

The evolution of citation graphs in artificial intelligence research

Morgan R. Frank¹, Dashun Wang^{2,3}, Manuel Cebrian¹ and Iyad Rahwan^{1,4,5*}

As artificial intelligence (AI) applications see wider deployment, it becomes increasingly important to study the social and societal implications of AI adoption. Therefore, we ask: are AI research and the fields that study social and societal trends keeping pace with each other? Here, we use the Microsoft Academic Graph to study the bibliometric evolution of AI research and its related fields from 1950 to today. Although early AI researchers exhibited strong referencing behaviour towards philosophy, geography and art, modern AI research references mathematics and computer science most strongly. Conversely, other fields, including the social sciences, do not reference AI research in proportion to its growing paper production. Our evidence suggests that the growing preference of AI researchers to publish in topic-specific conferences over academic journals and the increasing presence of industry research pose a challenge to external researchers, as such research is particularly absent from references made by social scientists.

Today's artificial intelligence (AI) has implications for the future of work¹, the stock market^{2,3}, medicine^{4,5}, transportation^{6,7}, the future of warfare⁸ and the governance of society^{9–11}. On one hand, AI adoption has the positive potential to reduce human error and human bias¹². As examples, AI systems have balanced judges towards more equitable bail decisions¹³, AI systems can assess the safety of neighbourhoods from images¹⁴ and AI systems can improve hiring decisions for board directors while reducing gender bias¹⁵. On the other hand, recent examples suggest that AI technologies can be deployed without understanding the social biases they possess or the social questions they raise. Consider the recent reports of racial bias in facial recognition software^{16,17}, the ethical dilemmas of autonomous vehicles⁶ and income inequality from computer-driven automation^{18–20}.

These examples highlight the diversity of today's AI technology and the breadth of its application; an observation leading some to characterize AI as a general-purpose technology^{1,21}. As AI becomes increasingly widespread, researchers and policymakers must balance the positive and negative implications of AI adoption. Therefore, we ask: how tightly connected are the social sciences and cutting-edge machine intelligence research?

Here, we employ the Microsoft Academic Graph (MAG) to explore the research connections between AI research and other academic fields through citation patterns. The MAG data offer coverage for both conference proceedings, where AI papers are often published, and academic journals, where other fields prefer to publish. Although early AI research was inspired by the several other fields, including some social sciences, modern AI research is increasingly focused on engineering applications—perhaps due to the increasingly central role of the technology industry. Furthermore, the most central research institutions within the AI research community are increasingly based in industry rather than academia.

Modern AI research

The effort to create human-like intelligence has dramatically advanced in recent decades thanks to improvements in algorithms

and computers. However, engineering the entirety of human intelligence has proved difficult. Instead, progress has come from engineering specific human capabilities. While we often use the term AI today in reference to machine learning, the meaning of AI has fluctuated in the past 60 years to variably emphasize vision, language, speech and pattern recognition.

To study the nature of AI research, we use the MAG to identify relevant computer science (CS) subfields from the citations of academic publications from 1950 to 2018. The MAG uses natural language processing (NLP), including keyword analysis, to identify the academic field of each publication according to a hierarchy of academic fields. These data have been particularly useful for studying bibliometric trends in CS^{22–25}. Our analysis relies strongly on the MAG's field of study classifications and, thus, our analysis is potentially limited in its accounting of more specific research areas within CS and within AI-related fields. These data enable us to study the paper production and referencing behaviour of different academic fields. For example, CS has risen to the fourth most productive academic field according to annual paper production (see Supplementary Fig. 1) with AI being the most prominent subfield of CS in recent decades²⁶ (see also Fig. 1d).

To identify the CS subfields that are most relevant to AI research, we construct a citation network using all CS papers published within each decade from 1950 to 2018. We consider CS subfields to represent AI research if they are strongly associated with AI, which is itself a CS subfield, throughout a significant proportion of the time period under analysis. Examples include computer vision, machine learning and pattern recognition. Interestingly, NLP, which is colloquially thought of as a specific problem area in AI²⁷, is strongly associated with AI research before the mid 1980s, after which NLP becomes more strongly associated with information retrieval and data mining for text-based data (Fig. 1a–c,e). In the remainder, we use papers published in AI, computer vision, machine learning, pattern recognition and NLP to approximate AI research from the 1950s to today.

¹Media Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA. ²Kellogg School of Management, Northwestern University, Evanston, IL, USA. ³Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL, USA. ⁴Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁵Center for Humans and Machines, Max Planck Institute for Human Development, Berlin, Germany.

