



LAII


artificial intelligence index

2019 annual report



Stanford
Human-Centered
Artificial Intelligence





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

| | | |
|------------------|----------------------------------------------------------------|------------|
| | Introduction | 4 |
| | Report Highlights | 5 |
| | Acknowledgements | 9 |
| Chapter 1 | Research and Development | 12 |
| Chapter 2 | Conferences | 37 |
| Chapter 3 | Technical Performance | 47 |
| Chapter 4 | The Economy | 71 |
| Chapter 5 | Education | 106 |
| Chapter 6 | Autonomous Systems | 127 |
| Chapter 7 | Public Perception | 136 |
| Chapter 8 | Societal Considerations | 146 |
| Chapter 9 | National Strategies and Global AI Vibrancy | 156 |
| | Technical Appendix | 182 |



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.

(c) 2019 by Stanford University, "The AI Index 2019 Annual Report" is made available under a Creative Commons Attribution-NoDerivatives 4.0 License (International)

<https://creativecommons.org/licenses/by-nd/4.0/legalcode>

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

McKinsey&Company



AI21labs



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 Commission

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 Tembhornikar

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, Maaïke 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); Fangzhzhen 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.

