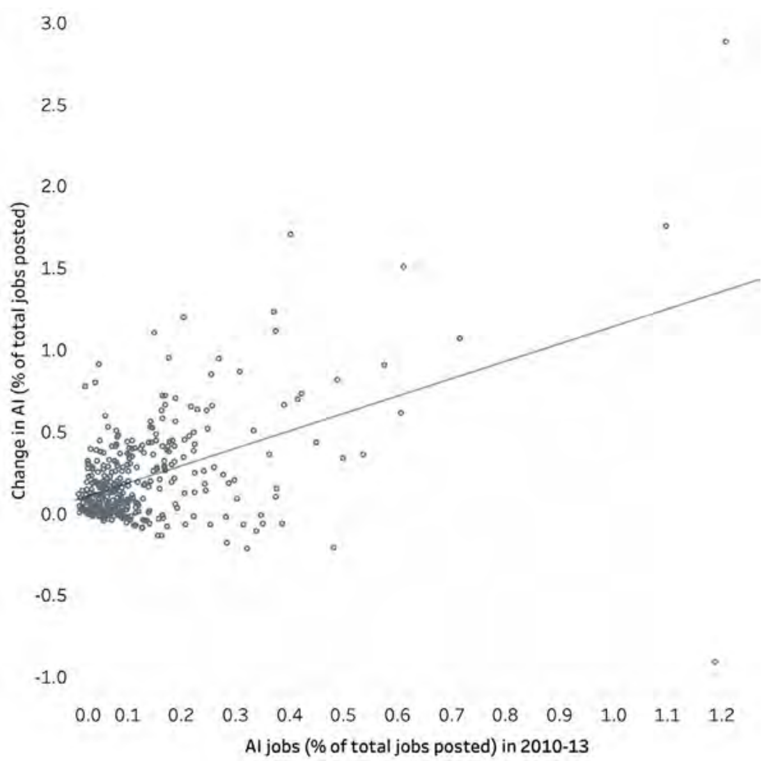

Fig. A4.13b.

Return to Economy - Jobs: Measurement Questions and Policy Implications

References for AI Measurement and Policy Implications

Brynjolfsson, Erik, Daniel Rock, Prasanna Tambe. "How Will Machine Learning Transform the Labor Market?". Governance in an Emerging New World, Hoover Institution, Spring Series, Issue 619, May 6, 2019.

Goldfarb, A., Taska, B., Teodoridis, F. "Machine Learning in Healthcare?". (2019)

Goldfarb, A., Taska, B., Teodoridis, F. "Could Machine Learning Be a General-Purpose Technology? Evidence from Online Job Postings" (2019)

Hershbein, Brad, and Lisa B. Kahn. "Do recessions accelerate routine-biased technological change? Evidence from vacancy postings." American Economic Review 108, no. 7 (2018): 1737-72.

Tambe, Prasanna. "Big data investment, skills, and firm value." Management Science 60, no. 6 (2014): 1452-1469.

Tambe, Prasanna, and Lorin M. Hitt. "Job hopping, information technology spillovers, and productivity growth." Management science 60, no. 2 (2013): 338-355.

Tambe, Prasanna, Lorin Hitt, Daniel Rock and Erik Brynjolfsson. "IT, AI and the Growth of Intangible Capital." (2019).

4.2 Investment Activity

Table A4.7. Summary Table on Private Investment Metrics

Metadata	Metric	Investment Amount	Number of Companies	Per Capita Investment Amount	M&A and IPO amount	% Focus areas
	Definition	Private Investment received by the AI startups who have received over $400,000 investment for the last ten years	Number of AI startups who have been founded within the last ten years	Private Investment received by AI startups divided by population of their headquartered country	M&A or IPO deal amount for AI startups who have received over $400,000 investment for the last ten years	Percentage of focus areas within all AI startups who have received funding within the last 1 year
	Source	CapIQ, Crunchbase, Quid	CapIQ, Crunchbase, Quid	CapIQ, Crunchbase, World Bank, Quid	CapIQ, Crunchbase, Quid	CapIQ, Crunchbase, Quid
	Unit of Measurement	investment event	a company	population of startup headquarter country	M&A, IPO event	number of companies
Coverage	# of countries, states, regions - data available	84	123	84	84	80
	# of sectors available	46	46	46	46	36
	# of years available	10	10	10	10	1
Methodology and Evaluation	Methodology	Used boolean search query in Quid's NLP technology to search for all global AI startups who have received over $400,000 funding for the last 10 years and created panel data showing investment time trend with headquartered countries and clusters	Used boolean search query in Quid's NLP technology to search for all global AI startups who have been founded within the last ten years and created panel data showing founding year time trend with headquarter countries	Used Quid's NLP technology to search for all global AI startups and thier financial activities data as long as they are disclosed in CapIQ and Crunchbase. Used World Bank's 2018 population data to divide the annual investment amount per country by population.	Used boolean search query in Quid's NLP technology to search for all global AI startups who have received over $400,000 funding for the last 10 years and created panel data showing all financial event time trend	Used boolean search query in Quid's NLP technology to search for all global AI startups who have received funding within the last one year, and created network map based on neuroscience technology segmented into different focus areas based on NLP algorithm

Private Investment Activity

Source

Quid is a big data analytics platform that inspires full picture thinking by drawing connections across massive amounts of unstructured data. The software applies advanced natural language processing technology, semantic analysis, and artificial intelligence algorithms to reveal patterns in large, unstructured datasets, and generate visualizations to allow users to gain actionable insights. Quid uses Boolean query to search for focus areas, topics, and keywords within the archived news and blogs, companies and patents database, as well as any custom uploaded datasets. This can filter out the search by published date time frame, source regions, source categories, or industry categories on the news; and by regions, investment amount, operating status, organization type (private/public), and founding year within the companies database. Quid then visualizes these data points based on the semantic similarity.

Search, Data Sources, and Scope

Here 1.8M public and private company profiles from multiple data sources are indexed in order to search across company descriptions, while filtering and including metadata ranging from investment information to firmographic information such as founded year, HQ location, and more. Company information is updated on a weekly basis. Trends are based on reading any text to identify key words, phrases, people, companies, and institutions; then compare different words from each document (news article, company descriptions…etc) to develop links between these words based on similar language. This process is repeated at an immense scale which produces a network that shows how similar all documents are.

Data

Organizational data from CapIQ and Crunchbase are embedded together. These companies include all types of companies (private, public, operating, operating as subsidiary, out of business) in the world; The investment data include private investment, M&A, public offering, minority stake made by PE/VCs, corporate venture arms, governments, institutions both in and out of the US. Some data is simply unreachable when the investors are undisclosed, or the funding amounts by investors are undisclosed. We also embed firmographic information such as founded year, HQ location.

We embed CapIQ data as a default, and add in data from Crunchbase for the ones that are not captured in CapIQ. This way we not only have comprehensive and accurate data on all global organizations, but also capture early-stage startups and funding events data. Company information is uploaded on a weekly basis.

Data Sourced From:

Quid indexes 1.8M public and private company profiles from multiple data sources allowing you to search the company descriptions, while filtering by included metadata ranging from investment information to firmographic information such as founded year, HQ location, and more. Company information is updated on a weekly basis. 7,000 companies can be analyzed within one network. Company information is updated on a weekly basis.

Boolean Search

"artificial intelligence" OR "AI" OR "machine learning" OR "deep learning"

Return to Economy - Investment Activity: Startup Activity

Global AI companies invested within last one year (06/27/2018~ 06/27/2019)

<Companies>

Chart 4.2.1, 4.2.2., 4.2.3, 4.2.4, 4.2.5, 4.2.7:

1. Global AI & ML companies who have been invested over 400k for the last 10 years (01/01/2019 to 11/04/2019) – 7000 companies out of 7.5k companies have been selected through Quid's relevance algorithm

Chart 4.2.6.

2. Global AI & ML companies who have been invested (private, IPO, M&A) from 06/27/2018 to 06/27/2019

Visualization in Quid software:

We use Boolean query to search for focus areas, topics, and keywords within the archived company database, within their business descriptions and websites. We can filter out the search results by HQ regions, investment amount, operating status, organization type (private/public), founding year. Quid then visualizes these companies. If there are more than 7000 companies from the search result, Quid selects 7000 most relevant companies for visualization based on the language algorithm.

Target Event Definitions:

•Private investments: A Private Placement is a private sale of newly issued securities (equity or debt) by a company to a selected investor or a selected group of investors. The stakes that buyers take in private placements are often minority stakes (under 50%), although it is possible to take control of a company through a private placement as well, in which case the private placement would be a majority stake investment.

•Minority investment: These refer to Minority Stake Acquisitions in Quid, which are where the buyer acquires less than 50% of the existing ownership stake in entities, asset product and business divisions.

•M&A: These refer to events where a buyer acquires more than 50% of the existing ownership stake in entities, asset product and business divisions.

Network Methodology

The algorithm uses textual similarities to identify documents that are similar to each other. It then creates a network based on these similarities so that the user can visualize these similarities as a network of clusters. To do this, Quid leverages proprietary NLP algorithms and unsupervised machine learning to automate the topical generation.

The dots, or nodes, represents individual companies (or articles), and the links represent semantic similarity between two nodes, with the clusters (groupings) differentiated by colors representing the topics.

Readers can refer to Olivia Fischer, Jenny Wang (2019). *Innovation and Convergence/Divergence: Searching the Technology Landscape* for more details on the methodology.

Return to Economy - Investment Activity: Startup Activity

How to Read a Quid Map in Companies

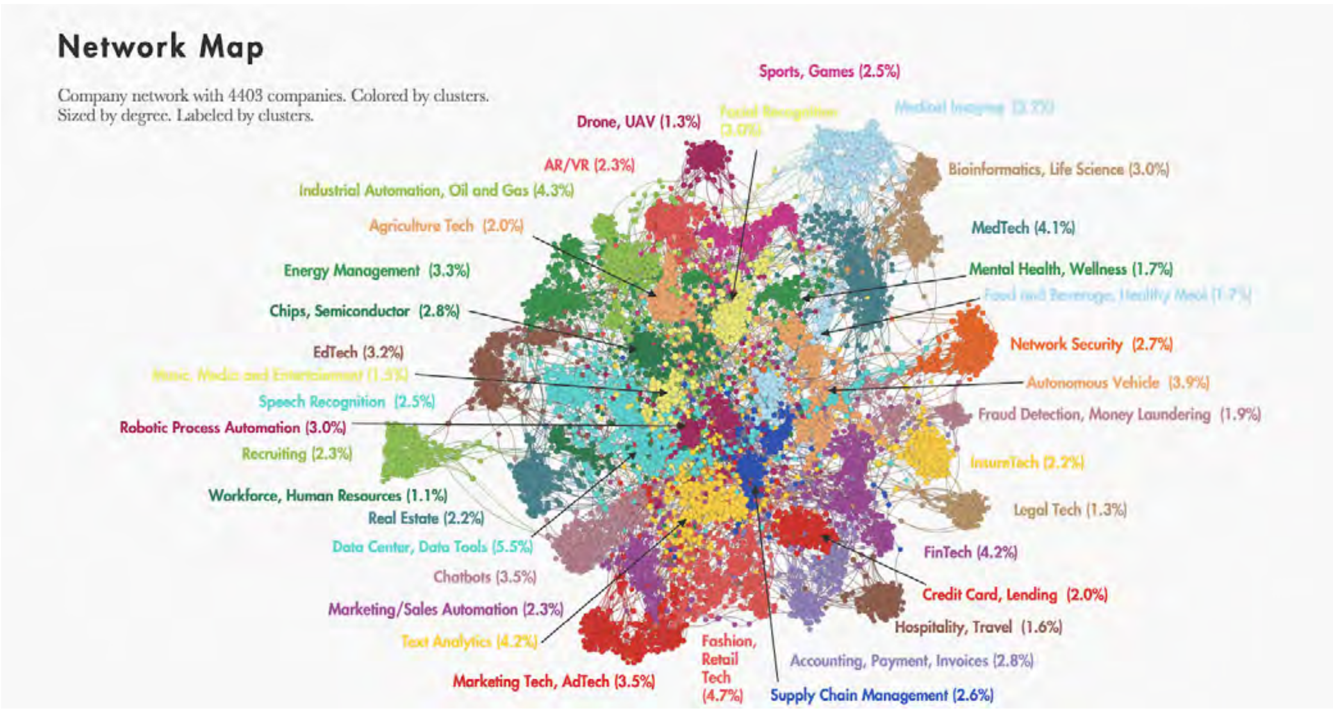

Fig. A4.14.

*Reading map visualization: Each node represents a company. Links connect companies sharing similar languages in their business descriptions and websites. Clusters form when many nodes share strong similarity, revealing focus areas.

When considering the network, cardinal directions (e.g. North, South, East, West) does not matter – what does matter is proximity. Two clusters which are close together (e.g. Medical Imaging, MedTech and Bioinformation, Life Science) share more common language than the ones that are far away (e.g. Fashion Retail, Tech). Centrality also matters – those clusters that are more central to network (e.g. Robotic Process Automation) are more core to the market versus those on the periphery (e.g. Recruiting).

List of European Countries

United Kingdom, France, Germany, Finland, Switzerland, Sweden, Spain, Belgium, Ireland, the Netherlands, Luxembourg, Norway, Denmark, Portugal, Austria, Italy, Poland, Iceland, Czech Republic, Serbia, Estonia, Romania, Slovenia, Latvia, Croatia, Greece, Bulgaria, Lithuania, Malta