*e-mail: irahwan@mit.edu

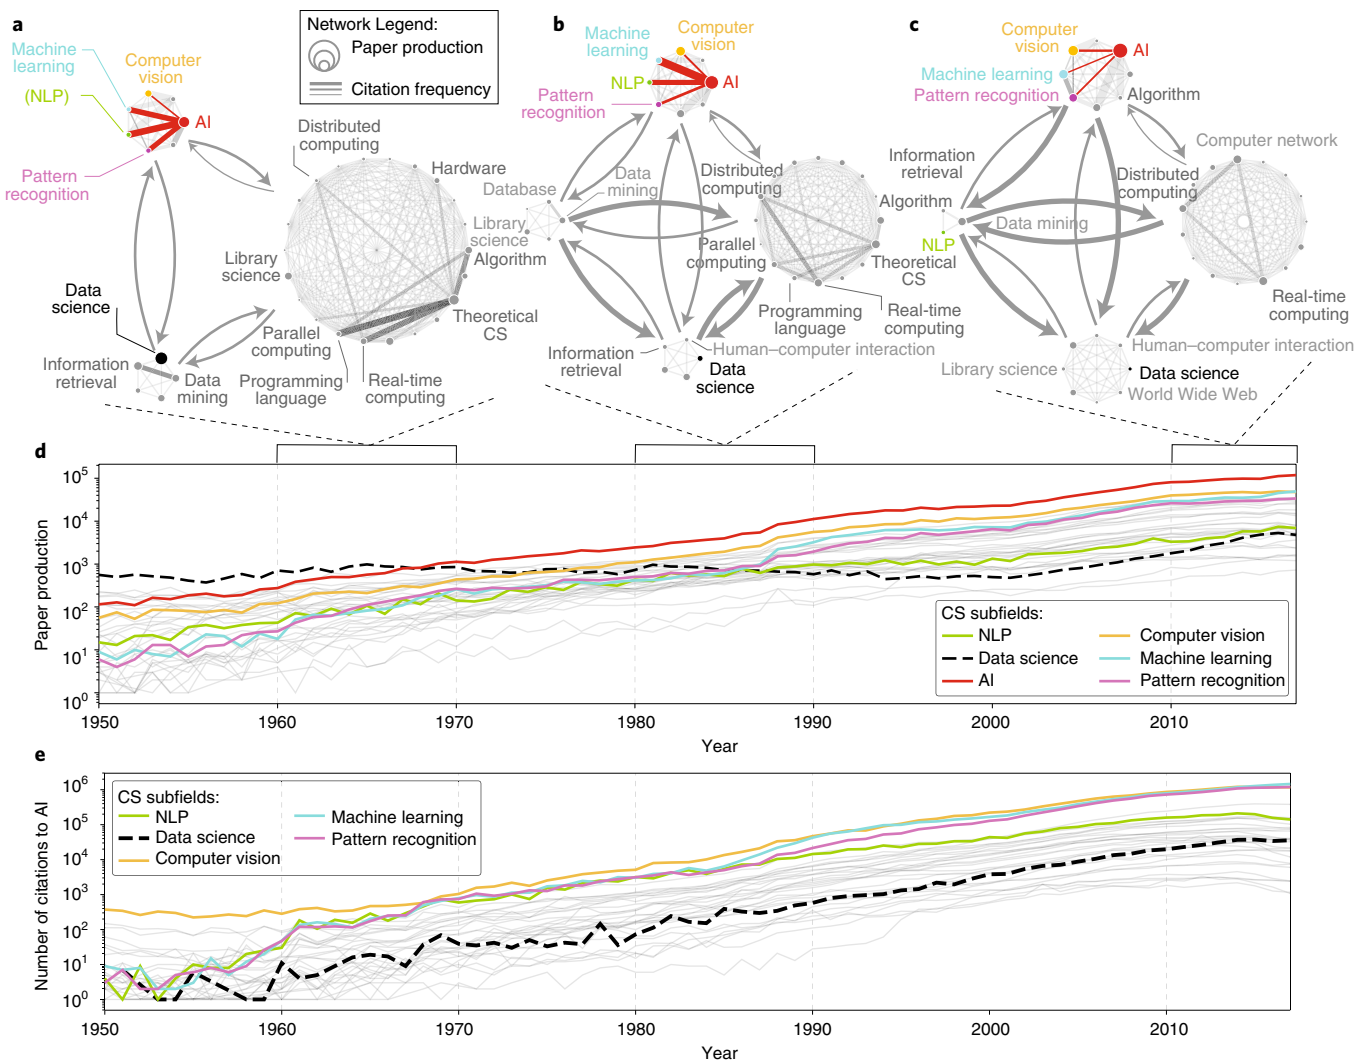


Fig. 1 | Citation patterns among CS subfields identify areas of AI-related research. a–c, We examine the rate of citations between CS subfields based on journal and conference publications from three different decades: the 1960s (**a**), the 1980s (**b**) and the 2010s through 2017 (**c**). For each network, the nodes (circles) correspond to CS subfields according to the MAG data, and the node size corresponds to the number of papers published in each subfield (note, the same paper may belong to multiple subfields). The width of the links connecting the nodes corresponds to the number of references made between papers published in those subfields. After constructing the complete network, we apply topological clustering⁴⁵ and report the number of citations made between these clusters using weighted arrows. Networks with labels for each subfield are provided in Supplementary Section 2. **d,** Annual paper production by CS subfield. Subfields related to AI are coloured, as well as data science (black) because of its notable decline in relative paper production. **e,** The annual number of references from papers in each CS subfield to papers in the AI subfield, and vice versa (that is, (subfield → AI) + (subfield ← AI)).

The paper production of CS subfields has varied over the past half-century. For example, data science has gradually diminished in relative paper production and theoretical CS has been replaced by increased focus on real-time and distributed computing. However, AI-related research areas have experienced steadily growing paper production since 1950 and account for the largest share of paper production in CS today (Fig. 1d).

Shaping the study of intelligent machines

Just as early myths and parables emphasized the social and ethical questions around human-created intelligence^{28–30}, today's intelligent machines provide their own interesting social questions. For example, how responsible are the creators, the manufacturers and the users for the outcomes of an AI system? How should regulators handle distributed agency^{1,31}? How will AI technologies reduce instances of human bias? As AI systems become more widespread^{1,21}, it becomes increasingly important to consider these social, ethical

and societal dynamics to completely understand the impact of AI systems^{9–11,32,33}. However, the developers of new AI systems are often separate from the scientists who study social questions. Therefore, we might hope to see increasing research interest between these fields of study and AI.

To investigate, we study the association between various academic fields and AI research through the referencing relationship of papers published in each academic field. External fields reference AI research for a number of reasons. Some fields, such as engineering or medicine, reference AI research because they use AI methods for optimization or data analysis. Other fields, such as philosophy, reference AI research because they explore its consequences for society (for example, moral and/or ethical consequences). Similarly, AI researchers reference other fields, such as mathematics or psychology, because AI research incorporates methods and models from these areas. AI researchers may also cite other fields because they use them as application domains to benchmark AI techniques.

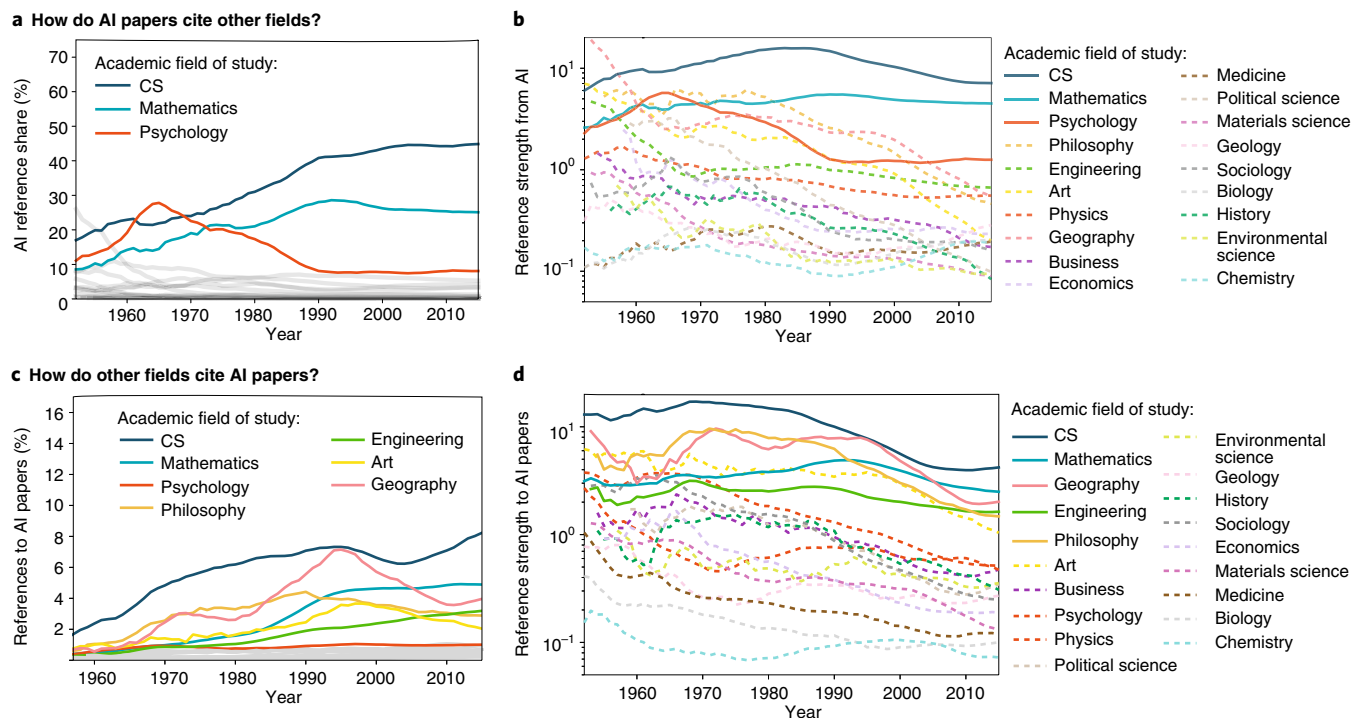


Fig. 2 | The referencing strength between AI and other sciences is declining. **a**, The share of references made by AI papers in each year to papers published in other academic fields. **b**, The reference strength (see equation (2)) from AI papers to papers published in other academic fields. **c**, The share of references made by each academic field to AI papers in each year. **d**, The reference strength from each other academic field to AI papers in each year. All lines are smoothed using a five-year moving average. In **b,d**, dashed lines indicate academic fields exhibiting lower reference strength than would be expected under random referencing behaviour in 2017.

In Fig. 2a,c, we examine the share of references made from AI papers to other fields, and from papers published in other fields to AI. The reference share from academic field A to field B according to

$$\text{share}_{\text{year}}(A, B) = \frac{\# \text{ refs from } A \text{ papers to } B \text{ papers in year}}{\# \text{ refs made by } A \text{ papers in year}} \quad (1)$$

controls for the total paper production of the referencing field over time, and has been used in other bibliometric studies³⁴. However, temporal changes in reference share may be explained by paper production in the referenced field; therefore, we consider another measure that also controls for the total paper production in the referenced field as well (Fig. 2b,d). We calculate the reference strength from field A to field B according to

$$\begin{aligned} \text{strength}_{\text{year}}(A, B) &= \frac{\left(\frac{\# \text{ refs from } A \text{ papers to } B \text{ papers in year}}{\# \text{ refs made by } A \text{ papers in year}} \right)}{\left(\frac{\text{no. of } B \text{ papers published from 1950 to year}}{\text{no. of papers published from 1950 to year}} \right)} \\ &= \frac{(\text{reference share from } A \text{ to } B \text{ in year})}{(B's \text{ share of all papers from 1950 to year})} \end{aligned} \quad (2)$$

A reference strength of $\text{strength}_{\text{year}}(A, B) > 1$ indicates that the rate of referencing from field A to field B is greater than would be expected by random referencing behaviour given the number of published papers in field B . Both reference share and reference strength capture the aggregate referencing behaviour between fields of study, but these calculations may obfuscate other dynamics from sub-communities within larger academic fields.

Before 1980, AI research made relatively frequent reference to psychology in addition to CS and mathematics (Fig. 2a). Controlling for the paper production of the referenced fields, we find that early AI's reference strengths towards philosophy, geography and art were comparable to the field's strength of association with mathematics (Fig. 2b) suggesting that early AI research was shaped by a diverse set of fields. However, AI research transitioned to strongly relying on mathematics and CS soon after 1987, which suggests an increasing focus on computational research.

How important is AI research to other academic fields? Unsurprisingly, CS, which includes all of the AI-related subfields in our analysis, steadily increased its share of references made to AI papers throughout the entire period of analysis (Fig. 2c). Surprisingly, mathematics experienced a notable increase in reference share to AI only after 1980. Meanwhile, several fields that are not often cited in today's AI research played an important role in the field's development, but may not have reciprocated this interest. For example, psychology was relatively important to early AI research, but psychology did not reciprocate as strong of an interest at any point from 1990 onwards (that is, $\text{strength}(\text{psychology}, \text{AI}) < 1$ in recent years). Instead, philosophy, art, engineering and geography have increased their share of references to AI papers up to 1995. On aggregate, when we control for AI paper production over time, we observe decreasing reference strength towards AI from all external academic fields. This suggests that other fields have difficulty keeping track of increasing AI paper production in recent decades (see Supplementary Fig. 3). This result may in part be explained by the increased complexity of AI-related research that is not relevant to the study of other scientific disciplines.

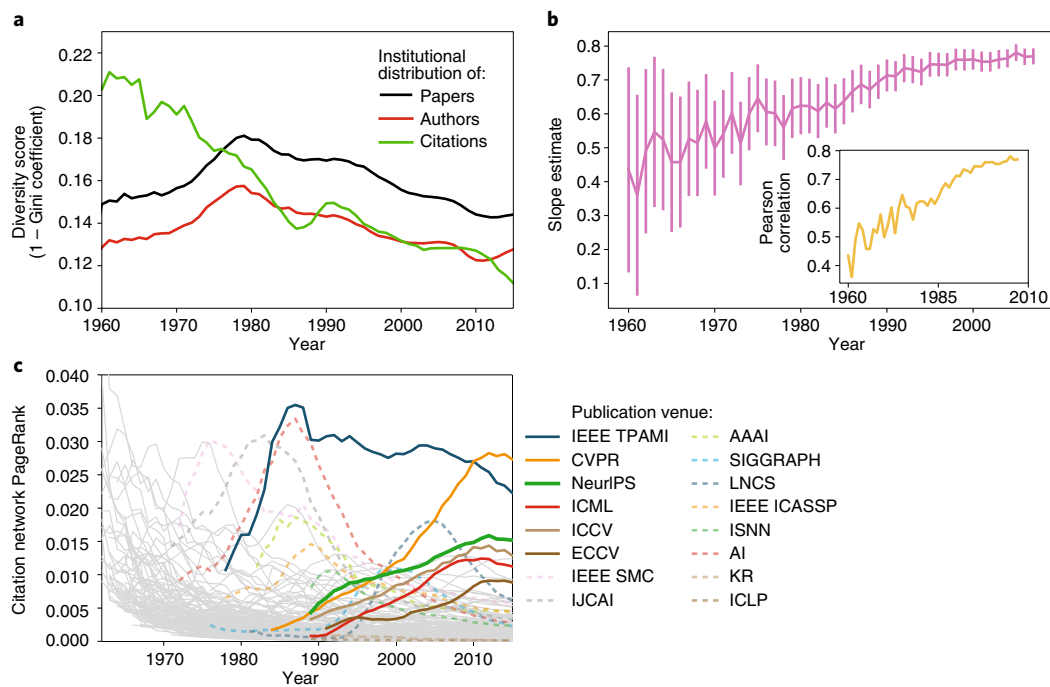


Fig. 3 | AI research is increasingly dominated by only a few research institutions and AI-specific conferences. **a**, The diversity of the annual distribution of all AI papers (black), AI authors (red) and all citations to AI papers (green) across research institutions according to the Gini coefficient. Example distributions of AI paper production and AI citation share are provided in Supplementary Section 3. **b**, To see whether preferential attachment explains citation dynamics, we include only AI papers with at least one citation and estimate the linear relationship between each research institution's cumulative citation count from 1950 to the institution's citation count in each year (see equation (3)). The model's slope estimation steadily rises throughout the period of analysis to around 0.70 as the model increasingly captures variance in the citation accumulation of institutions according to Pearson correlation (inset). The error bars are 95% confidence intervals for our estimate of the linear model's slope in each year (m in equation (3)). **c**, The PageRank of each publication venue for AI papers using the number of references from AI papers published in each venue to papers published in each other venue. The lines of notable publication venues are highlighted with colour. Dashed lines indicate venues whose PageRank has declined during the period of analysis. In all plots, lines are smoothed using a five-year moving average. More recent citation results may change as recent publications continue to accumulate citations. LNCS, Lecture Notes in Computer Science; ICLP, International Conference on Logic Programming; ISNN, International Symposium on Neural Networks; ICCV, International Conference on Computer Vision; ECCV, European Conference on Computer Vision; ICML, International Conference on Machine Learning; CVPR, Conference on Computer Vision and Pattern Recognition; IJCAI, International Joint Conference on AI; KR, Principles of Knowledge Representation and Reasoning.

The consolidation of AI research

How do leading research institutions shape AI research? On one hand, the prestige of an academic university can boost the scientific impact of CS publications³⁵. On the other hand, although scientific research is often undertaken at universities, major AI advances have emerged from industry research centres as well. For example, the AI start-up DeepMind received recent attention for their AlphaGo project³⁶ and Google has been acknowledged as a leader in the development of autonomous vehicles^{37–39}. With increased industrial and regulatory involvement, recent work suggests that areas of AI, including deep learning²¹, are undergoing a consolidation of research and deployment worldwide. While CS on the whole has become increasingly diverse⁴⁰, what can be said about AI research?

If the AI research community is experiencing a consolidation of influence, then what types of citation dynamics might indicate such a phenomenon? We investigate by examining the distribution of AI paper production and the distribution of citations made to AI papers by research institution (see Supplementary Section 3 for visualization of the distributions by decade). Since 1980, the diversity of AI paper production, authorship and citations to AI papers across institutions have decreased by 30% according to the Gini coefficient applied to annual distributions (Fig. 3a). Repeating this analysis for other academic fields, we find that this decreasing diversity is not simply a reflection of aggregate academic trends since most other

fields of study actually exhibit increasing diversity over time according to these metrics (see Supplementary Section 5).