Geography of AI Startup Activity

US, Europe, China, and India, State Level Startup Activity

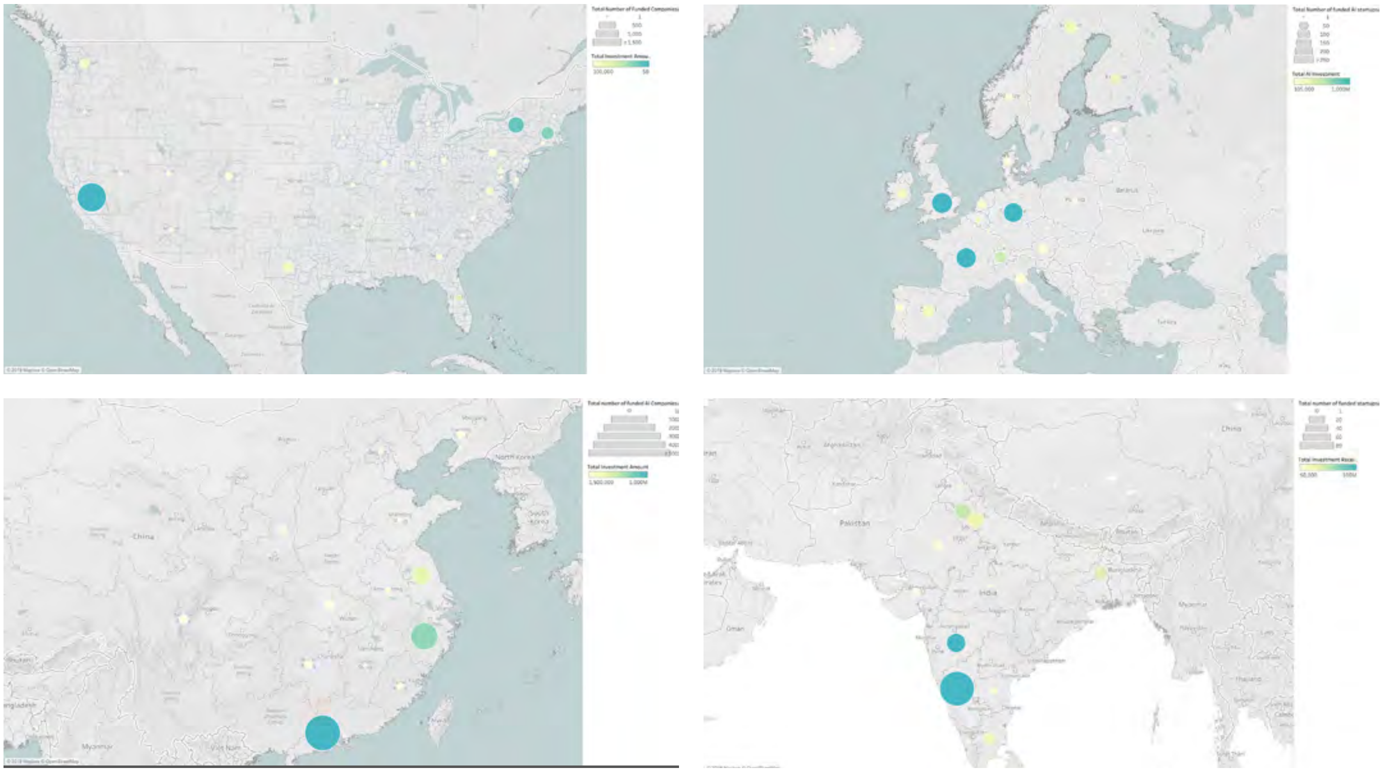

US, Europe, China, and India, City Level Startup Activity

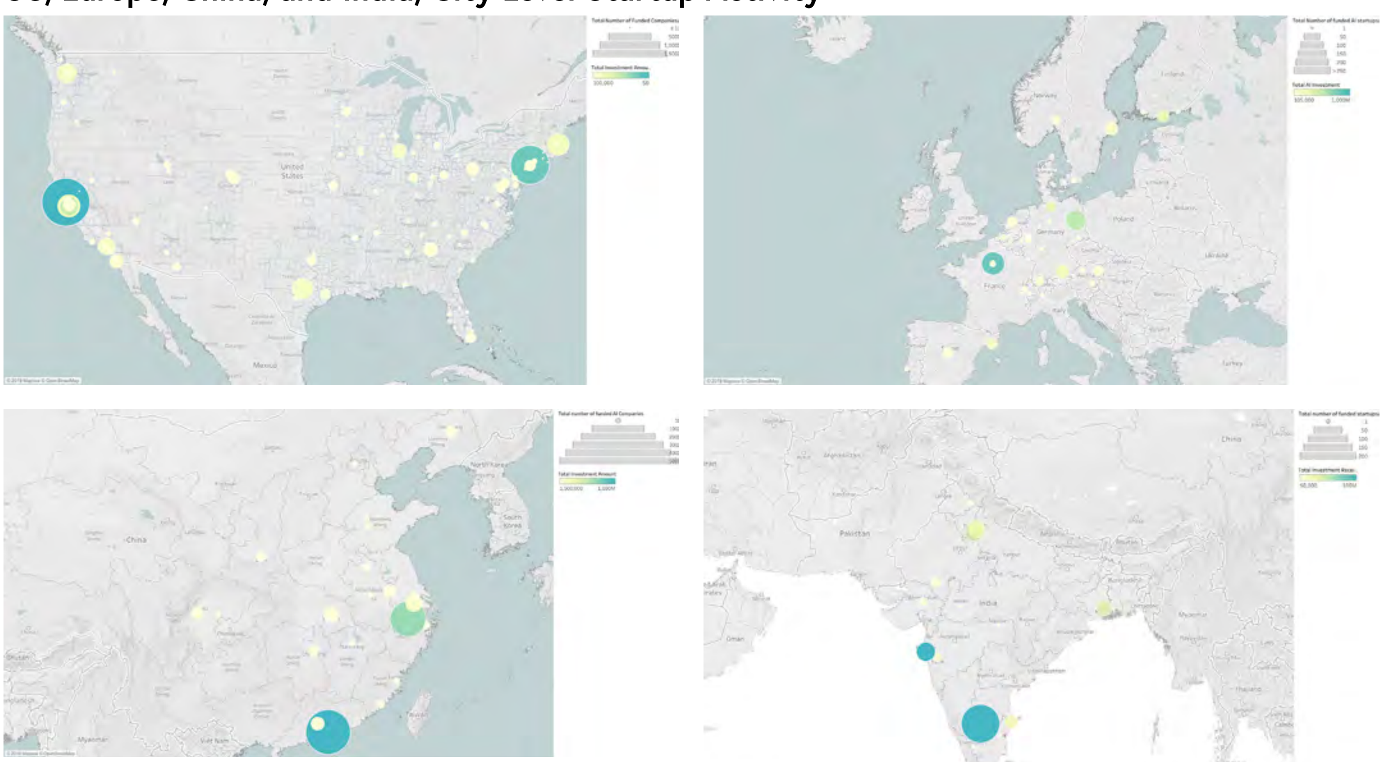

Fig. A4.15.
Investment deals in billion dollars

Bloomberg Government: Public Investment

Source

Bloomberg Government (BGOV) is a subscription-based market intelligence service designed to make US government budget and contracting data more accessible to business development and government affairs professionals. BGOV's proprietary tools ingest and organize semi-structured government data sets and documents, enabling users to track and forecast investment in key markets.

Methodology

The BGOV data included in this section was drawn from three original sources:

•The Defense Department's FY 2020 Research, Development, Test & Evaluation (RDT&E) Budget Request, which is available at https://comptroller.defense.gov/Budget-Materials/. BGOV built a custom dashboard and data visualization using the software tool, Tableau, to organize all 6,664 budget line items included in the FY 2020 RDT&E budget request and make them text-searchable. For the section "Department of Defense (DOD) Budget," BGOV used a set of more than a dozen AI-specific keywords to identify 346 unique budget activities related to artificial intelligence and machine learning worth a combined $4.0 billion in FY 2020.

•The Federal Procurement Data System-Next Generation (FPDS-NG), which is available at https://www.fpds.gov/fpdsng_cms/index.php/en/. BGOV's Contracts Intelligence Tool ingests on a twice-daily basis all contract spending data published to FPDS-NG, and structures the data to ensure a consistent picture of government spending over time. For the section "US Government Contract Spending," BGOV analysts used FPDS-NG data, organized by the Contracts Intelligence Tool, to build a model of government spending on artificial intelligence-related contracts in the fiscal years 2000 through 2019. BGOV's model used a combination of government-defined Produce Service Codes and more than 100 AI-related keywords and acronyms to identify AI-related contract spending.

•The Congressional Record, which is available at https://www.congress.gov/congressional-record. BGOV maintains a repository of congressional documents, including bills, amendments, bill summaries, Congressional Budget Office assessments, reports published by congressional committees, Congressional Research Service (CRS), and others. For the section "US Government Perception," BGOV analysts identified all legislation (passed or introduced), congressional committee reports, and CRS reports that referenced one or more of a dozen AI-specific keywords. Results are organized by two-year congressional session.

4.3 Corporate Activity
Industry Adoption: Mckinsey Global Survey

Return to Economy - Corporate Activity: Industry Adoption

Source

This survey was written, filled, and analyzed by McKinsey & Company (McKinsey). You can find additional results from the Global AI Survey here.

Methodology

The survey was conducted online and was in the field from March 26, 2019 to April 5, 2019. It garnered responses from 2,360 participants who represent the full ranges of regions, industries, company sizes, functional specialties, and tenures within McKinsey's Online Executive Panel. All survey participants are members of the online panel, a group of more than 30,000 registered users of McKinsey.com who opted in to participate in proprietary McKinsey research and represent a worldwide sample of executives, managers, and employees at all levels of tenure. 115 countries are represented in the survey sample; to adjust for differences in response rates, the data are weighted by the contribution of each respondent's nation to global GDP.

Notes

Survey respondents are limited by their perception of their organizations' AI adoption.

Robotic Installations

Return to Economy - Corporate Activity: Robotic Installations

Source

Data was pulled directly from the International Federation of Robotics' (IFR) 2014, 2015, and 2017 World Robotics Reports. See links to the reports below. Learn more about IFR.

Methodology

The data displayed is the number of industrial robots installed by country. Industrial robots are defined by the ISO 8373:2012 standard. See more information on IFR's methodology.

Nuance

• It is unclear how to identify what percentage of robot units run software that would be classified as "AI" and it is unclear to what extent AI development contributes to industrial robot usage.
• This metric was called robot imports in the 2017 AI Index report

Return to Education - [Coursera]

Coursera

http://www.coursera.org/gsi (page 46)

Global AI Skill Index

Building the GSI involves data from several components: **Coursera's Skills Graph, Skills Benchmarking, Competency Growth,** and **Trending Skills**. Below we provide more insight into how we calculate each piece. This is our first look into the global skills landscape using our unique data, and we are constantly evolving our methodology to maximize its usefulness for our learners and customers.

This GSI report focuses on the 60 countries with the most learners on the Coursera platform and 10 of the largest industries that have both seen major shifts in their skill landscapes and are primed for future workforce development. The 60 countries account for 97% of learners on the Coursera platform, and for about 80% of the world's population and 95% of global GDP (based on 2017 World Bank Data).

Skills Graph

Coursera's Skills Graph maps the connections among skills, content, careers, and learners on the Coursera platform. For GSI, in particular, we leverage the following edges of the Skills Graph:

is_parent_of	This edge describes the connections among skills. It generates a skills taxonomy where broad, higher-level skills are parents of more granular, lower-level skills.
is_taught_by	This edge maps skills to the Coursera courses that teach them.
is_assessed_by	This edge maps skills to the graded items that assess them Graded items on Coursera can be of several types: multiple choice quizzes, peer review assignments like essays and projects, or programming assignments.
is_outcome_of	This edge connects competencies to learners who have demonstrated them by passing relevant graded items. We measure this using Coursera's Skills Benchmarking methodology, described further below.

Identifying the set of skills and relationships among skills, Is_parent_of

We assemble a vast skills taxonomy of over 40,000 skills in the subject areas of Business, Technology, and Data Science through a combination of open-source taxonomies like Wikipedia and crowdsourcing from Coursera educators and learners. Guided by open-source data combined with knowledge from industry experts, we assemble a structured taxonomy that connects Coursera domains to the set of skills within them, ranging from competencies down to very specific skills ('Level 1+ skills'). For the GSI, we focus on measuring performance at the competency level.

To illustrate the mapping among domains, competencies, and skills, we have a sample snapshot of a subsection of Coursera's Skills Taxonomy below:

Return to Education - Coursera

Coursera AI Taxonomy

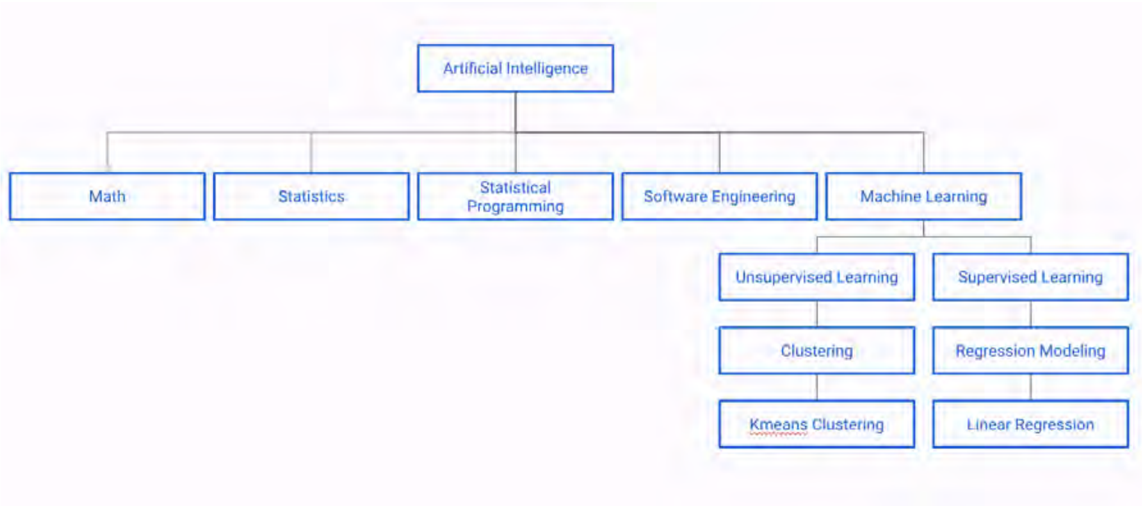

The full set of competencies for which we measure learner proficiency in the GSI, grouped by domain, are listed in the following table:

Business	Technology	Data Science
Accounting	Computer Networking	Data Management
Communication	Databases	Data Visualization
Finance	Human Computer Interaction	Machine Learning
Management	Operating Systems	Math
Marketing	Security Engineering	Statistical Programming
Sales	Software Engineering	Statistics

Mapping skills to courses and assessments, is_taught_by and is_assessed_by

The skills in Coursera's Skills Taxonomy are mapped to the courses that teach them using a machine learning model trained on a data set of university instructor and learner-labeled skill-to-course mappings. Features of the model include occurrence counts (e.g., in the lecture transcripts, assignments, and course descriptions), NLP embeddings, and learner feedback.

With over 1,500 courses in Business, Technology, and Data Science from top ranked university and industry partners around the world, our catalog spans the wide variety of skills that are relevant to competencies in the GSI. For each skill-course pair, this machine learning model outputs a score that captures how likely it is that the skill is taught in the course. To define the set of skill-tocourse tags that power GSI, we tune a cutoff threshold based on expert feedback from our content strategy team.

When a skill within a competency is tagged to a course, we extract the graded items in that course as being relevant for assessing a given competency. These competency-to-assessment mappings were reviewed with industry experts to ascertain their fidelity and adjusted as needed. This final set serves as the pool we use to measure individual learners' skill proficiencies.

Return to Education - Coursera

Skills Benchmarking

Measuring individual learners' skill proficiencies, is_outcome_of

With the set of assessments for each competency defined, we consider grades for all learners taking relevant assessments and train machine learning models to simultaneously estimate individual learners' skill proficiencies (i.e., how proficient each learner is in each competency) and individual assessment difficulties (i.e., how challenging each assessment is). Each domain and competency has its own model to estimate these parameters, resulting in 21 separate models. This methodology allows us to measure learner skill proficiencies adjusting for item difficulty. This is essential because the Coursera platform contains a wide variety of courses ranging from the introductory college level to the advanced graduate level. Adjusting for item difficulty ensures we neither penalize learners for taking difficult courses nor over-reward learners for strong performance in easy courses.

Measuring individual learners' skill proficiencies, is_outcome_of

Because learners attempt various numbers of graded items at various levels of difficulty, we also assess the precision with which we are measuring skill proficiency for each learner through the calculation of standard errors. We use the skill proficiency estimated above as a measure of the relative ability of each learner within a domain or competency. Aggregating across learners in an entity reveals the average proficiency in that group.

We calculate the weighted average of skill proficiency estimates, where the weights are the inverse of the standard error for that learner. To avoid undue influence of any individual learner, weights are trimmed to be at or below the median value of the overall distribution of weights within each domain/competency. This weighted average for each domain and competency is the GSI estimate of an entity's skill proficiency. We then compare groups to each other via a percentile ranking of all GSI estimates. Performance bands for a group's skill proficiency are computed by segmenting skill proficiencies into quartiles:

Cutting-Edge for 76th percentile or above, rank #1–15
Competitive for 51st to 75th percentile, rank #16–30
Emerging for 26th to 50th percentile, rank #31–45
Lagging for 25th percentile or below, rank #46–60

Our 38 million registered learners span the globe and myriad industries, and the GSI reflects the average skill proficiency of learners in each entity on the Coursera platform, accounting for the precision with which we measure each individual's skill proficiency. Note that the GSI estimate may not reflect the average skill proficiency of all entity members because Coursera learners are not necessarily representative of a country or industry.

Competency Popularity by Enrollments

We measure competency growth by enrollments on the Coursera platform in courses teaching related skills between 2017 and 2018. Competency Popularity provides high-level insight into which direction learners are increasingly investing their time for skill development, and provides an additional signal as to which skills are trending within the labor market.

Return to Education - Coursera

Trending Skills

We measure trending skills within each domain (Business, Technology, and Data Science) on a quarterly basis, incorporating several measures of internal and external demand for each skill into a single, weighted index:

Learner Enrollments The average enrollments per course by learners in content tagged to a particular skill.
Search Trends The number of searches on Coursera by logged in learners for a particular skill.
Google Trends The Google Trends Index for a particular skill, which provides a measure of search activity on Google pertaining to specific keywords and topics.
Labor Market Value The estimated dollar value of a skill based on the relative frequency in job postings, career salary, and general return to skills from the literature,3 based on US data only.

For a given domain we calculate the above fields for each skill. To ensure all metrics are on the same scale, we first compute the z-score of each attribute within its domain and then generate a weighted average of the four z-scores above to calculate the index value for a skill in a particular quarter.

Tracking the value of this index over time allows us to see what is increasing and decreasing in popularity.

We can calculate this index for particular demographic groups as well by restricting the set of learners used in it. As an example, we calculate the trending skills for each GSI region subgroup by finding the consumer enrollments, enterprise enrollments, and search impressions on the Coursera platform by learners within each GSI region, weighing the z-scores together to compute the index.

Cutting Edge

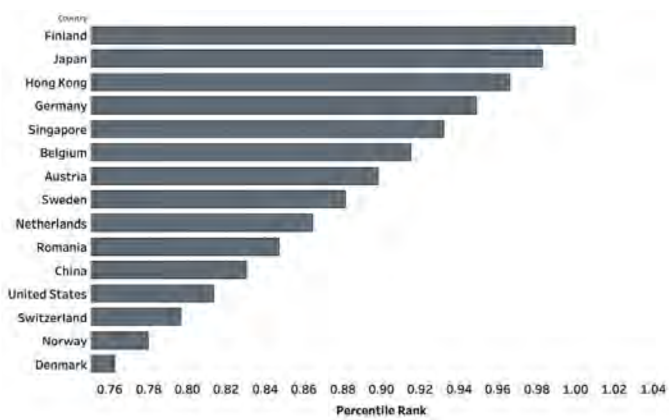

Competitive

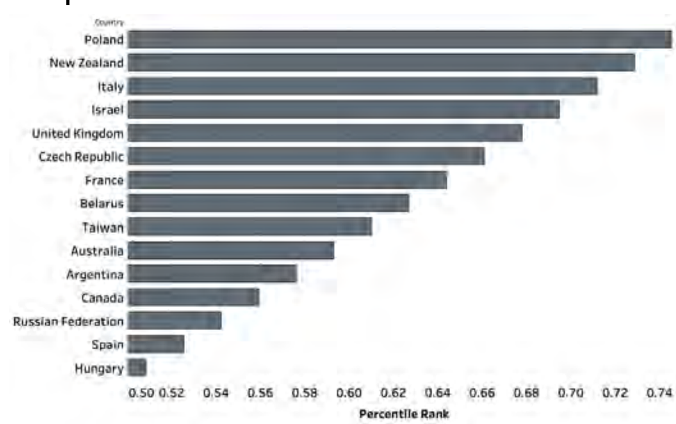

Emerging

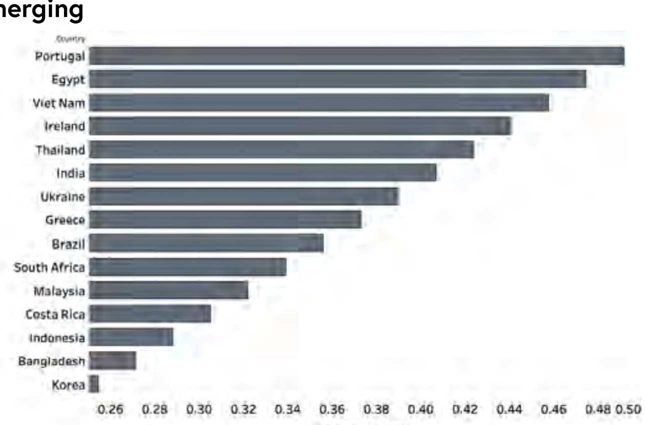

Lagging

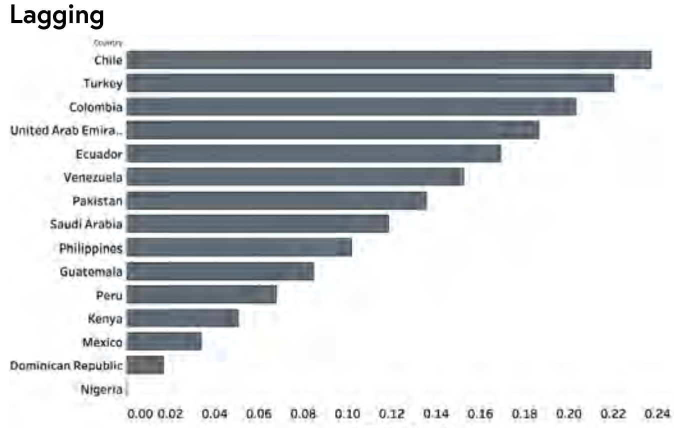

Fig. A5.1a, b, c, & d.

Return to Education - Coursera

Trending Skills

Competency Id

Avg. Percentile Rank: 0.000 – 1.000

Country	artificial-intelligence	fields-of-mathematics	machine-learning	software-engineering	statistical-programmi..	statistics
Argentina	0.576	0.356	0.475	0.559	0.576	0.610
Australia	0.593	0.644	0.508	0.271	0.458	0.508
Austria	0.898	0.593	0.898	0.898	0.542	0.627
Bangladesh	0.271	0.169	0.169	0.119	0.068	0.254
Belarus	0.627	0.322	0.966	0.966	0.983	0.966
Belgium	0.915	0.814	0.949	0.797	0.729	0.797
Brazil	0.356	0.085	0.373	0.983	0.661	0.407
Canada	0.559	0.424	0.593	0.949	0.814	0.441
Chile	0.237	0.475	0.305	0.305	0.593	0.322
China	0.831	0.949	0.407	0.000	0.000	0.644
Colombia	0.203	0.068	0.237	0.458	0.339	0.136
Costa Rica	0.305	0.373	0.288	0.593	0.271	0.271
Czech Republic	0.661	0.881	0.797	0.475	0.780	0.898
Denmark	0.763	0.525	0.814	0.915	0.508	0.695
Dominican Republic	0.017	0.136	0.017	0.847	0.356	0.153
Ecuador	0.169	0.339	0.051	0.678	0.746	0.288
Egypt	0.475	0.627	0.559	0.390	0.051	0.305
Finland	1.000	0.746	0.932	0.746	0.949	0.949
France	0.644	0.763	0.881	0.508	0.881	0.864
Germany	0.949	0.712	0.915	0.695	0.932	0.831
Greece	0.373	0.559	0.525	0.661	0.220	0.525
Guatemala	0.085	0.153	0.153	0.525	0.390	0.186
Hong Kong	0.966	0.983	0.746	0.169	0.153	0.780
Hungary	0.508	0.508	0.695	0.932	0.288	0.458
India	0.407	0.780	0.220	0.034	0.119	0.390
Indonesia	0.288	0.203	0.186	0.153	0.559	0.169
Ireland	0.441	0.441	0.441	0.492	0.186	0.373
Israel	0.695	0.610	0.780	0.339	0.644	0.915
Italy	0.712	0.542	0.661	0.763	0.797	0.712
Japan	0.983	1.000	0.424	0.441	0.169	0.847
Kenya	0.051	0.017	0.068	0.051	0.136	0.034
Korea	0.254	0.915	0.322	0.068	0.034	0.339
Malaysia	0.322	0.288	0.339	0.186	0.627	0.356
Mexico	0.034	0.237	0.136	0.424	0.322	0.085
Netherlands	0.864	0.797	0.831	0.610	0.898	0.678
New Zealand	0.729	0.271	0.542	0.780	0.475	0.763
Nigeria	0.000	0.000	0.000	0.102	0.102	0.000
Norway	0.780	0.661	0.847	0.542	0.831	0.661
Pakistan	0.136	0.034	0.085	0.136	0.305	0.068
Peru	0.068	0.390	0.102	0.644	0.610	0.102
Philippines	0.102	0.305	0.119	0.085	0.203	0.119
Poland	0.746	0.678	0.712	0.576	0.712	0.729
Portugal	0.492	0.220	0.576	0.729	0.678	0.542
Romania	0.847	0.186	0.678	0.712	0.254	0.932
Russian Federation	0.542	0.831	1.000	1.000	1.000	1.000
Saudi Arabia	0.119	0.119	0.254	0.407	0.407	0.017
Singapore	0.932	0.898	0.729	0.203	0.424	0.814
South Africa	0.339	0.102	0.390	0.373	0.373	0.203
Spain	0.525	0.864	0.610	0.814	0.441	0.593
Sweden	0.881	0.576	0.864	0.627	0.763	0.881
Switzerland	0.797	0.966	0.983	0.831	0.847	0.983
Taiwan	0.610	0.695	0.492	0.017	0.017	0.424
Thailand	0.424	0.458	0.458	0.254	0.915	0.492
Turkey	0.220	0.051	0.203	0.288	0.237	0.220
Ukraine	0.390	0.254	0.627	0.864	0.966	0.746
United Arab Emirates	0.186	0.492	0.356	0.322	0.864	0.237
United Kingdom	0.678	0.729	0.644	0.356	0.695	0.576
United States	0.814	0.847	0.763	0.220	0.492	0.559
Venezuela	0.153	0.932	0.034	0.881	0.085	0.051
Viet Nam	0.458	0.407	0.271	0.237	0.525	0.475

Fig. A5.2.

US AI and ML course enrollment

Source

Course enrollment data was collected directly by AI Index from each university. Total student population was collected from school archives (typically housed on Office of the Registrar sites). The following universities are included in our analysis:

University of California-Berkeley, Stanford University, University of Illinois at Urbana–Champaign, University of Washington-Seattle, Carnegie Mellon University

Methodology

We requested enrollment in introductory AI and introductory ML courses over time at many leading computer science universities in the US Several schools participated. Enrollment counts were not included in our analysis if the school did not include sufficient historical data, or if the data was overly nuanced.

Some schools provided enrollment by semester, and some provided it by year. In order to compare schools to each other, we collapsed semester enrollment data to show full academic years. Additionally, some schools had multiple courses that were considered "introductory" while others just had one. When appropriate and relevant, multiple courses were combined to show one "introductory AI" trend line.

For enrollment as a percent of the undergraduate population, each year's AI / ML enrollment was divided by the undergraduate population for that same year. This is a calculated field intended to show trends in enrollment on an even playing field across schools.

Nuance

- Nearly every school noted that enrollment, particularly in recent years, is a function of supply, rather than student demand. Our data shows that the number of students that were successfully enrolled in a course, and does not account for waitlists or other demand metrics.
- Courses are generally open to undergraduates only, and can typically be taken by majors and non-majors. Some courses have changed their names over time, we show course names as of 2017 below. We also list any additional details / nuances that school administrators were able to provide on the enrollment data.
- Any nuance that was not mentioned by school administrators is not captured below.

AI courses:
Berkeley CS 188
Stanford CS 221
UIUC CS 440
- UIUC representatives attribute growth to larger classrooms / additional sections to meet some of the excess student demand.

UW CSE
- 415 / 416 (non-majors) & CSE 473(Majors)
- CSE 416 is new as of AY 2017 and accounts for some growth in AI course enrollment in 2017
- 415, 473, 573 were used for 2018

ML courses:
Berkeley CS 189
- Representatives at Berkeley speculate that growth is due to a combination of novelty, subject interest, and growth in majors that allow Intro ML as a way to fill requirements.
- 189/289 were included in 2018.

Stanford CS 229
- The reason for the drop in ML enrollment in 2016 from 2015 is due to two factors. First, in 2015, CS229 was offered twice, but then in 2016 and 2017 it was only offered once. So that might explain part of the drop from 2015 to 2016. The other (bigger) factor was that in 2016 the course was mostly taught by an instructor other than Andrew Ng (although Andrew was still listed as an instructor, but only gave a few of the lectures) who's primary appointment was not in CS. So this was really an exogenous event. In 2017, even though the class was only offered once, there was pent up demand from the year before and the instructors were Andrew Ng and another very popular CS instructor (Dan Boneh), so enrollment bounced back. In 2018, CS229 was again offered twice a year.

UIUC CS 446
UW CSE 446

International Course Enrollment

Source

Course enrollment data was collected directly from each university. The following universities are included in our analysis:

Tsinghua University (China), National Institute of Astrophysics, Optics and Electronics (Mexico), University of British Columbia (Canada), University of Toronto (Canada), University of Edinburgh (Scotland), Pontificia Universidad Católica de Chile (Chile), Universidad Tecnica Federico Santa Maria (UTFSM) (Chile), Universidad Nacional Andrés Bello (UNAB) (Chile), High School of Economics (HSE) (Russia), University of Melbourne (Australia), Universidade Federal do Rio Grande do Sul (UFRGS) (Brazil), Peking (China

Methodology

— See methodology in US Course Enrollment Appendix.

Nuance

- Nearly every school noted that enrollment, particularly in recent years, is a function of supply, rather than student demand. Our data shows the number of students that were successfully enrolled in a course, and does not account for waitlists or other demand metrics.
- Unlike the US schools studied, international schools significantly varied on whether courses were only open to undergraduates or not.
- Visual one shows growth in AI and ML courses combined. Visual two shows just AI course enrollment. We did this in order to show like for like data on each graph. In some cases, we had access to additional data on a school but did not show it because we wanted to have parallel information across schools. Additional data in located in the underlying data link in the top right corner.
- Some courses have changed their names over time, we show course names as of 2017 below. We also list any additional details / nuances that school administrators were able to provide on the enrollment data. Any nuance that was not mentioned by school administrators is not captured below.

INAOE — Courses: C141 (AI) and C142 (computational learning)
Notes: INAOE AI / ML enrollment is greatly affected by the number of students accepted into the INAOE graduate program as a whole. INAOE representatives say that there is a decreasing number of INAOE students, thus affecting AI / ML course enrollment.
USTC — Courses: USTC listed several introductory AI / ML courses across various departments including the Department of Computer Science and Technology, The Department of Automation, the Department of Information Science and Technology and the Department of Data Science.
University of Edinburgh — Courses: Intro applied ML (undergraduate and graduate students) and Informatics 2D — Reasoning and Agents (undergraduate only)
SJTU — Course: CS410 (undergraduate intro to AI)
PUC — Course: Intro to AI
Prior to 2017, the course was only taught once a semester. The large demand in 2017, relative to 2018, is due to the transition from one course to two courses.
Tsinghua — Courses: AI (60240052 & 40250182) and ML (00240332 & 70240403 & 80245013 & 80250943)
Open to undergraduates and graduate students
Toronto — Courses: AI (CSC384) and ML:(CSC411)
2016 was the first year that a summer AI course was open. decision to open two semesters of ML in 2015 — due to increased demand
University of Melbourne (Australia) — Two undergraduate subjects (one on machine learning, one on more general AI) was extracted.
UFRGS (Brazil) — UFRGS offers the courses up to two times a year. Typically, UFRGS has about 100 PhD students enrolled and 200 MSc students.
HSE (Russia) — Introduction courses in AI and ML
The key aim of HSE's Data Culture project is to provide all undergrads insight into the latest technologies used in data analysis. This way, students in management will be able to set clear tasks for analysts, while analysts, in turn, will be fast and efficient in building their models, and applied specialists will rely on the most cutting-edge data tools. Project Levels: Elementary, Basic, Advanced, Professional, Expert.
Peking (China) — Introduction to AI course.

Faculty Diversity

Source

Faculty diversity was collected manually via AI department websites on September 21st, 2018. Schools selected are leading computer science universities with easily accessible AI faculty rosters.