This decrease in scientific diversity suggests that notable research 'hubs' may be forming (similar to the industry use of deep learning²¹). This type of hierarchical structure can occur when referencing between institutions is well modelled by preferential attachment⁴¹. If preferential referencing explains the citation dynamics within AI research, then the proportion of citations gained by a research institution in each year will be proportional to the institution's total accumulation of citations. Figure 3b reports estimates of the slope m for the model

$$\log_{10}(\# \text{ of citations}) = m \times \log_{10}(\text{cumulative } \# \text{ of citations}) + b \quad (3)$$

as well as 99% confidence intervals for those slope estimates using linear regression. Both the annual slope estimates and the performance of this model (see inset) rise steadily throughout the period of analysis. Combined, this evidence suggests that preferential referencing may be occurring among AI research institutions.

How have AI publication practices changed over time to enable preferential referencing? To investigate, we calculate the PageRank⁴² of each AI publication venue—including both academic journals and conferences—from the references of the AI papers published by each venue in each year (Fig. 3c). Publications venues with larger PageRank are more central to AI research. In the

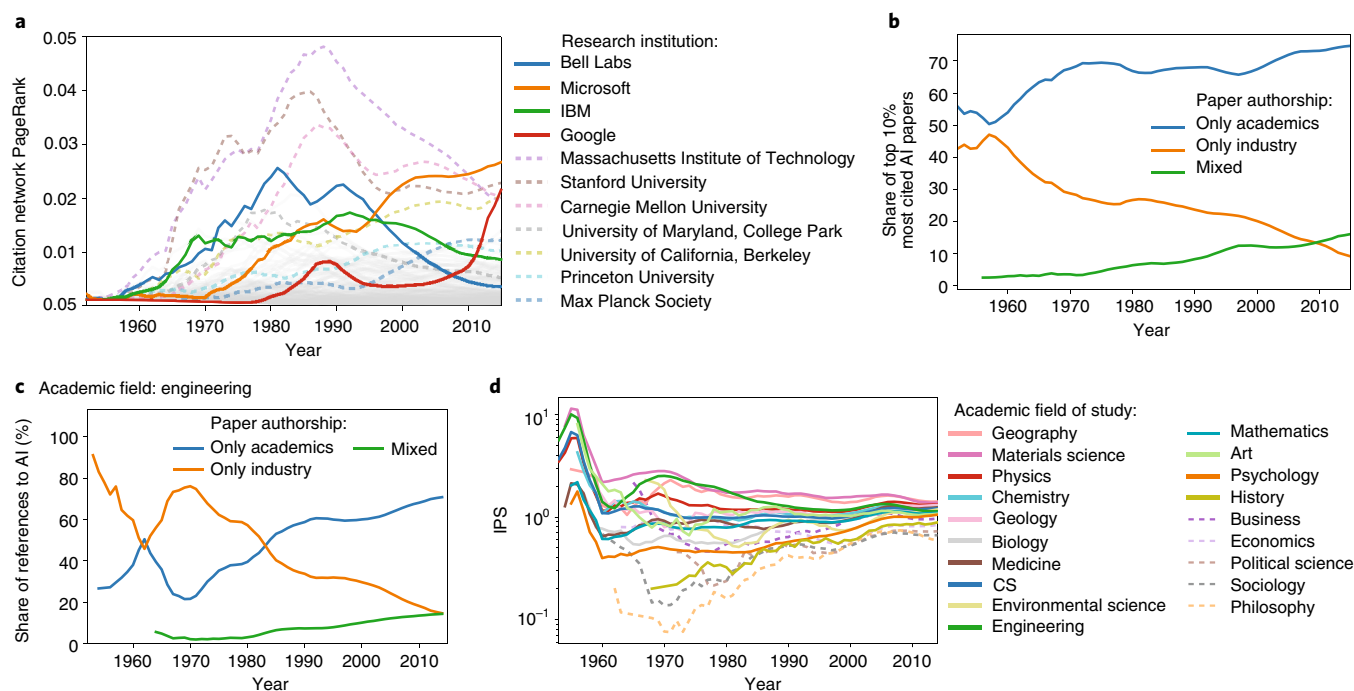


Fig. 4 | Industry is increasingly central to AI research, but industry-authored AI papers are referenced less often by other academic fields. **a**, The PageRank of each research institution using the number of references from AI papers published by each institution to papers published by each other research institution. The lines of notable research institutions are coloured for visualization. Dashed lines indicate academic institutions while solid lines indicate industry. **b**, The share of the top 10% most cited AI papers published in each year with academic-only, industry-only and mixed authorship. **c**, Similarly to **b**, we examine the referencing behaviour of engineering towards AI papers according to the authorship of the AI papers. Analogous plots for each other academic field are provided in Supplementary Section 4. **d**, Generalizing on **c**, the IPS calculated from each academic field's referencing behaviour towards AI papers (see equation (4)). The solid (dashed) lines indicate fields that reference AI papers with industry-only authorship more (less) than would be expected according to random referencing behaviour. In all plots, lines are smoothed using a five-year moving average.

late 1980s, several specific conferences, including the Conference on Computer Vision and Pattern Recognition, the Conference on Neural Information Processing Systems and the International Conference on Machine Learning, rise in prominence, while more general AI conferences, including the National Conference on Artificial Intelligence and the International Joint Conference on Artificial Intelligence, decline in prominence for AI researchers. Meanwhile, very few academic journals maintain high citation PageRank with the exception of the IEEE Transactions on Pattern Analysis and Machine Intelligence, which remains one of the most central publication venues for AI research.

If preferential referencing is producing research hubs, then which research institutions enjoy a privileged role in the AI research community? To investigate, we calculate the citation PageRank of each institution from the references of the AI papers published by each institution in each year (Fig. 4a). Before 1990, the most prominent research institutions were academic, including the Massachusetts Institute of Technology, Stanford University and Carnegie Mellon University, and included only a few industry-based research institutions, such as Bell Labs and IBM. However, the late 1980s again marks a transition point that reshaped the field. While universities dominate scientific progress across all academic fields⁴³, industry-based organizations, including Google and Microsoft, are increasingly central to modern AI research, and the PageRank scores of academic institutions are on the decline. Chinese research institutions at today's forefront of AI research are notably absent from Fig. 4a because their rise in prominence is recent in the 65-year time span of our analysis. However, the increasing prominence of Chinese research institutions, as well as other non-US-based institutions, is apparent when focusing on recent years (see Supplementary Section 8).

While academia has remained the largest source of AI papers throughout the entire period of analysis, the increased presence of industry can be seen from the authorship of AI papers over time (Fig. 4b). Out of the 10% of AI papers with the most citations after 10 years, the relative number of papers with industry-only authorship is on the decline. Meanwhile, collaborations between academia and industry are becoming more abundant.

How are other fields of study responding to the increased presence of industry in AI research? As an example, references from engineering showed preference for AI papers with industry-only authorship until the late 1980s, which is contrary to the aggregate trend (Fig. 4c; and see Supplementary Section 3 for similar plots for all academic fields). Similar to reference strength, temporal changes in a field's preference for AI papers with industry authorship (that is, at least one author has an industry affiliation) may result from the abundance of industry-based AI paper production over time. Therefore, we examine each field's industry preference score, which is given for field *A* by

$$\text{IPS}_{\text{year}}(A) = \frac{(\text{ref. share of } A \text{ to industry AI papers})}{(\text{industry share of AI papers from 1950 to year})} \quad (4)$$

Here, an AI paper has industry authorship if at least one co-author has an affiliation with an industry-based institution. Fields with $\text{IPS}(A) > 1$ exhibit stronger preference for industry AI papers than would be expected under random referencing behaviour towards AI papers. Academic fields that may be interested in the application of AI technology, such as materials science, engineering, chemistry and physics, tend to have greater preference for industry AI papers. However, many of the social sciences and fields that

study social and societal dynamics, such as sociology, economics, philosophy and political science, tend to have lower preference for industry AI papers.

Discussion

Humanity's long-standing quest²⁸ for AI is rapidly advancing in areas such as vision, speech and pattern recognition. However, as we deploy AI systems, their complete impact includes their social, ethical and societal implications in addition to capabilities and productivity gains. Understanding these implications requires an ongoing dialogue between the researchers who develop new AI technology and the researchers who study social and societal dynamics. Therefore, it is concerning to find a gap between AI research and the research conducted in other fields (Fig. 2).

AI paper production has increased quickly and steadily throughout the past half-century (Fig. 1), which suggests that the remarkable and seemingly sudden progress in AI is rooted in decades of research. Although AI research found as much early inspiration in psychology as CS and mathematics, it has since transitioned towards computational research. Conversely, several other academic fields are dedicating relatively more references to AI research. For example, engineering and mathematics research cite AI papers with increasing relative abundance throughout the period of analysis—making more frequent references to AI papers than would be expected under random referencing behaviour (Fig. 2c,d). However, the decreasing reference strength towards AI papers that we observe on aggregate suggests that most researchers are unable to keep up with the explosion of AI paper production (Fig. 2d). These findings may help explain why recent AI technologies have only recently revealed important (and largely unintentional) social consequences, such as racial bias in facial recognition software^{16,17}, the ethical dilemmas that have arisen from autonomous vehicles⁶ and income inequality in the age of AI^{18–20}. If current trends persist, then it may become increasingly difficult for researchers in any academic fields to keep track of cutting-edge AI technology.

The bibliometric gap between AI and other sciences grew with the advent of AI-specific conferences and the increased prominence of industry within AI research. In general, CS conferences can bolster the importance of publications⁴⁴ and enable major players to disproportionately influence the entire area of research⁴⁰. Although CS is becoming more diverse on the whole⁴⁰, the scientific impact of AI research institutions is becoming less diverse (Fig. 3a). In particular, Microsoft and Google have taken away the central role from universities according to citation PageRank (Fig. 4a), perhaps through preferential referencing of publications within AI (Fig. 3b).

This transition towards industry is challenging for studying the social and societal dynamics of AI technologies. Social science research is less likely to reference AI publications with authors who have industry-based affiliations. Combined with AI's decreasing reference strength towards social sciences, these observations suggest that this gap between research areas will continue to grow. The fields that study social bias, ethical concerns and regulatory challenges may be ignorant of new AI technology—especially when deployed in industry. While our interpretation of these results is speculative, we believe that our observations may highlight an important dynamic within the AI research community that merits further investigation.

Conclusion

The gap between social science and AI research means that researchers and policymakers may be ignorant of the social, ethical and societal implications of new AI systems. While this gap is concerning from a regulatory viewpoint, it also represents an opportunity for researchers. The academic fields that typically inform policymakers on social issues have the opportunity to fill this gap. While our study is a step towards this goal, further work may explicitly quantify the

social and societal benefits and consequences of today's AI technology as well as identifying the mechanisms that limit communication between research domains.

Received: 6 September 2018; Accepted: 15 January 2019;

Published online: 11 February 2019

References

1. Brynjolfsson, E. & Mitchell, T. What can machine learning do? Workforce implications. *Science* **358**, 1530–1534 (2017).
2. Kirilenko, A., Kyle, A. S., Samadi, M. & Tuzun, T. The flash crash: high-frequency trading in an electronic market. *J. Finance* **72**, 967–998 (2017).
3. Brogaard, J. et al. *High Frequency Trading and its Impact on Market Quality* Working Paper No. 66 (Northwestern University Kellogg School of Management, 2010).
4. Verghese, A., Shah, N. H. & Harrington, R. A. What this computer needs is a physician: humanism and artificial intelligence. *J. Am. Med. Assoc.* **319**, 19–20 (2018).
5. Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L. H. & Aerts, H. J. Artificial intelligence in radiology. *Nat. Rev. Cancer* **18**, 500–510 (2018).
6. Bonnefon, J.-F., Shariff, A. & Rahwan, I. The social dilemma of autonomous vehicles. *Science* **352**, 1573–1576 (2016).
7. *The Road to Zero: A Vision of Achieving Zero Roadway Deaths by 2050* (National Safety Council and the RAND Corporation, 2018).
8. Russell, S., Hauer, S., Altman, R. & Veloso, M. Ethics of artificial intelligence. *Nature* **521**, 415–416 (2015).
9. Rahwan, I. Society-in-the-loop: programming the algorithmic social contract. *Ethics Inf. Technol.* **20**, 5–14 (2018).
10. Crandall, J. W. et al. Cooperating with machines. *Nat. Commun.* **9**, 233 (2018).
11. Taddeo, M. & Floridi, L. How AI can be a force for good. *Science* **361**, 751–752 (2018).
12. Miller, A. P. Want less-biased decisions? Use algorithms. *Harvard Business Review* <https://hbr.org/2018/07/want-less-biased-decisions-use-algorithms> (2018).
13. Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J. & Mullainathan, S. Human decisions and machine predictions. *Q. J. Econ.* **133**, 237–293 (2017).
14. Naik, N., Kominers, S. D., Raskar, R., Glaeser, E. L. & Hidalgo, C. A. Computer vision uncovers predictors of physical urban change. *Proc. Natl Acad. Sci. USA* **114**, 7571–7576 (2017).
15. Erel, S. L. H. T. C., Isil & Weisbach, M. S. Could machine learning help companies select better board directors? *Harvard Business Review* <https://hbr.org/2018/04/research-could-machine-learning-help-companies-select-better-board-directors> (2018).
16. Buolamwini, J. & Gebru, T. Gender shades: intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency* 77–91 (MLR, 2018).
17. Buolamwini, J. How I'm fighting bias in algorithms. *TED Talks* https://www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms (2016).
18. Frank, M. R., Sun, L., Cebrian, M., Youn, H. & Rahwan, I. Small cities face greater impact from automation. *J. R. Soc. Interface* **15**, 20170946 (2018).
19. Frey, C. B. & Osborne, M. A. The future of employment: how susceptible are jobs to computerisation? *Technol. Forecast. Soc. Change* **114**, 254–280 (2017).
20. Acemoglu, D. & Restrepo, P. Robots and jobs: evidence from US labor markets (National Bureau of Economic Research, 2017).
21. Klinger, J., Mateos-Garcia, J. C. & Stathoulopoulos, K. Deep learning, deep change? Mapping the development of the artificial intelligence general purpose technology. Preprint at <https://arxiv.org/abs/1808.06355> (2018).
22. Sinha, A. et al. An overview of Microsoft Academic Service (MAS) and applications. In *Proc. 24th International Conference on World Wide Web* 243–246 (ACM, 2015).
23. Effendy, S. & Yap, R. H. Analysing trends in computer science research: a preliminary study using the microsoft academic graph. In *Proceedings of the 26th International Conference on World Wide Web Companion*, 1245–1250 (International World Wide Web Conferences Steering Committee, 2017).
24. Hug, S. E. & Brändle, M. P. The coverage of Microsoft academic: analyzing the publication output of a university. *Scientometrics* **113**, 1551–1571 (2017).
25. Burd, R. et al. GRAM: global research activity map. In *Proc. 2018 International Conference on Advanced Visual Interfaces* 31 (ACM, 2018).
26. Fiala, D. & Tutokly, G. Computer science papers in web of science: a bibliometric analysis. *Publications* **5**, 23 (2017).
27. Russell, S. J. & Norvig, P. *Artificial Intelligence: A Modern Approach* (Pearson Education Limited, London, 2016).
28. McCorduck, P. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence* (CRC, Natick, 2009).
29. Kieval, H. J. Pursuing the Golem of Prague: Jewish culture and the invention of a tradition. *Mod. Jud.* **17**, 1–20 (1997).