Methodology

In order to get the gender diversity breakdown of AI faculty, professor names were collected on school websites, and then genders were assigned (see first nuance below) using both names and pictures. Please see below for more specific details on each school:

UC Berkeley — See faculty link
Includes *Assistant Professors, Primary, Secondary Faculty*
Stanford University — See faculty link
Includes *Faculty and Research Scientists* and *Affiliated Faculty*
University of Illinois at Urbana–Champaign — See faculty link
Includes *CS Faculty* and *Affiliate Faculty*
Carnegie Mellon University — See faculty link
Includes all faculty listed
University College, London — See faculty link
Includes all faculty under the *People* link
University of Oxford — See faculty link
Includes *Faculty* section only
ETH Zurich — See faculty link
Includes only those with "*Dr.*" in their title
Georgia Tech — See faculty link
Includes all faculty under the *Machine Learning* link
NUS Singapore — See faculty link
Includes AI Faculty in the *Computing* section
University of Toronto — See faculty link
Includes Faculty in the *Machine Learning Department*
IIT Madras — See faculty link
Includes *Current Faculty* in the *Department of Computer Science and Engineering*

Nuance

- We assigned genders using professor names and pictures. In doing so, the AI index may have misgendered someone. We regret that we could not include non-binary gender categories into this analysis. We hope the metric is still useful in showing a broad view of gender representation in AI academia today, and look forward into expanding into other types of gender diversity in future reports.
- School data was pulled September 21st, 2018. School faculty could be altered by the time the 2018 AI Index report is published.
- Data is pulled from schools' AI faculty rosters and does not account for visiting professors or professors housed in other departments. Similarly, it will count a professor that is listed as an active member of AI faculty, even if that professor belongs to a different department.
- Gender representation in academia does not imply representation in industry (in fact, the proportion of women working in AI in industry may be lower).
- The metric provides a snapshot of representation today, and does not account for improvements over time, see below for a statement from a Stanford AI faculty member, Dr. James A. Landay:

  "*We are very focused on hiring more diverse faculty. Most of the women on that list have been hired in just the last 2—3 years, so we have been making progress*"
  - Dr. Landay, Stanford University

CRA Taulbee Survey

Source

Computing Research Association (CRA) is a 200+ North American organizations active organization in computing research: academic departments of computer science and computer engineering; laboratories and centers in industry, government, and academia; and affiliated professional societies (AAAI, ACM, CACS/AIC, IEEE Computer Society, SIAM USENIX). CRA's mission is to enhance innovation by joining with industry, government and academia to strengthen research and advanced education in computing. Learn more about CRA here.

Methodology

CRA Taulbee survey gathers survey data during the fall reaching out to over 200 Ph.D.-granting departments. Details about the Taulbee Survey can be found here. Taulbee doesn't directly survey the students. The department identifies each new PhD's area of specialization as well as their type of employment. Data is collected in September - January of each academic year for PhDs awarded in the previous academic year. Results are published in May after data collection closes. So the 2018 data points we provided were newly available last month and 2019 data will be available in May 2020.

The CRA Taulbee Survey goes only to doctoral departments of computer science and computer engineering. Historically, (a) Taulbee covers 1/4 to 1/3 of total BS CS recipients in the US, (b) the percent of women earning bachelor's is lower in the Taulbee schools than overall, and (c) Taulbee tracks the trends in overall CS production.

Nuances

- Of particular interest in PhD job market trends are the metrics on AI PhD area of specialization. The categorization of specialty areas changed in 2008 and was clarified in 2016. From 2004-7, AI and Robotics were grouped; 2008-present AI is separate; 2016 clarified to respondents that AI included ML.
- Notes about the trends in new tenure-track hires (overall, and particularly at AAU schools): In the 2018 Taulbee, for the first time, we asked how many new hires had come from the following sources: new PhD, postdoc, industry, and other academic. 29% of new assistant professors came from another academic institution.
- Some may have been teaching or research faculty previously rather than tenure-track, but there is probably some movement between institutions, meaning the total number hired overstates the total who are actually new.

Return to Education - PhD Hires

Undergraduate CS Enrollment

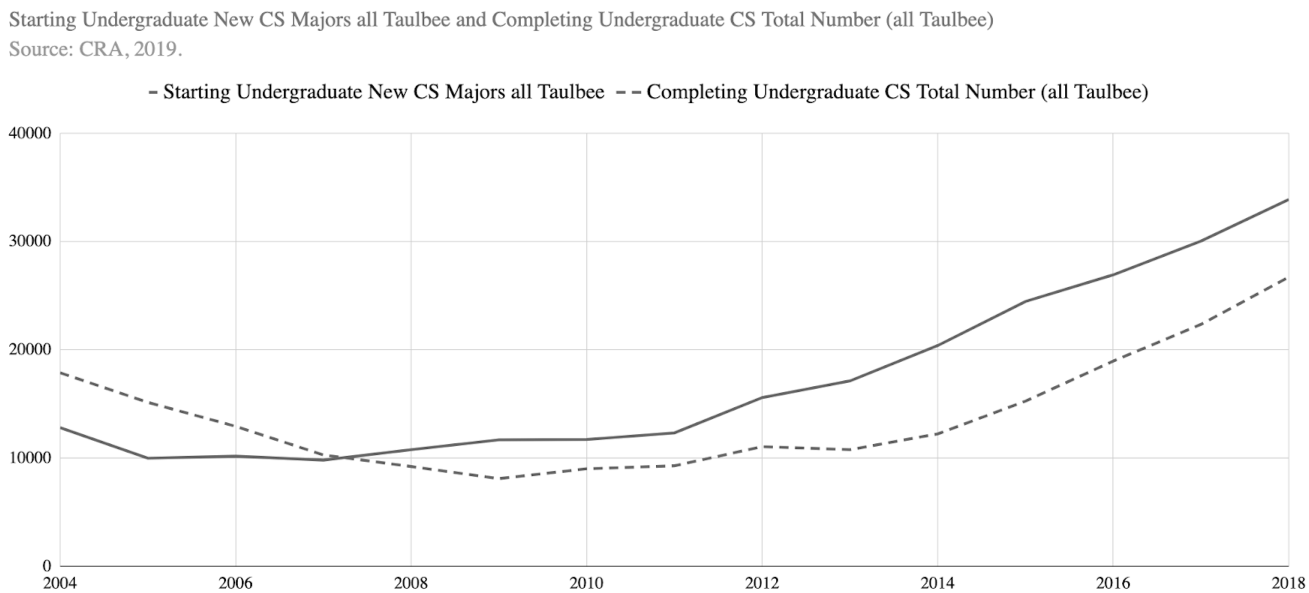

Fig. A5.3a.

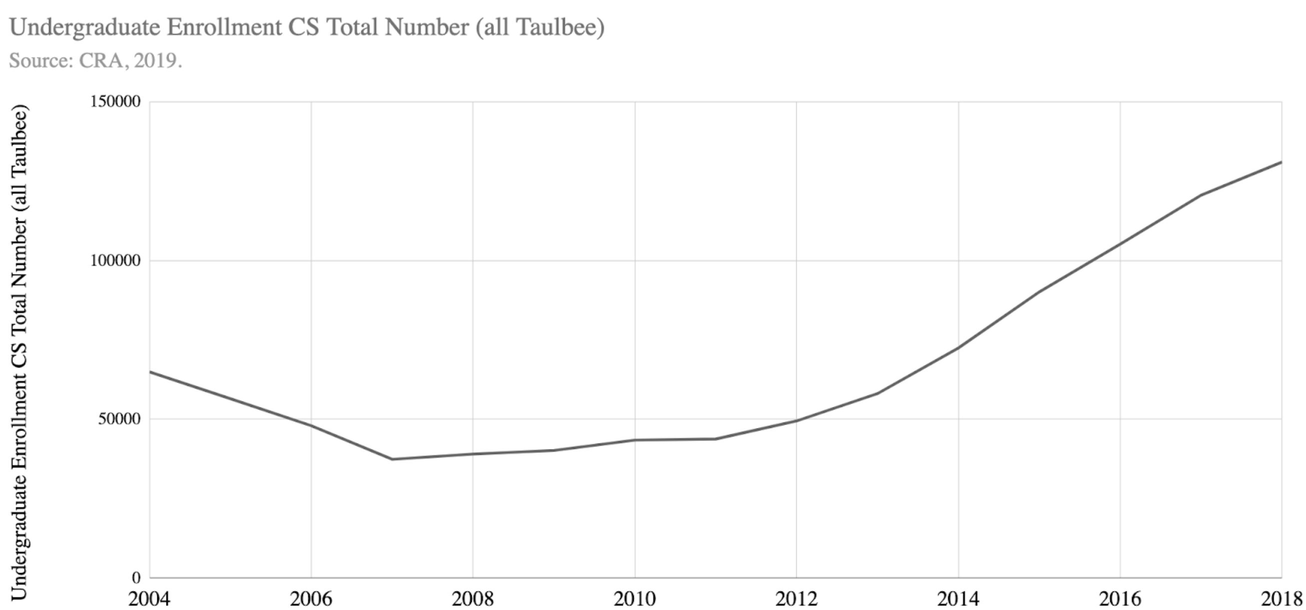

Fig. A5.3b.

Joint Research Center - EU Academic Offering

Source

The Joint Research Centre (JRC) is the European Commission's science and knowledge service. The JRC employs scientists to carry out research in order to provide independent scientific advice and support to EU policy. Learn more about JRC here.

Methodology

In this analysis, all the universities across Europe having a website have been considered (as listed by the Webometrics initiative) and, by applying text mining and machine learning techniques, the content related to study programmes addressing specified AI domains has been extracted. The aim was manifold: to collect independently a first set of results, to have therefore a suitable term of comparison when considering third party sources, and to be able to measure strengths and weaknesses of a (semi)automatic classification system for programmes' content in view of a systematisation of the exercise. The identification of programmes related to AI, HPC and CS from web pages has several challenges: (a) inconsistency in terminology used by universities to refer to study programmes (e.g. a "course" may refer both to a part of a study programme, or the whole programme); (b) troublesome identification of individual programmes in the entire webpage (header, footer, menu items), especially in webpages showing lists with the whole education offer.

Additionally, only English language content has been selected, due to limited resources to undergo a multilingual approach in data harvesting and text mining (mainly related to the amount of data to treat). The basic assumption, tested on randomly selected pages, is that the majority of master programmes are announced in English, while it is not the rule with undergraduate studies. Under these assumptions, the final product was a list of universities potentially focusing on AI, HPC and CS by announcing their bachelor and master studies. However, the identification of individual study programmes did not provide trustworthy data. As a consequence, another source to study education offer has been investigated.

In order to rely on a validated source and have access to more detailed information, StudyPortals data on bachelor and master studies has been collected. Worldwide, StudyPortals covers over 170,000 programmes at 3,050 educational institutes across 110 countries, out of which over 50,000 correspond to programmes taught in European universities. Programme information is collected by StudyPortals from institutions' websites; their database is kept updated, with new programmes added at least once a year. The consideration of this source increases the precision and the coverage of academic programmes by EU universities with respect to what offered by the the approach followed in the previous exploratory phase (in more than 90% of countries, the exploratory approach based on text mining Universities' websites resulted in lower university coverage than that provided by the selected source).

Return to Autonomous Systems: Autonomous Vehicles

6.1 Autonomous Vehicles

Six Levels of Autonomy by SAE

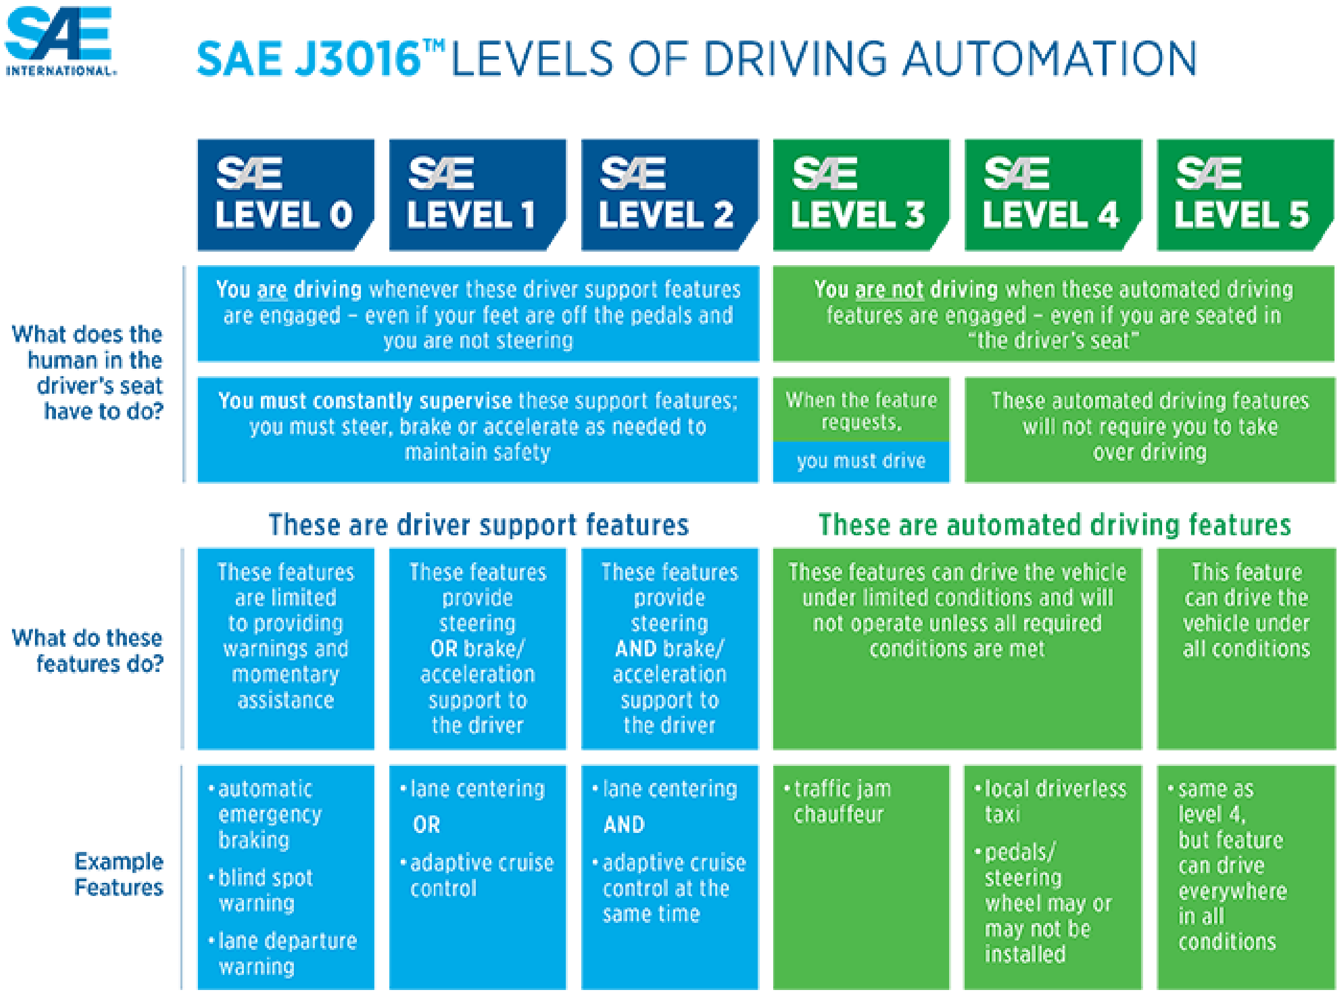

Fig. A6.1.

Note on Collision Reports

There is evidence that the AV drivers are learning to drop the vehicle out of the autonomous mode just before a crash, which then causes the crash to be classified as "conventional," not autonomous. If one look at Waymo 2017 the rate appears almost down to the "human" level. That is because Waymo only had ONE crash in 2017 that was classified as "autonomous." But actually they had three crashes, but two others were coded as "conventional." Here is the description of one of these "conventional" accidents so you can see what might be terms Waymo "gaming" the reporting system:

SECTION 5 — ACCIDENT DETAILS - DESCRIPTION

☐ Autonomous Mode ☑ Conventional Mode

A WAYMO LEXUS-MODEL AUTONOMOUS VEHICLE ("WAYMO AV") MADE CONTACT WITH A CURB WHILE IN MANUAL MODE ON MIDDLEFIELD ROAD AT OREGON EXPRESSWAY IN PALO ALTO, CA. THE WAYMO AV WAS TRAVELING EASTBOUND IN AUTONOMOUS MODE IN THE RIGHTMOST LANE OF MIDDLEFIELD ROAD. AS THE VEHICLE CROSSED OREGON EXPRESSWAY, THE WAYMO AV AUTONOMOUS SYSTEM DETECTED THE VEHICLE IN THE LEFT ADJACENT LANE BEGIN TO DRIFT TO THE RIGHT, TOWARD THE WAYMO AV. THE WAYMO AV NUDGED TO THE RIGHTMOST SIDE OF ITS LANE. AS THE LEFT ADJACENT VEHICLE CONTINUED TO DRIFT TOWARDS THE WAYMO AV, THE WAYMO AV TEST DRIVER TOOK MANUAL CONTROL. THE WAYMO AV'S FRONT PASSENGER-SIDE TIRE THEN MADE CONTACT WITH THE RIGHT CURB, CAUSING IT TO DEFLATE. THE OTHER VEHICLE THEN STRAIGHTENED ITS TRAJECTORY IN ITS LANE AND CONTINUED ON. THERE WERE NO INJURIES REPORTED.