30. Pollin, B. R. Philosophical and literary sources of Frankenstein. *Comp. Lit.* **17**, 97–108 (1965).
31. Floridi, L. Distributed morality in an information society. *Sci. Eng. Ethics* **19**, 727–743 (2013).
32. Plant, S. *Zeros and ones* (Doubleday Books, 1997).
33. David, A. H. Why are there still so many jobs? The history and future of workplace automation. *J. Econ. Perspect.* **29**, 3–30 (2015).
34. Sinatra, R., Deville, P., Szell, M., Wang, D. & Barabási, A.-L. A century of physics. *Nat. Phys.* **11**, 791 (2015).
35. Morgan, A. C., Economou, D., Way, S. F. & Clauset, A. Prestige drives epistemic inequality in the diffusion of scientific ideas. *EPL Data Sci.* **7**, 40 (2018).
36. Silver, D. et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484 (2016).
37. Bergholz, R., Timm, K. & Weisser, H. Autonomous vehicle arrangement and method for controlling an autonomous vehicle. US patent 6,151,539 (2000).
38. Pilutti, T. E., Rupp, M. Y., Trombley, R. A., Waldis, A. & Yopp, W. T. Autonomous vehicle identification. US patent 9,552,735 (2017).
39. Herbach, J. S. & Fairfield, N. Detecting that an autonomous vehicle is in a stuck condition. US patent 8,996,224 (2015).
40. Pham, M. C., Klamma, R. & Jarke, M. Development of computer science disciplines: a social network analysis approach. *Soc. Netw. Anal. Min.* **1**, 321–340 (2011).
41. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
42. Page, L., Brin, S., Motwani, R. & Winograd, T. *The PageRank Citation Ranking: Bringing Order to the Web* (Stanford InfoLab, 1999).
43. Larivière, V., Macaluso, B., Mongeon, P., Siler, K. & Sugimoto, C. R. Vanishing industries and the rising monopoly of universities in published research. *PLoS ONE* **13**, 1–10 (2018).
44. Freyne, J., Coyle, L., Smyth, B. & Cunningham, P. Relative status of journal and conference publications in computer science. *Commun. ACM* **53**, 124–132 (2010).
45. Newman, M. E. Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* **74**, 036104 (2006).

Acknowledgements

The authors would like to thank E. Moro and Z. Epstein for their comments.

Author contributions

M.R.F. and D.W. processed data and produced figures. All authors wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s42256-019-0024-5>.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence should be addressed to I.R.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© Springer Nature Limited 2019

Supplementary Materials: The evolution of citation graphs in artificial intelligence research

Morgan R. Frank¹, Dashun Wang^{2,3}, Manuel Cebrian¹, and Iyad Rahwan^{1,4,5,*}

¹Media Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

²Kellogg School of Management, Northwestern University, Evanston, IL, USA

³Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL, USA

⁴Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA, USA

⁵Center for Humans and Machines, Max Planck Institute for Human Development, Lentzeallee 94, 14195, Berlin, Germany

*To whom correspondence should be addresses: irahwan@mit.edu

February 12, 2019

1 Microsoft Academic Graph Fields of Study




Figure 1: Annual paper produce by top-level field of study. In the legend, fields of study are ordered according to their share of overall paper production in the final year of analysis.

The Microsoft Academic Graph (MAG) data assigns fields of study (FOS) to each publication in the dataset. FOS are selected from a hierarchical taxonomy of fields, including biology, mathematics, and, in particular, computer science as FOS at the top of the hierarchy. Figure 1 demonstrates the share of annual paper production assigned to each top-level FOS from 1900 to 2018. Computer science has risen to the fourth most productive FOS in the last few decades beginning around 1950.

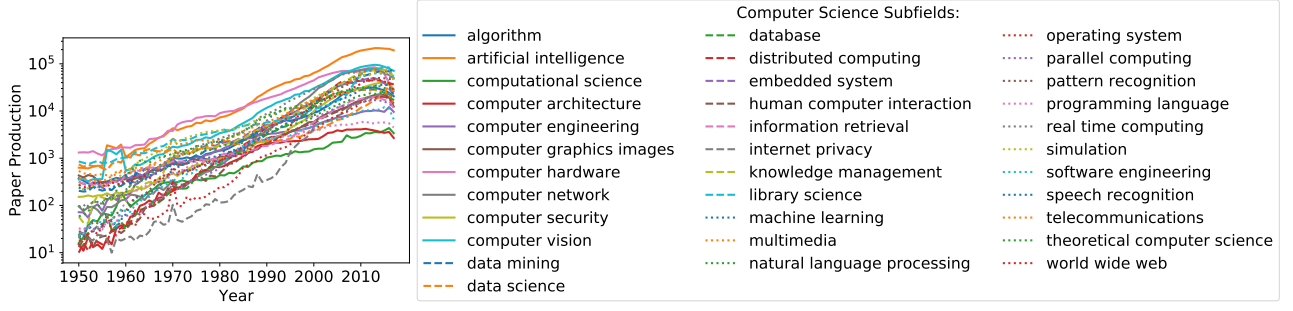


Figure 2: Annual paper production by subfield of Computer Science.




Figure 3: Annual artificial intelligence paper production.

1.1 Computer Science Subfields

Each FOS is divided into subfields. We are particularly interested in the subfields of Computer Science. Figure 2 demonstrates the annual paper production by Computer Science subfield.

2 Computer Science Subfield Citation Networks by Decade

Today, the phrases *machine intelligence* and *artificial intelligence* (AI) are most commonly used in reference for machine learning, but this was not always the case. Over the past 60 years, AI has been closely related to various Computer Science subfields, including Computer Vision, Machine Learning, Natural Language Processing, and Pattern Recognition.

To see this, we construct citation networks from the papers published in each Computer Science subfield (see Figures 4-10). In these networks, nodes are CS subfields and node size corresponds to paper production in that subfield (note: one paper may belong to multiple subfields). The connections between subfields have width proportional to the number of references made between papers in a pair of fields. After constructing this raw citation network, we apply community detecting (according to [1]) to identify clusters of Computer Science subfields based on how these fields reference each other. In the citation networks, we use color to identify these clusters and encode the number of references between clusters in the width of the arrows.

The strength of association between AI and Computer Vision, Natural Language Processing (NLP), Machine

Learning, and Pattern Recognition change dynamically over time. In fact, we can see the number of references between AI papers and NLP papers slowly diminish over the past several decades until NLP is actually contained in a separate community of Computer Science subfields. We also observe interesting dynamics around the subfield of Theoretical Computer Science and the emergence of the World Wide Web. Based on this analysis, we use papers in the following Computer Science subfields as a proxy for publications on AI-related fields of study: Artificial Intelligence, Machine Learning, Natural Language Processing, Computer Vision, and Pattern Recognition.