Had these two crashes been coded correctly (in my opinion) as "autonomous," then Waymo's 2017 rate would have been very similar to their other years.

6.2 Autonomous Weapons

Source

tockholm International Peace Research Institute (SIPRI) is an international institute based in Sweden, dedicated to research into conflict, armaments, arms control and disarmament. Learn more about SIPRI here.

Methodologies, and Nuances:

However, there are a number of caveats we would like you to consider if you would like to use it:

1) This is not a dataset listing Lethal Autonomous Weapon Systems (LAWS), but a dataset intended to map out the *development of autonomy in military systems*. Many of the systems that are included in the dataset are not weapon systems but unarmed military systems that feature some notable autonomous capabilities.

2) The dataset is neither truly global nor comprehensive. For obvious practical reasons covering all countries and all types of weapon systems is not feasible. Some countries are also not transparent about their weapon development and acquisition programmes, which means that there is no way for us to guarantee that it is a representative dataset.

3) The data is a few years old. It has not been updated since 2017. Please make that clear if you intend to use the data in some way.

4) The dataset is not meant to measure the 'level of autonomy' of weapon systems. The dataset explore autonomy by functions. the scores by functions are not meant to be added together to create a total score of autonomy for systems as a whole. (Within each binary function, there are very different levels of autonomy and different autonomous functions do not necessarily have the same weight).

5) Beware of comparisons. comparing categories of systems, countries or applications can be tricky. The world is producing more UAVs than AUVs, so if one sees more autonomous UAVs than AUVs that does not necessarily mean that UAVs are more autonomous.

6) Reliability of information: we used a colour code to indicate how confident we were with the sources we used.

Table A6.1.

Color	Meaning
Orange	The sources are very bare, with little detail; the systems can be not autonomous (most likely) or their autonomy is not described, but it is difficult to assess as there is so little information available.
Blue	The information is not very reliable or realistic and comes across as propaganda or marketing
Grey	Competing information from different sources
Yellow	The system has been assessed with 1 or 2 descriptive sources, but there are not enough independent sources available (e.g. Non-copied press releases)
Green	Good, 3 different sources presenting information on the system

Return to Autonomous Systems - Autonomous Weapons

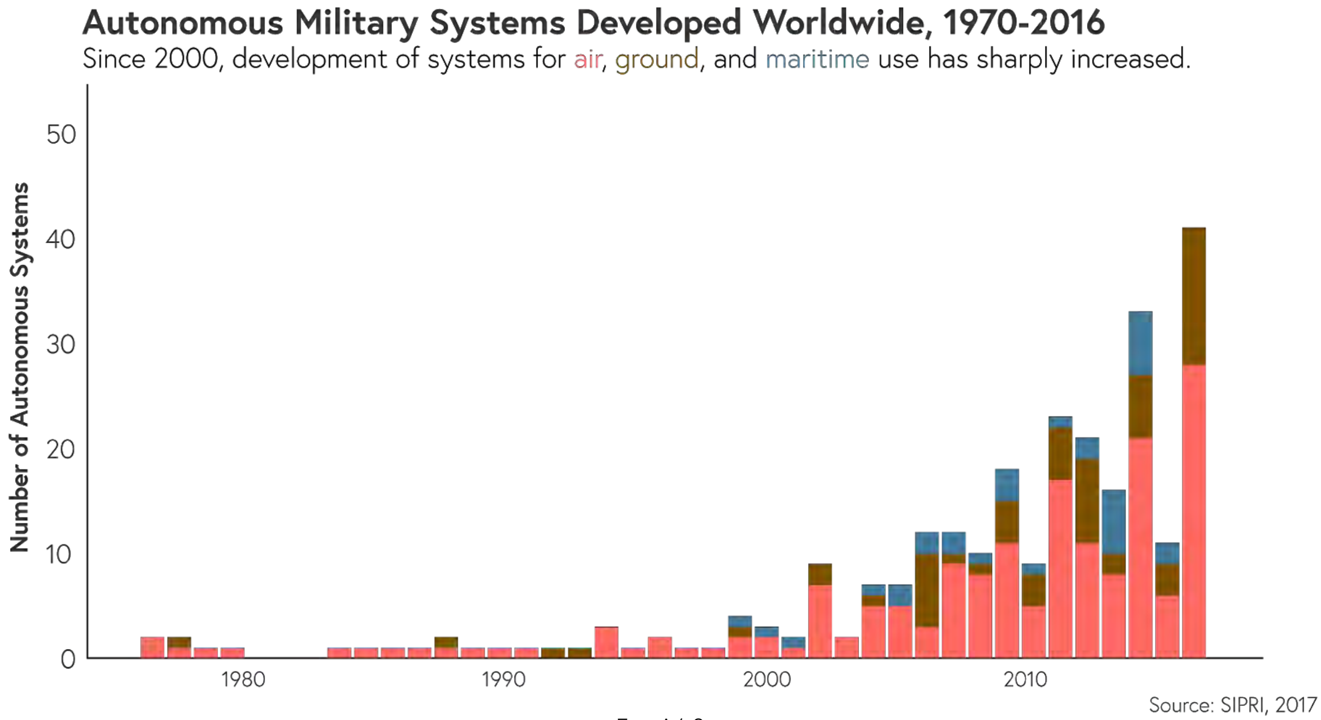

Fig. A6.2.

Central Bank and Corporate Perception

Source

Prattle provides sentiment data that predicts the market impact of central bank and corporate communications. Learn more about Prattle here.

Examples

Here are some examples of how AI is mentioned by central banks: in the first case, China uses a geopolitical environment simulation and prediction platform that works by crunching huge amounts of data and then providing foreign policy suggestions to Chinese diplomats or the Bank of Japan use of AI prediction models for foreign exchange rates. For the second case, many central banks are leading communications through either official documents, for example on 25 July 2019 the Dutch Central Bank (DNB) published Guidelines for the use of AI in financial services and launched its six "SAFEST" principles for regulated firms to use AI responsibly, or a speech on 4 June 2019 by the Bank of England's Executive Director of UK Deposit Takers Supervision James Proudman, titled "Managing Machines: the governance of artificial intelligence," focused on the increasingly important strategic issue of how Boards of regulated financial services should use AI.

Return to Public Perception - Corporate Perception

Government Perception Example

The 110th congress presented the congressional hearing on "The Use of Artificial Intelligence To Improve the US Department of Veteran Affairs' Claim Processing System" as well as the House Judiciary Committee records and reports related to US Code in National and Commercial Space.

Corporate Perception

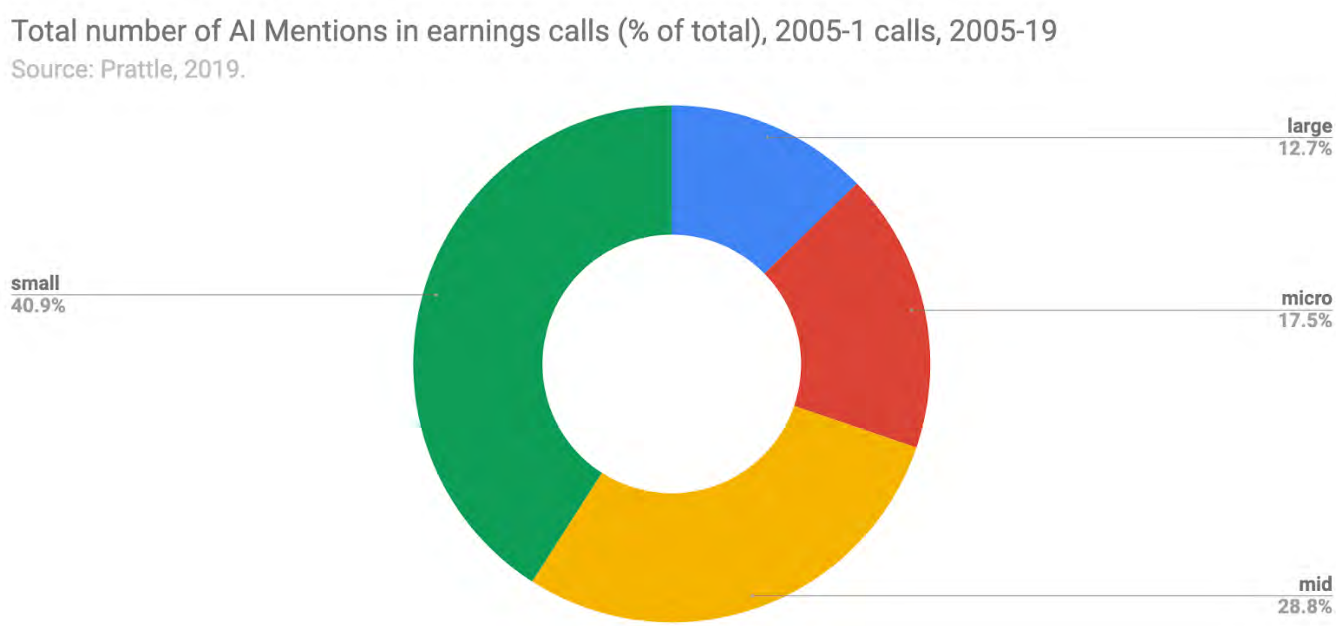

Fig. A7.1.

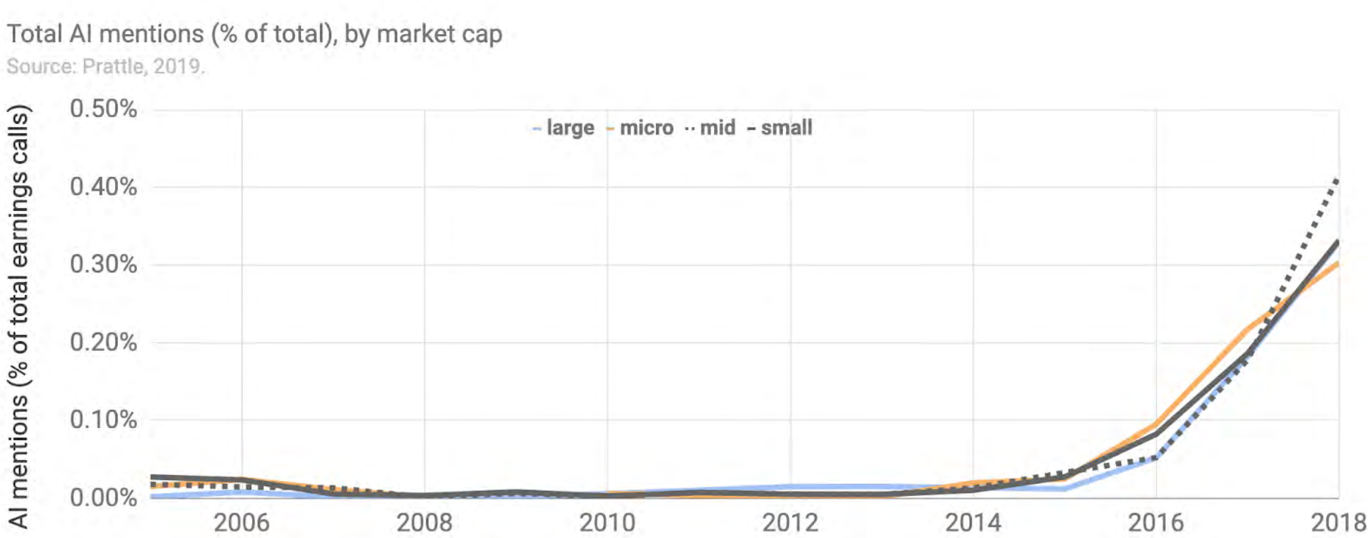

Fig. A7.2a.

Return to Public Perception - Corporate Perception

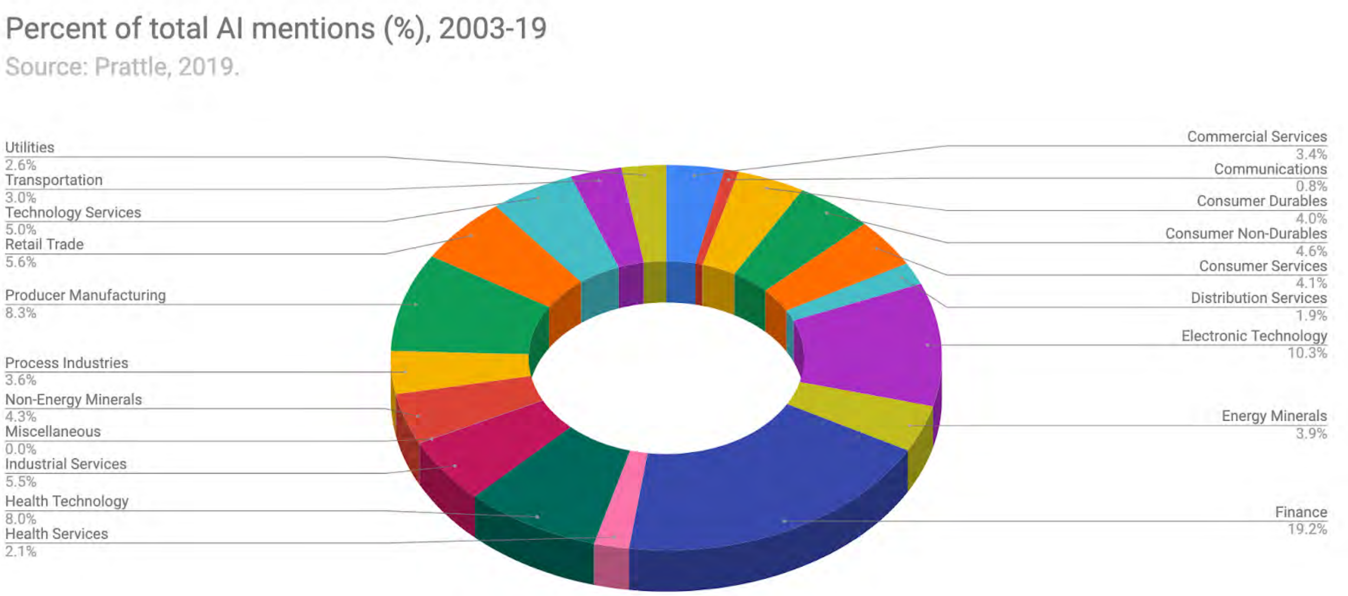

Fig. A7.2b.

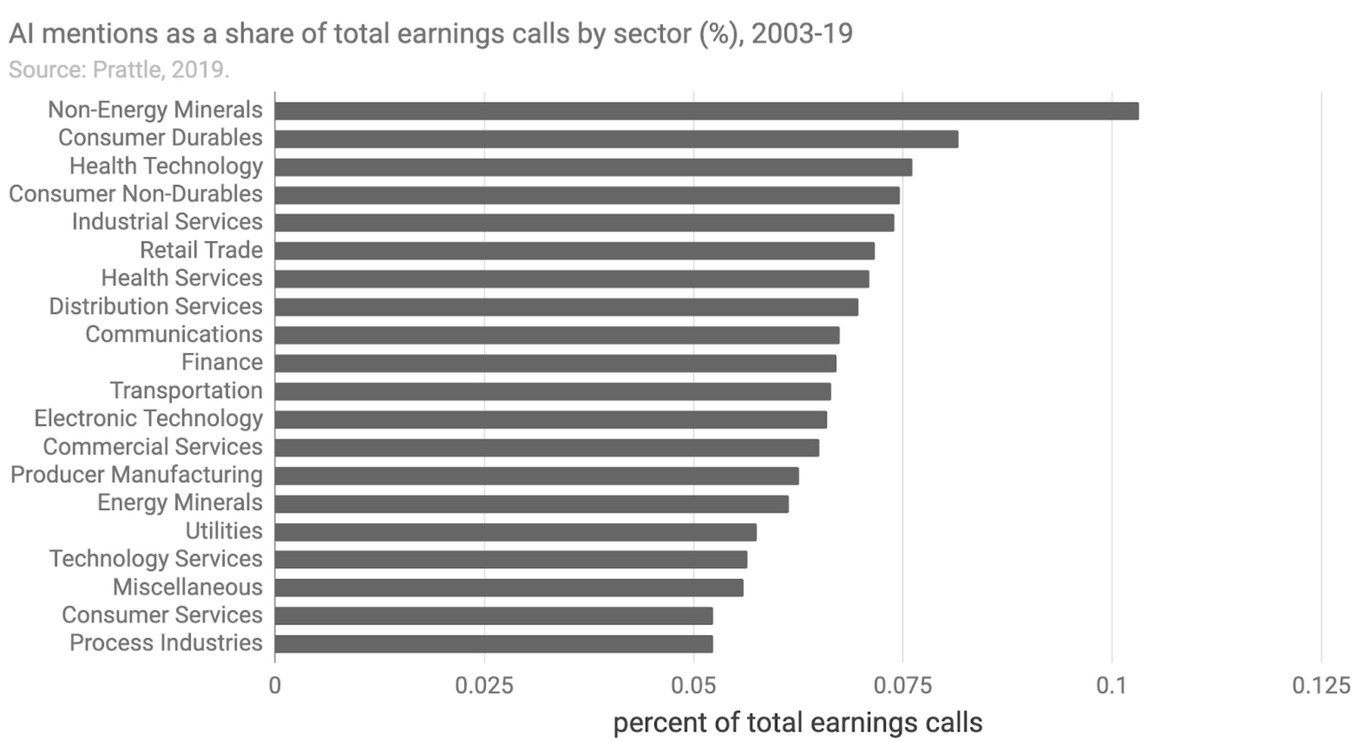

Fig. A7.2c.

Government mentions

Sources

Data collection and analysis was performed by the McKinsey Global Institute (MGI).

Methodologies, and Nuances, by country:

Canada (House of Commons):
Data was collected using the Hansard search feature on Parliament of Canada website. MGI searched for the terms "*Artificial Intelligence*" and "*Machine Learning*" (quotes included) and downloaded the results as a CSV. The date range was set to "All debates." Data is as of 11/20/2018. Data are available online from 08/31/2002.

Each count indicates that *Artificial Intelligence* or *Machine Learning* was mentioned in a particular comment or remark during the proceedings of the House of Commons. This means that within an event or conversation, if a member mentions *AI* or *ML* multiple times within their remarks, it will appear only once. However if, during the same event, the speaker mentions *AI* or *ML* in separate comments (with other speakers in between) it will appear multiple times. Counts for *Artificial Intelligence* and *Machine Learning* are separate, as they were conducted in separate searches. Mentions of the abbreviations "*AI*" or "*ML*" are not included.

United Kingdom (House of Commons, House of Lords, Westminster Hall, and Committees)
Data was collected using the Find References feature of the Hansard website of the UK Parliament. MGI searched for the terms "*Artificial Intelligence*" and "*Machine Learning*" (quotes included) and catalogued the results. Data is as of 11/20/2018. Data are available online from 1/1/1800 onwards. Contains Parliamentary information licensed under the Open Parliament Licence v3.0.

Like in Canada, each count indicates that *Artificial Intelligence* or *Machine Learning* was mentioned in a particular comment or remark during a proceeding. Therefore, if a member mentions *AI or ML* multiple times within their remarks, it will appear only once. However if, during the same event, the same speaker mentions *AI or ML* in separate comments (with other speakers in between) it will appear multiple times. Counts for *Artificial Intelligence* and *Machine Learning* are separate, as they were conducted in separate searches. Mentions of the abbreviations *"AI" or "ML"* are not included.

United States (Senate and House of Representatives)
Data was collected using the advanced search feature of the US Congressional Record website. MGI searched the terms "*Artificial Intelligence*" and "*Machine Learning*" (quotes included) and downloaded the results as a CSV. The "word variant" option was not selected, and proceedings included Senate, the House of Representatives, and Extensions of Remarks, but did not include the Daily Digest. Data is as of 11/20/2018, and data is available online from the 104th Congress onwards (1995).

Each count indicates that *Artificial Intelligence* or *Machine Learning* was mentioned during a particular event contained in the Congressional Record, including the reading of a bill. If a speaker mentioned *AI or ML* multiple times within remarks, or multiple speakers mentioned *AI or ML* within the same event, it would appear only once as a result. Counts for *Artificial Intelligence* and *Machine Learning* are separate, as they were conducted in separate searches. Mentions of the abbreviations *"AI" or "ML"* are not included.

Ethical Challenges

Sources

The data on ethical challenges and principles is curated by experts and topic modeling by PwC. Organizations globally, both private and public, are releasing core sets of ethical AI principles by which AI should operate. However, these principles vary organization by organization. To create a common set of principles, PwC has analyzed and categorized existing ethical AI principles documents for comparison. To learn more about PwC and PwC's work in Responsible AI, please see here.

Methodology

Candidate documents are updated on an ongoing basis. Team members then review the document for relevance; if the document is considered relevant for scrutiny, it is assigned a three letter acronym. The document is then reviewed, with principles identified and categorized according to principle definitions (see Appendix tables). This is a **living document**, and new entries are continuously added. Future documents will be categorized automatically, using AI/NLP methods. This document is not meant to be all inclusive; while extensive, **we recognize our list may not be fully exhaustive** given the frequency of release, breadth of organizations releasing such documents, and language considerations. The complete list of aggregated ethical principles are presented in the table below.

Table A8.1.

Ethical Challenges	Definition
Data Privacy	Users must have the right to manage their data which is used to train and run AI systems.
Beneficial AI	The development of AI should promote the common good.
Fairness	The development of AI should refrain from using datasets that contain discriminatory biases.
Accountability	All stakeholders of AI systems are responsible in the moral implications of their use and misuse.
AI Understanding	Designers and users of AI systems should educate themselves about AI.
Human Agency	A fully autonomous power should never be vested in AI technologies.
Diversity & Inclusion	Understand and respect the interests of all stakeholders impacted by your AI technology.
Safety	Throughout their operational lifetime, AI systems should not compromise the physical safety or mental integrity of humans.
Transparency	An AI system should be able to explain its decision making process in a clear and understandable manner.
Human Rights & Values	AI systems should be designed such that their behaviour and actions are aligned with human rights and values.
Lawfulness & Compliance	All the stakeholders in design of an AI system must always act in accordance with the law and all relevant regulatory regimes.
Reliability	AI systems should be development such that they will operate reliably over long periods of time using the right models and datasets.
Sustainability	The AI development must ensure the sustainability of our planet is preserved for future

Return to Societal Considerations - Ethical Challenges

Table A8.2. List of Organizational Documents

Acronym	Document title	Document Categorization	Issuer
MTL	Montreal Declaration for Responsible AI	Academia	Université de Montréal
ASM	Asilomar AI Principles	Associations & Consortiums	Future of Life Institute
IEE	IEEE Ethically Aligned Design v2	Associations & Consortiums	IEEE
EGE	Statement on AI, Robotics and 'Autonomous' Systems	Think Tanks / Policy Institutes	European Group on Ethics in Science and New Technologies
UKL	AI in the UK: ready, willing and able?	Official Government/Regulation	UK House of Lords
PAI	Tenets	Associations & Consortiums	Partnership on AI
OXM	Oxford-Munich Code of Conduct for Professional Data Scientists	Academia	University corsortium
MIG	Ethics Framework	Associations & Consortiums	Digital Catapult's Machine Intelligence Garage
GOO	AI at Google: our principles	Tech Companies	Google
MSF	Microsoft AI Principles	Tech Companies	Microsoft
ACC	Universal principles of data ethics - 12 guidelines for developing ethics codes	Industry & Consultancy	Accenture
IBM	Trusting AI	Tech Companies	IBM
KPM	Guardians of Trust	Industry & Consultancy	KPMG
DMN	Exploring the real-world impacts of AI	Tech Companies	DeepMind
COM	Community Principles on Ethical Data Practices	Associations & Consortiums	Datapractices.org - The Linux Foundation Projects
FWW	TOP 10 PRINCIPLES FOR ETHICAL ARTIFICIAL INTELLIGENCE	Associations & Consortiums	Future World of Work
I4E	The Responsible Machine Learning Principles	Associations & Consortiums	The Institute for Ethical AI & Machine Learning
A4P	AI4APEOPLE Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and	Associations & Consortiums	AI4PEOPLE - ATOMIUM
SGE	The Ethics of Code: Developing AI for Business with Five Core Principles	Tech Companies	Sage
PHS	Phrasee's AI Ethics Policy	Tech Companies	Phrasee
JAI	Japanese Society for Artificial Intelligence (JSAI) Ethical Guidelines	Associations & Consortiums	Japanese Society for Artificial Intelligence
DKN	Ethical principles for pro bono data scientists	Associations & Consortiums	Data Kind
ACM	ACM Code of Ethics and Professional Conduct	Associations & Consortiums	Association for Computing Machinery (ACM)

Table A8.2. List of Organizational Documents

Acronym	Document Title	Document Categorization	Issuer
COE	European ethical Charter on the use of Artificial Intelligence in judicial systems and their environment	Official Government/Regulation	EUROPEAN COMMISSION FOR THE EFFICIENCY OF JUSTICE (CEPEJ)
EUR	European Guidelines for Trustworthy AI	Official Government/Regulation	AI HLEG
AUS	Artificial Intelligence - Australia's Ethics Framework	Official Government/Regulation	Australian Government - Department of Industry, Innovation & Science
DUB	SMART DUBAI AI ETHICS PRINCIPLES & GUIDELINES	Official Government/Regulation	Smart Dubai
OEC	OECD Principles on AI	Think Tanks / Policy Institutes	OECD
G20	G20 Ministerial Statement on Trade and Digital Economy	Official Government/Regulation	G20
PDP	Singapore Personal Data Protection Commission	Official Government/Regulation	Singapore PDPC
DLT	AI Ethics: The Next Big Thing In Government	Industry & Consultancy	Deloitte
MEA	Work in the age of artificial intelligence. Four perspectives on the economy, employment, skills and ethics	Official Government/Regulation	Finland - Ministry of Economic Affairs and Employment
TIE	Tieto's AI ethics guidelines	Tech Companies	Tieto
OPG	Commitments and principles	Industry & Consultancy	OP Financial Group
FDP	How can humans keep the upper hand? Report on the ethical matters raised by AI algorithms	Think Tanks / Policy Institutes	France - Commission Nationale de l'Informatique et des Libertés
DTK	AI Guidelines	Industry & Consultancy	Deutsche Telekom
SAP	SAP's guiding principles for artificial intelligence	Tech Companies	SAP
AGI	L'intelligenzia artificiale al servizio del cittadino	Official Government/Regulation	[Agenzia per l'Italia Digitale]
ICP	Draft AI R&D Guidelines for International Discussions	Associations & Consortiums	Japan - Conference toward AI Network Society
SNY	Sony Group AI Ethics Guidelines	Tech Companies	Sony
TEL	AI Principles of Telefónica	Industry & Consultancy	Telefonica
IBE	Business Ethics and Artificial Intelligence	Think Tanks / Policy Institutes	Institute of Business Ethics
UKH	Initial code of conduct for data-driven health and care technology	Official Government/Regulation	UK - Department of Health and Social Care
IAF	Unified Ethical Frame for Big Data Analysis. IAF Big Data Ethics Initiative, Part A	Associations & Consortiums	The Information Accountability Foundation

Return to Societal Considerations - Ethical Challenges

Table A8.2. List of Organizational Documents

Acronym	Document Title	Document Categorization	Issuer
AMA	Policy Recommendations on Augmented Intelligence in Health Care H-480.940	Associations & Consortiums	AMA (American medical Association)
UNT	Introducing Unity's Guiding Principles for Ethical AI – Unity Blog	Tech Companies	Unity Technologies
GWG	Position on Robotics and Artificial Intelligence	Official Government/Regulation	The Greens, European Parliament
SII	Ethical Principles for Artificial Intelligence and Data Analytics	Associations & Consortiums	Software and Information Industry Association
ITI	ITI AI Policy Principles	Think Tanks / Policy Institutes	Information Technology Industry Council
WEF	White Paper: How to Prevent Discriminatory Outcomes in Machine Learning	Think Tanks / Policy Institutes	World Economic Forum
ICD	Declaration on ethics and data protection in Artificial Intelligence	Associations & Consortiums	International Conference of Data Protection and Privacy Commissioners
TPV	Universal Guidelines for Artificial Intelligence	Associations & Consortiums	The Public Voice Coalition
FAT	Principles for Accountable Algorithms and a Social Impact Statement for Algorithms	Associations & Consortiums	FATML
MAS	Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector	Official Government/Regulation	Monetary Authority of Singapore
VOD	Artificial Intelligence framework	Industry & Consultancy	Vodafone
DNB	General Principles for the use of Artificial Intelligence in the Financial Sector	Industry & Consultancy	DeNederlandsche Bank
IND	Artificial Intelligence in the Governance Sector in India	Associations & Consortiums	The centre for Internet & Society
DEK	Opinion of the Data Ethics Commission	Official Government/Regulation	Daten Ethik Kommission

Return to Societal Considerations - Ethics and AI: Global News Media

Global News Media

Sources

Quid is a data analytics platform that applies advanced natural language processing technology, semantic analysis, and artificial intelligence algorithms to reveal patterns in large, unstructured datasets, and generate visualizations to allow users to gain actionable insights. Quid uses Boolean query to search for focus areas, topics, and keywords within the archived news and blogs, companies and patents database, as well as any custom uploaded datasets. Users can then filter their search by published date time frame, source regions, source categories, or industry categories on the news; and by regions, investment amount, operating status, organization type (private/public), and founding year within the companies database. Quid then visualizes these data points based on the semantic similarity.

Network:

Searched for [AI technology keywords + Harvard ethics principles keywords] global news from 08/12/2018 ~ 08/12/2019

Search Query: (AI OR ["artificial intelligence"]("artificial intelligence" OR "pattern recognition" OR algorithms) OR ["machine learning"]("machine learning" OR "predictive analytics" OR "big data" OR "pattern recognition" OR "deep learning") OR ["natural language"]("natural language" OR "speech recognition") OR NLP OR "computer vision" OR ["robotics"]("robotics" OR "factory automation") OR "intelligent systems" OR ["facial recognition"]("facial recognition" OR "face recognition" OR "voice recognition" OR "iris recognition") OR ["image recognition"]("image recognition" OR "pattern recognition" OR "gesture recognition" OR "augmented reality") OR ["semantic search"]("semantic search" OR data-mining OR "full-text search" OR "predictive coding") OR "semantic web" OR "text analytics" OR "virtual assistant" OR "visual search") AND (ethics OR "human rights" OR "human values" OR "responsibility" OR "human control" OR "fairness" OR discrimination OR non-discrimination OR "transparency" OR "explainability" OR "safety and security" OR "accountability" OR "privacy")

News dataset data source:

Quid indexes millions of global-source English-language news articles and blog posts from LexisNexis. The platform has archived news and blogs from August 2013 to the present, updating every 15 minutes. Sources include over 60,000 news sources and over 500,000 blogs.

Visualization in Quid software:

Quid uses Boolean query to search for topics, trends, key words within the archived news database, with the ability to filter results by the published date time frame, source regions, source categories, or industry categories. (In this case, we only looked at global news published from 08/12/2018 to 08/12/2019) Quid then selects 10,000 most relevant stories using its NLP algorithm, and visualizes de-duplicated unique articles.