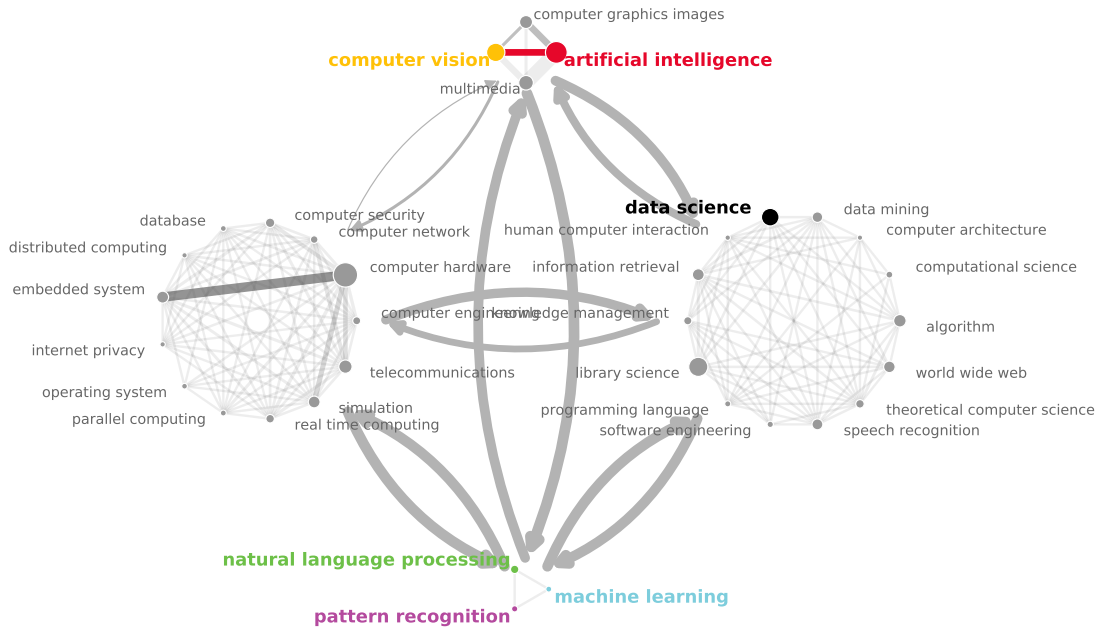


Figure 4: Citation network for Computer Science (CS) subfields constructed from papers published in the 1950's.

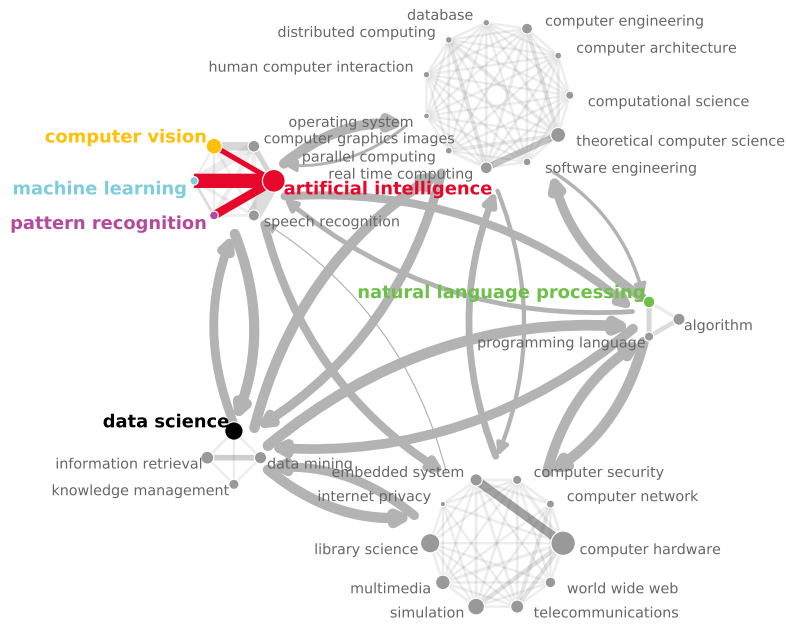


Figure 5: Citation network for Computer Science (CS) subfields constructed from papers published in the 1960's.

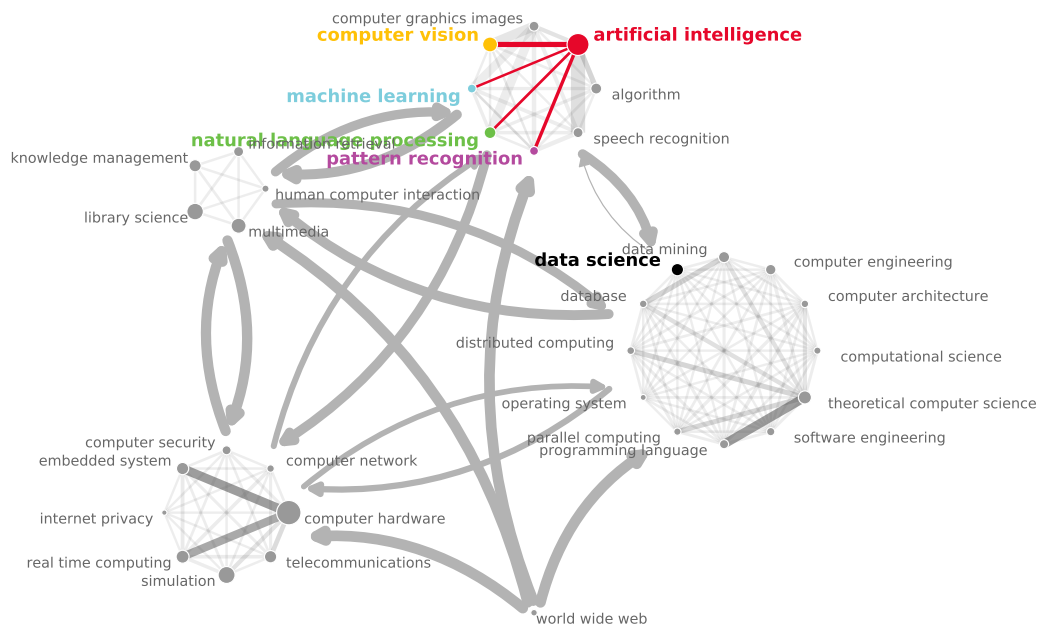


Figure 6: Citation network for Computer Science (CS) subfields constructed from papers published in the 1970's.

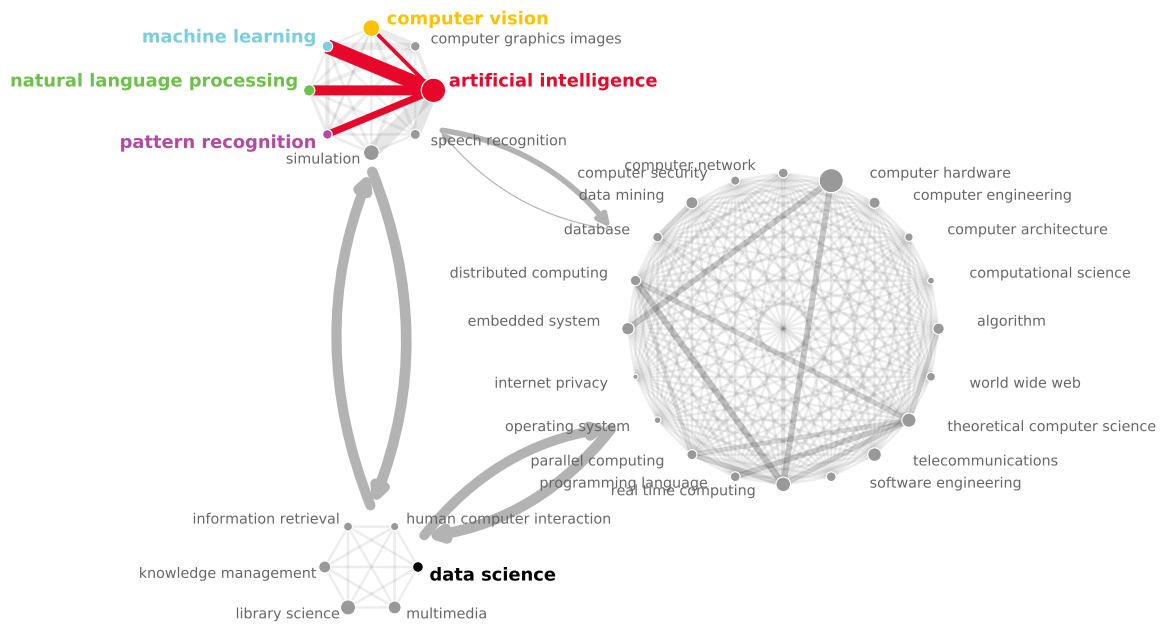


Figure 7: Citation network for Computer Science (CS) subfields constructed from papers published in the 1980's.