Return to Societal Considerations - Ethics and AI: Global News Media

How to read Quid Map in news:

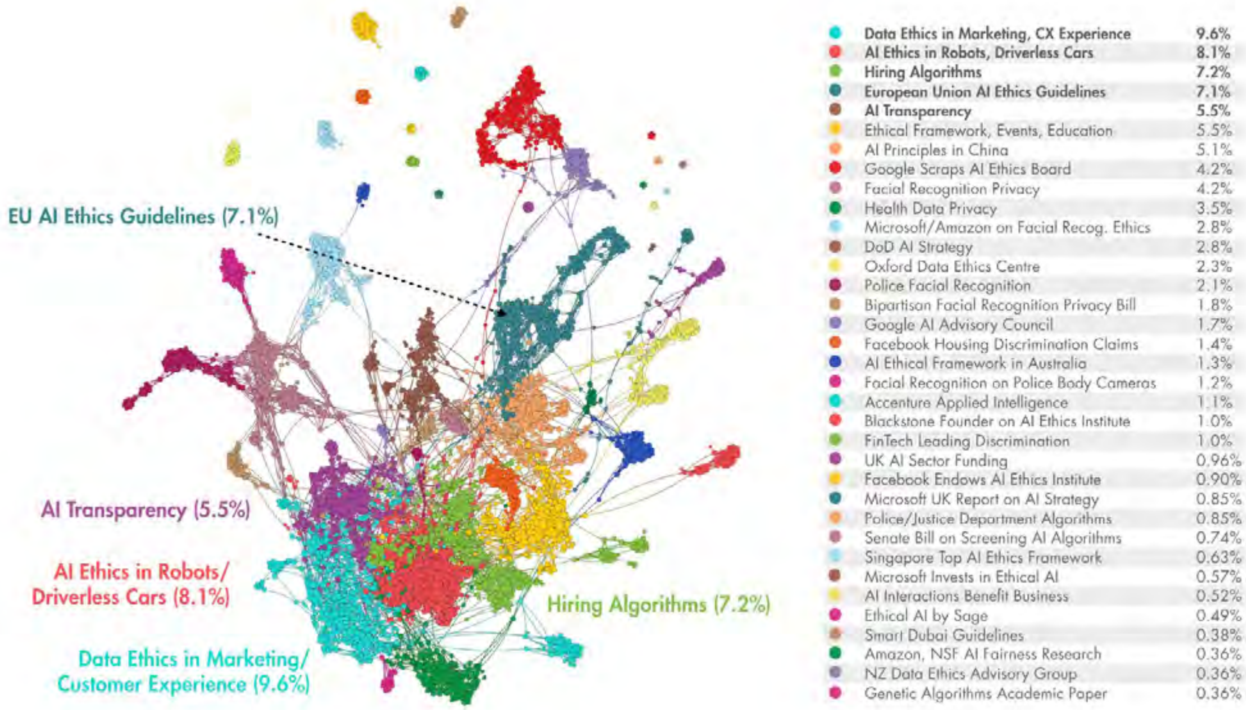

Fig. A8.1.

*How to read map visualization: Each node represents a news article. Links connect articles sharing similar languages. Clusters form when many articles share strong similarity, revealing topics.

When considering the network, cardinal directions (e.g. North, South, East, West) does not matter – what does matter is proximity. Two clusters which are close together (e.g. health data privacy and data ethics in marketing, customer experience) share more common language than the ones that are far away (e.g. Google Scraps AI ethics board). Centrality also matters – those clusters that are more central to network are more core to the narrative versus those on the periphery.

AI for Sustainable Development

Source

Data and analysis was provided by the McKinsey Global Institute. You can find additional details of MGI's research on AI for social good here.

Methodology

To build this use case library, MGI adopted a two-pronged approach, both societal and technological. From a societal point of view, MGI sought to identify key problems that are known to the social-sector community and determine where AI could aid efforts to resolve them. For a technological point of view, MGI took a curated list of 18 AI capabilities and sought to identify which types of social problems they could best contribute to solving.

For each use case, MGI tried to identify at least one existing case study. Where none were identified, they worked iteratively with experts to identify gaps and added corresponding use cases to our library. To guide their thinking, MGI conducted interviews with over 100 experts in the social sector, technology, business, and academia.

Each use case highlights a type of meaningful problem that can be solved by an AI capability or some combination of AI capabilities. The problem to solve was given a sufficiently broad definition so that similar solutions would be grouped together. The library is not comprehensive, but it nonetheless showcases a wide range of problems where AI can be applied for social good.

Return to National Strategies and Global AI Vibrancy

National AI Strategy Radar

Source

PwC's Global Data Analytics and AI consulting practices have been supporting government entities in their design of artificial intelligence national strategies , as well as enabling business globally to build, deploy and monitor enterprise AI. Some of these initiatives may be broad in mandate and difficult to define. Other countries have made strides to more clearly articulate their priorities, resulting in lengthy documents that can be a challenge to quickly consume and compare against others. To further those efforts, PwC created the National AI Strategy Radar (NAISR) to monitor advancements in and the changing landscape around how regulatory bodies discuss their priorities with respect to AI.

You can learn more about PwC's efforts working with national entities on AI here.

Methodology

The NAISR dashboard uses AI to monitor national AI strategies, by surfacing key priorities and topics that are discussed in policy documents and publications from regulatory bodies around the globe regarding AI and its implications. It helps assess what is being talked about where and the direction these discussions are taking.

Approach:

- Extracting all relevant PDFs along with metadata
- Extracting text from PDF by paragraphs
- Identifying all countries mentioned in the paragraphs using NER approach
- Identifying other entities/noun phrases in the paragraphs
- Topic modelling to cluster similar PDFs and get relevant themes for comparison
- Time series analysis and data visualization

Datasets Used:

- Recent government publications and related documents summarizing investments and priorities in the AI space

Table A9.1.

Country	Title	AuthorOrg	Publishing Date
Global	AI NOW 2017 Report	New York University	12/1/2017
UK	AI In The UK: Ready, Willing And Able?	UK Parliament – House of Lords	4/1/2018
Global	European Union regulations on algorithmic decision-making and a "right to explanation"	Oxford University	8/1/2016
Global	Smart Policies for Artificial Intelligence	Miles Brundage, Joanna Bryson	8/1/2016
Global	The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation	Future of Humanity Institute, Oxford University,Centre for the Study of Existential Risk, University of Cambridge,Center for a New American Security,Electronic Frontier Foundation,OpenAI	2/1/2018
Global	On the Promotion of Safe and Socially Beneficial Artificial Intelligence	AI & Society	10/1/2017
Global	Artificial Intelligence Index: 2017 Annual Report	Yoav Shoham, Raymond Perrault, Erik Brynjolfsson, Jack Clark, Calvin LeGassick	11/1/2017
Global	Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies	Harvard Journal of Law & Technology	1/1/2016
Global	Artificial General Intelligence	Foresight Institute	11/1/2017
Global	Artificial Intelligence and National Security	Harvard Kennedy School	7/1/2017
Global	Artificial Intelligence and Foreign Policy	Stiftung Neue Verantwortung	1/1/2018
Global	Artificial Intelligence and Life in 2030	Stanford University, AI100	9/1/2016
Global	Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability	AI Now	4/1/2018
Global	Regulating Artificial Intelligence Proposal for a Global Solution	Association for the Advancement of Artificial Intelligence	1/1/2018
Global	Policy Desiderata in the Development of Superintelligent AI	Future of Humanity Institute, Oxford University,Yale University	1/1/2017
Australia	Prosperity Through Innovation	Australian Government	1/11/2017
US	Federal Automated Vehicles Policy: Accelerating the Next Revolution In Roadway Safety	US Department of Transportation, National Highway Traffic Safety Administration	9/1/2016
Global	Data management and use: Governance in the 21st century	British Academy, The Royal Society	6/1/2017
Denmark	National Strategy for Artificial Intelligence	Danish Government	1/3/2019
Global	Destination unknown: Exploring the impact of Artificial Intelligence on Government September 2017 Working Paper	Center for Public Impact	9/1/2017

Table A9.1.

Country	Title	AuthorOrg	Date
Global	Existential Risk Diplomacy and Governance	Global Priorities Project	1/1/2017
Finland	Finland's Age of Artificial Intelligence	Ministry of Economic Affairs and Employment of Finland	1/12/2017
France	Machine Politics Europe and the AI Revolution	European Council on Foreign Relations	1/6/2019
Global	Artificial Intelligence: An Overview Of State Initiatives	Future Grasp	1/6/2019
Germany	Artificial Intelligence Strategy	The Federal Government Germany	1/11/2018
Global	Global Catastrophic Risks 2016	Global Challenges Foundation	1/11/2018
India	National Strategy for Artificial Intelligence #AI For All	NITI Aayog	1/6/2018
Global	International Cooperation vs. AI Arms Race	Foundational Research institute	12/1/2013
Japan	Artificial Intelligence Technology Strategy	Strategic Council for AI Technology	1/3/2017
Korea	Mid- to Long-Term Master Plan in Preparation for the Intelligent Information Society	Government of the Republic of Korea Interdepartmental Exercise	1/12/2016
Global	Making the AI revolution work for everyone	The Future Society, AI Initiative	1/1/2017
France	For A Meaningful Artificial Intelligence: Towards A French And European Strategy	French Parliament	3/1/2018
US	The National Artificial Intelligence Research And Development Strategic Plan	Executive Office of the President National Science and Technology Council Committee on Technology	10/1/2016
Global	Strategic Implications of Openness in AI Development	Oxford University, Future of Humanity Institute	1/1/2017
Poland	Map of the Polish AI	Digital Poland Foundation	1/1/2019
US	Preparing For The Future Of Artificial Intelligence	Executive Office of the President National Science and Technology Council Committee on Technology	10/1/2016
Qatar	National Artificial Intelligence Strategy For Qatar	Qatar Computing Research Institute	1/1/2018
Global	Racing To The Precipice: A Model Of Artificial Intelligence Development	Future of Humanity Institute	12/1/2013
Global	How Might Artificial Intelligence Affect the Risk of Nuclear War?	Edward Geist and Andrew J. Lohn, Security 2040, RAND Corporation	1/1/2018
Saudi Arabia	Vision 2030	Council of Economic and Development Affairs	1/1/2018
Sweden	National approach to artificial intelligence	Government Offices of Sweden	1/1/2018

Return to National Strategies and Global AI Vibrancy

Table A9.1.

Country	Title	AuthorOrg	Date
Switzerland	Digital Switzerland Strategy	Switzerland Federal Council	1/9/2018
Taiwan	AI Taiwan	Taiwan Cabinet	1/9/2018
Global	The MADCOM future: how artificial intelligence will enhance computational propaganda, reprogram human culture, and threaten democracy... And what can be done about it	Atlantic Council	9/1/2017
Global	The Future of Employment: how susceptible are jobs to computerisation?	Carl Benedikt Frey, Michael A. Osborne	9/1/2013
China	A Next Generation Artificial Intelligence Development Plan	China State Council	7/1/2017
UK	AI in the UK: Ready, Willing and Able?	Secretary of State for Business, Energy and Industrial Strategy	1/6/2018
Global	Artificial Intelligence and Robotics for Law Enforcement	United Nations Interregional Crime and Justice Research Institute	1/1/2019
Global	Unprecedented Technological Risks	Future of Humanity Institute, Oxford University,Centre for the Study of Existential Risk, University of Cambridge	9/1/2014
Global	Artificial Intelligence: The Race Is On The Global Policy Response To AI	FTI Consulting	2/1/2018

Topic Concept Graph of AI Strategy Documents

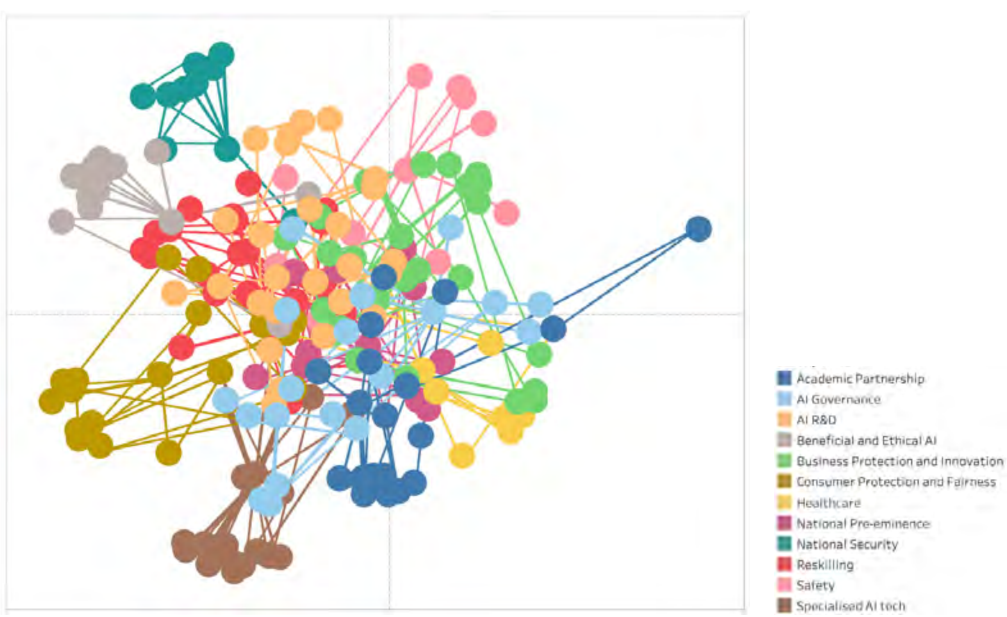

Fig. A9.1.

Non-exhaustive List of AI Strategies and Policies in place

Note: The listing below was manually generated, and used to inform the "Strategy In Place" portion of the NAISR dashboard. Not all documents below were included in the broader topic and thematic analysis; many were too high level or not directly relevant, and were therefore excluded from the topic modeling exercises.

a) **Australia**: Australia has dedicated $29.9 million in the country's annual budget to promote and guide the development of AI.

- AI and automation are already considered under the national Innovation Strategy and are also featured in several more recent initiatives
- The Australian Government unveiled a new Digital Economy Strategy on September 19, 2017

b) **Canada**: Canada has a national AI strategy called the Pan-Canadian Artificial Intelligence Strategy.

- Launched the Pan-Canadian Artificial Intelligence (AI) Strategy in its 2017 Budget with the allocation of $125 million.

c) **China**: China has a national AI strategy, defined under the "New Generation Artificial Intelligence Development Plan."

- July 2017, The State Council of China released the "New Generation Artificial Intelligence Development Plan" which outlines China's strategy to build a domestic AI industry worth nearly US$150 billion in the next few years and to become the leading AI power by 2030
- Back in 2016, the Chinese Three-Year Guidance for Internet Plus Artificial Intelligence Plan (2016-2018) indicated an intention to make AI a strong driving force in socioeconomic development. The Three-Year Action Plan for Promoting Development of a New Generation Artificial Intelligence Industry (2018–2020) reinforced this goal.

d) **Denmark**: Denmark has a digital strategy that includes a focus on AI along with other technologies.

- In January 2018, the Danish Government launched the "Strategy for Denmark's Digital Growth," which consists of seven main initiatives: Digital Hub Denmark; SME:Digital; The Technology Pact; Strengthened Computational Thinking in Elementary School; Data as a Driver of Growth; Agile Regulation for New Business Models; and Strengthened Cyber Security in Companies.
- In October 2017, Denmark published, "Towards a Digital Growth Strategy – MADE," which identified AI as a major growth area, with a Danish center for artificial intelligence (DCKAI) listed as one of the targeted strategies.

e) **Finland**: Finland has an Artificial Intelligence Programme guided by a steering group under the Ministry of Economic Affairs and Employment.

- first report in December 2017 titled, "Finland's Age of Artificial Intelligence: Turning Finland into a leading country in the application of artificial intelligence."
- report June 2018 titled, "Artificial Intelligence: Four Perspectives on the Economy, Employment, Knowledge and Ethics." The report provides 28 policy recommendations related to the effects of AI on economics and employment, the labor market, education and skills management, and ethics.

f) **France**: France has a national strategy for AI called "AI for Humanity," which is outlined in the "Villani Report".

- developed a national strategy for AI titled "AI for Humanity" outlined in the "Villani Report".
- Digital Republic Bill. Its objective included ensuring "characteristics that must be at the heart of the French AI model: respect for privacy, protection of personal data, transparency, accountability of actors and contribution to collective wellbeing."

g) **Germany**: The German Government adopted its Artificial Intelligence Strategy in November 2018.

- Adopted a national AI strategy (available to download here) and earmarked €3 billion for investment in AI research and development.
- launched a government aid campaign in the field of machine learning in 2017
- German Federal Ministry of Transport and Digital Infrastructure (BMVI) published ethical guidelines for self-driving cars in a report titled, "Ethics Commission: Automated and Connected Driving," which defined 20 ethical rules for automated and connected vehicular traffic.

Germany has since released findings from the Data Ethics Commission

Non-exhaustive List of AI Strategies and Policies in place

h) **India**: India defined a national policy on AI in a working paper titled, "National Strategy for Artificial Intelligence #AI forAll."

- defined a national policy on AI in a working paper titled, "National Strategy for Artificial Intelligence #AIforAll."

i) **Japan**: Japan has an "Artificial Intelligence Technology Strategy" and has also included AI in its "integrated innovation strategy."

- "Artificial Intelligence Technology Strategy" in March 2017
- On July 28, 2017, Japan published Draft AI R&D GUIDELINES for International Discussions in preparation for the Conference toward AI Network Society.

j) **Singapore**: Singapore has a national AI program called AI Singapore and is establishing an AI ethics advisory council.

- AI Singapore is the national program established in May 2017 to harness AI throughout the country.

k) **South Korea**: South Korea has an Artificial Intelligence Information Industry Development Strategy.

- Defined an Artificial Intelligence Information Industry Development Strategy in 2016 (70-page report) which lays out a "National Vision" which is "Realizing a Human-Centered Intelligent Information Society."

l) **Sweden**: The Swedish government has released a "National Approach for Artificial Intelligence."

- In May 2018 Sweden released their "National Approach for Artificial Intelligence," (translated to English here,) a 12-page guiding document outlining the governments' assessments of what is needed for the country to be at the forefront of AI development and use.

m) **United Arab Emirates**: The UAE has a national strategy for AI and was the first country to name an AI Minister.

- In October 2017, the UAE Government announced the UAE Strategy for Artificial Intelligence

n) **United States of America**: The US launched the American AI Initiative February 2019.

- US President Donald Trump issued an Executive Order launching the American AI Initiative on February 11, 2019
- The day after the Executive Order was released, the US Department of Defense followed up with its own Artificial Intelligence Strategy

o) **United Kingdom**: The UK government launched a Sector Deal for AI to advance the UK's ambitions in AI consistent with its Industrial Strategy, and taking into account the advice of the Parliament's Select Committee on AI.

- On March 6, 2018 the UK government launched a Sector Deal for AI led by Business Secretary Greg Clark. The Deal aims to take "immediate, tangible actions" to advance the UK's ambitions in AI that are consistent with the Industrial Strategy
- The UK Government's Industrial Strategy was published in November 2017. The section on the Grand Challenges (pg. 30) features AI.

Return to National Strategies and Global AI Vibrancy

In Consideration / Development in Progress:

p) **Estonia**: Estonia is developing a legal framework for the use of AI in its country, including a bill on AI liability.
- developing a bill for AI liability which will be ready in March 2019
- developing a legal framework around use of AI

q) **Italy**: Italy has an interdisciplinary AI Task Force launched by the Agency for Digital Italy which released a White Paper called "AI at the service of citizens," in March 2018

r) **Malaysia**: The Malaysian government is developing a National Artificial Intelligence Framework, and establishing Digital Transformation Labs.
- existing National Big Data Analytics Framework
- announced plans to develop a National Artificial Intelligence Framework

s) **Mexico**: The Mexican government supported the creation of the white paper, "Towards an AI Strategy in Mexico: Harnessing the AI Revolution."
- A white paper titled "Towards an AI Strategy in Mexico: Harnessing the AI Revolution" was published in June 2018
- IA2030 Coalition, which is a group of people helping to enhance understanding of AI and realize a Mexican AI strategy

t) **Russia**: The Russian government is currently developing an AI R&D national strategy.
- a 10-point plan for AI development in Russia

u) **Tunisia**: Tunisia has created an AI Task Force and Steering Committee to develop a national AI strategy.

Related, But No Mentions of a Strategy:

v) **Austria**: Austria has an advisory Robot Council that is developing a national AI strategy.
- established a Robot Council in August 2017
- In January 2018, the new government proposed the establishment of an Ethics Council for Digitalization
- established the National Robotics-Technology Platform (GMAR) in 2015 to promote robotics, automation, and AI technology

w) **Ireland**: The Irish government has hosted AI workshops and launched a national AI Masters program.

x) **Kenya**: The Kenyan government created a Blockchain & Artificial Intelligence task force.

y) **New Zealand**: New Zealand has an AI Forum to connect and advance the country's AI ecosystem. This report is not a national AI strategy, but it explores the New Zealand AI landscape and the potential impacts of AI on the economy and society.

z) **Saudi Arabia**: Saudi Arabia was the first country to grant citizenship to a robot.

Return to [National Strategies]

Traceability Matrix of National AI documents

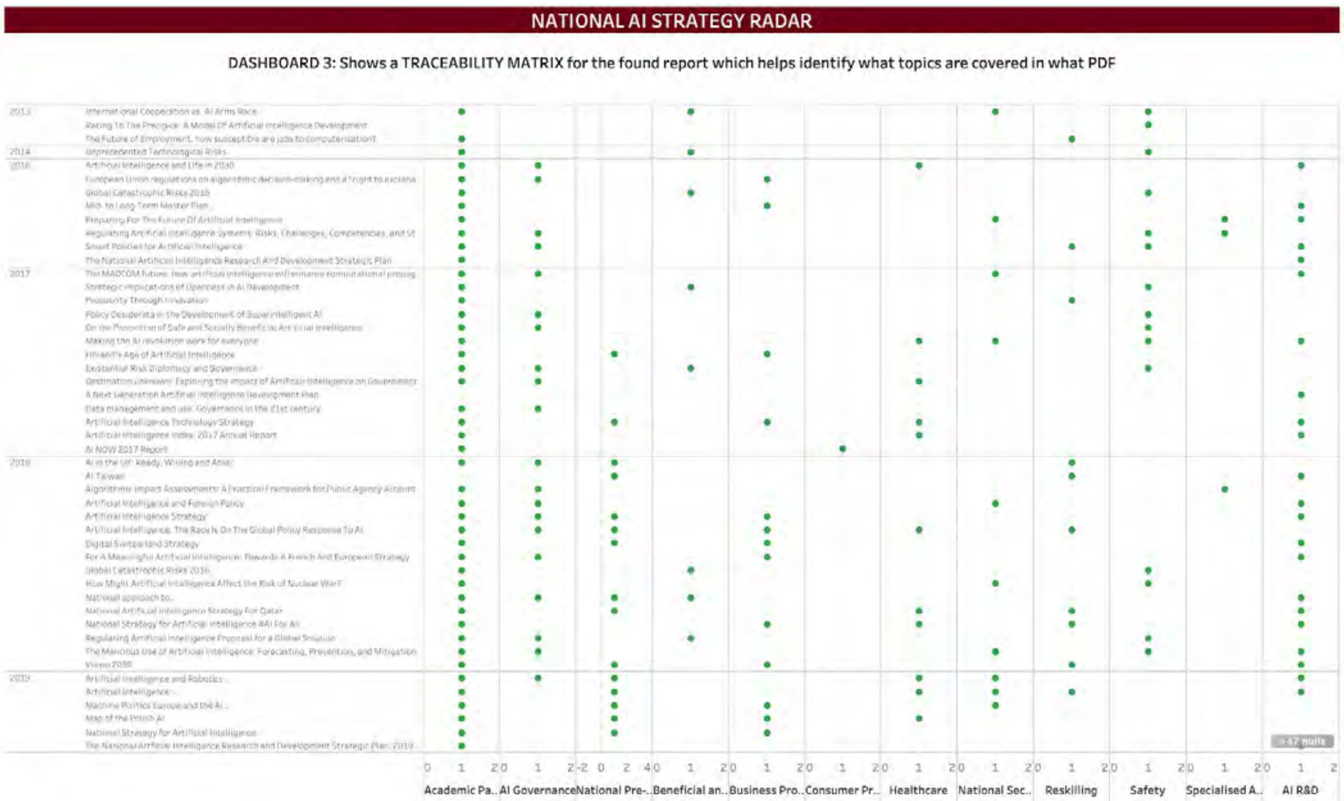

Construction the AI Vibrancy Index: Composite Measure

Go to Global AI Vibrancy: Country Weighting Tool

Source

The data is collected by AI Index using diverse datasets that are referenced in the 2019 AI Index Report Chapters.

Methodology

Step 1: Obtain, harmonizing, and integrating data on individual attributes across countries and time
Step 2: Use Min-Max Scalar to normalize each indicator between 0-100
Step 3: Arithmetic Mean per country over years
Step 4: Build Modular Weighted by high and low level categories

Aggregate Measure

The overall AI Vibrancy Index: Composite Measure is composed of the following high level pillars. This can be represented in the following simple equation:

$$AI\ Vibrancy_{c,t} = W_1 * (R\&D) + W_2 * (Economy) + W_3 * (Inclusion)$$

The approach can be improved by assigning error-bands to each metric associated with the raw data and measurement related uncertainties.

Normalization

To adjust for differences in units of measurement and ranges of variation, all 36 variables were normalised into the [0, 100] range, with higher scores representing better outcomes. A min-max normalisation method was adopted, given the minimum and maximum values of each variable respectively. For variables where higher values indicate better outcomes, the following normalisation formula was applied:

$$Min\text{-}max\ scalar\ (MS100) = 100 * \frac{(value) - (min)}{(max) - (min)}$$

Note all variables currently used have higher value corresponding to better outcome.

Scaled Sub-pillar Weighting

The score for each pillar is a weighted sum of its components.

Return to Global AI Vibrancy: Country Weighting Tool

Research and Development

id	Pillar	Sub-Pillar	Name	Definition	Source
1	**Research and Development**	Conference Publications	Number of AI conference papers*	Total count of published AI conference papers attributed to institutions in the given country.	Microsoft Academic Graph (MAG)
2	**Research and Development**	Conference Publications	Number of AI conference papers per capita	Total count of published AI conference papers attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
3	**Research and Development**	Conference Publications	Number of AI conference citations*	Total count of AI conference citations attributed to institutions in the given country.	Microsoft Academic Graph (MAG)
4	**Research and Development**	Conference Publications	Number of AI conference citations per capita	Total count of AI conference citations attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
5	**Research and Development**	Conference Publications	Number of AI conference references*	Total count of AI conference references attributed to institutions in the given country.	Microsoft Academic Graph (MAG)
6	**Research and Development**	Conference Publications	Number of AI conference references per capita	Total count of AI conference references attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
7	**Research and Development**	Journal Publications	Number of AI journal papers*	Total count of published AI journal papers attributed to institutions in the given country.	Microsoft Academic Graph (MAG)
8	**Research and Development**	Journal Publications	Number of AI journal papers per capita	Total count of published AI journal papers attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
9	**Research and Development**	Journal Publications	Number of AI journal citations*	Total count of AI journal citations attributed to institutions in the given country.	Microsoft Academic Graph (MAG)
10	**Research and Development**	Journal Publications	Number of AI journal citations per capita	Total count of AI journal citations attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
11	**Research and Development**	Journal Publications	Number of AI journal references*	Total count of AI journal references attributed to institutions in the given country.	Microsoft Academic Graph (MAG)
12	**Research and Development**	Journal Publications	Number of AI journal references per capita	Total count of AI journal references attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
13	**Research and Development**	Innovation > Patents	Number of AI patents*	Total count of published AI patents attributed to institutions in the given country.	Microsoft Academic Graph (MAG)

Return to [Global AI Vibrancy: Country Weighting Tool]

Research and Development

14	**Research and Development**	Innovation > Patents	Number of AI patents per capita	Total count of published AI patents attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
15	**Research and Development**	Innovation > Patents	Number of AI patent citations*	Total count of published AI patents citations attributed to institutions of originating patent filing.	Microsoft Academic Graph (MAG)
16	**Research and Development**	Innovation > Patents	Number of AI patent citations per capita	Total count of published AI patent citations attributed to institutions in the given country of originating patent filing, in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	Microsoft Academic Graph (MAG)
17	**Research and Development**	Innovation > Patents	Number of AI patent references*	Total count of published AI patent references attributed to institutions in the given country of originating patent filing, in per capita terms.	Microsoft Academic Graph (MAG)
18	**Research and Development**	Innovation > Patents	Number of AI patent references per capita	Total count of published AI patent references attributed to institutions in the given country of originating patent filing, in per capita terms. The denominator is population in millions for a given year to obtain appropriately scaled values.	Microsoft Academic Graph (MAG)
19	**Research and Development**	Journal Publications > Deep Learning	Number of Deep Learning papers*	Total count of arXiv papers on Deep Learning attributed to institutions in the given country.	arXiv, NESTA
20	**Research and Development**	Journal Publications > Deep Learning	Number of Deep Learning papers per capita	Total count of arXiv papers on Deep Learning attributed to institutions in the given country in per capita terms. The denominator is population in millions for a given year to obtain scaled values.	arXiv, NESTA
21	**Research and Development**	Journal Publications > Deep Learning	Revealed Comparative Advantage (RCA) of Deep Learning Papers on arXiv	Measure of relative specialization in Deep Learning papers based on arXiv at the country level.	arXiv, NESTA

Return to Global AI Vibrancy: Country Weighting Tool

Economy

id	Pillar	Sub-Pillar	Name	Definition	Source
22	**Economy**	Skills	Percentile Rank of AI Skills on Coursera	Coursera AI Global Skill Index Percentile Rank	Coursera
23	**Economy**	Skills	AI (% of total enrollment)	Percent of online students enrolled in AI courses in the given country.	Coursera
24	**Economy**	Skills	Relative Skill Penetration	Relative skill penetration rate (this is a method to compare how prevalent AI skills are at the average occupation in each country against a benchmark (here the global average), controlling for the same set of occupations	LinkedIn Economic Graph
25	**Economy**	Skills	Number of unique AI occupations (job titles)	Number of unique AI occupations (or job titles) with high AI skill penetration	LinkedIn Economic Graph
26	**Economy**	Labor	AI hiring index	AI hiring rate is the percentage of LinkedIn members who had any AI skills (see appendix for the AI skill grouping) on their profile and added a new employer to their profile in the same month the new job began, divided by the total number of LinkedIn members in the country. This rate is then indexed to the average month in 2015-2016; for example, an index of 1.05 indicates a hiring rate that is 5% higher than the average month in 2015-2016.	LinkedIn Economic Graph
27	**Economy**	Investment	Total Amount of Funding*	Total amount of Private Investment Funding received for AI startups (nominal US$).	Crunchbase, CapIQ, Quid
28	**Economy**	Investment	Total per capita Funding	Total amount of Private Investment Funding received for AI startups in per capita terms. The denominator is population in millions for a given year to obtain appropriately scaled values.	Crunchbase, CapIQ, Quid
29	**Economy**	Investment	Number of Startups Funded*	Total number of AI companies founded in the given country.	Crunchbase, CapIQ, Quid
30	**Economy**	Investment	Number of funded startups per capita	Total number of AI companies founded in the given country in per capita terms.	Crunchbase, CapIQ, Quid
31	**Economy**	Robot Installations	Robot Installations (in thousands of units)	Number of industrial robots installed in the given country (in 1000's of units).	International Federation of Robotics (IFR)

Inclusion

id	Pillar	Sub-Pillar	Name	Definition	Source
32	**Inclusion**	Gender Diversity	Proportion of female AI authors	Percentage of AI papers on arXiv where one author is attributed to be female.	arXiv, NESTA
33	**Inclusion**	Gender Diversity	AI Skill Penetration (female)	Relative skill penetration rate (this is a method to compare how prevalent AI skills are at the average occupation in each country against a benchmark (here the global average), controlling for the same set of occupations. The female AI skill penetration measure is a relative measure of female AI skill penetration in a country to global female AI skill penetration.	LinkedIn Economic Graph
34	**Inclusion**	Gender Diversity	Number of unique AI occupations (job titles), female	Number of unique AI occupations (or job titles) with high AI skill penetration for females in a given country.	LinkedIn Economic Graph