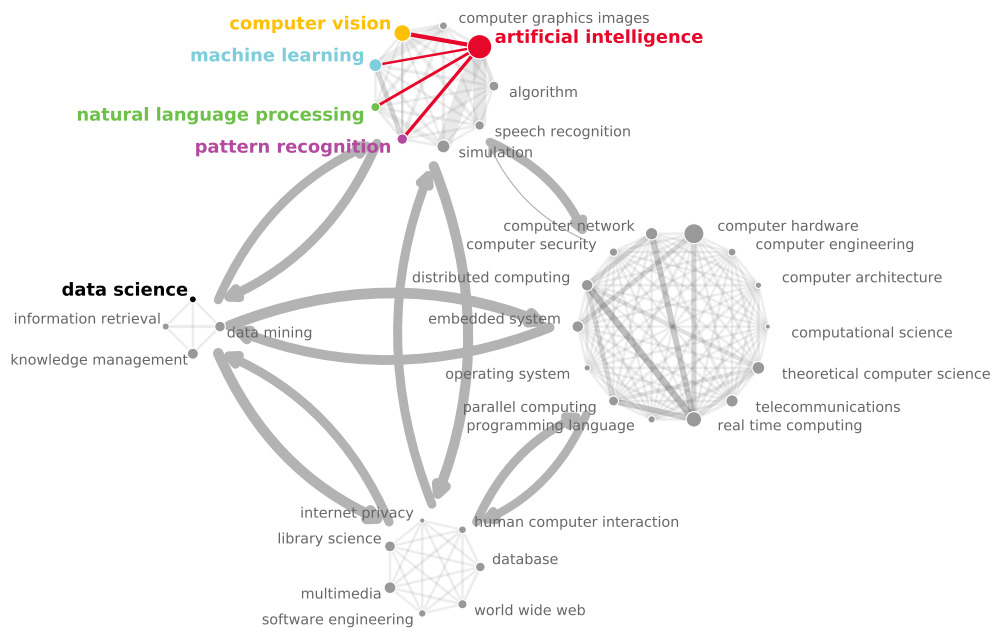


Figure 8: Citation network for Computer Science (CS) subfields constructed from papers published in the 1990's.

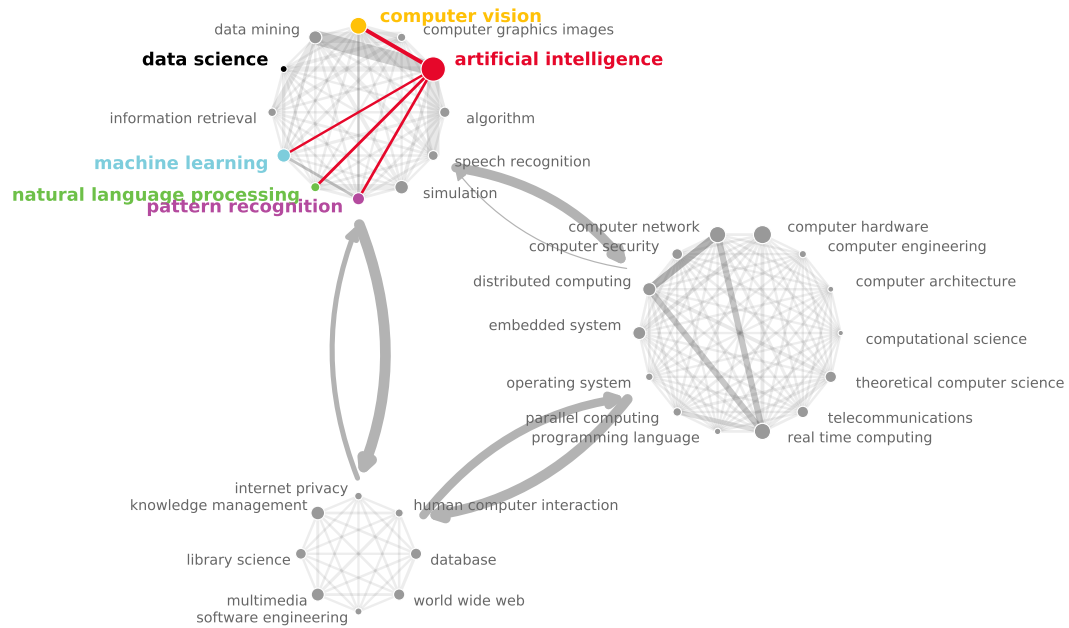


Figure 9: Citation network for Computer Science (CS) subfields constructed from papers published in the 2000's.

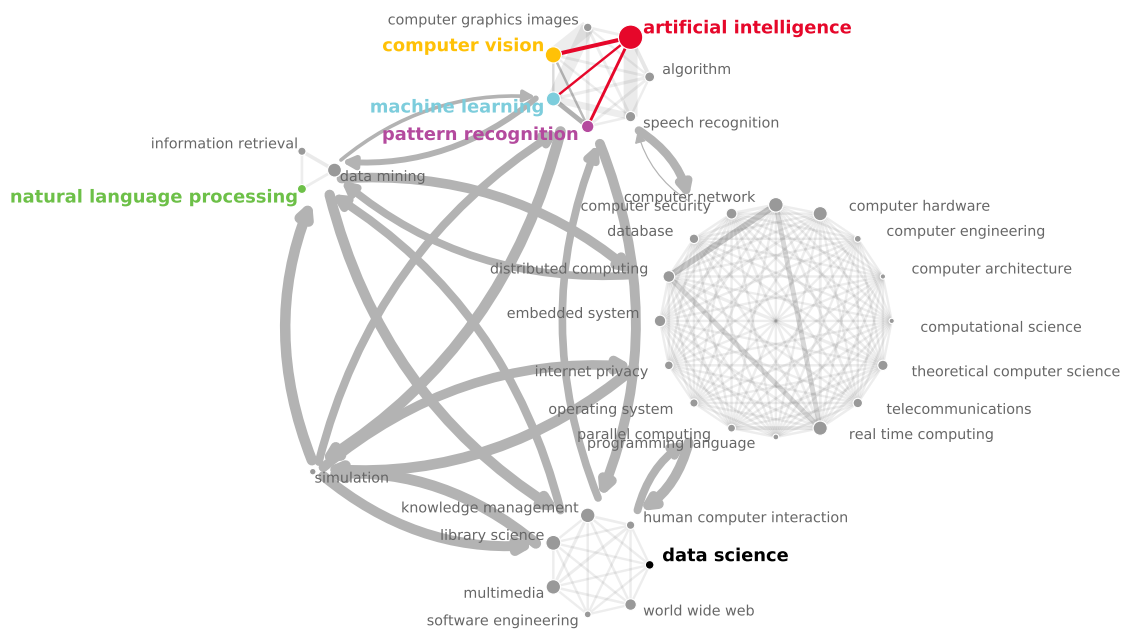


Figure 10: Citation network for Computer Science (CS) subfields constructed from papers published in the 2010's.

3 Distribution of Artificial Intelligence Productivity by Research Institution

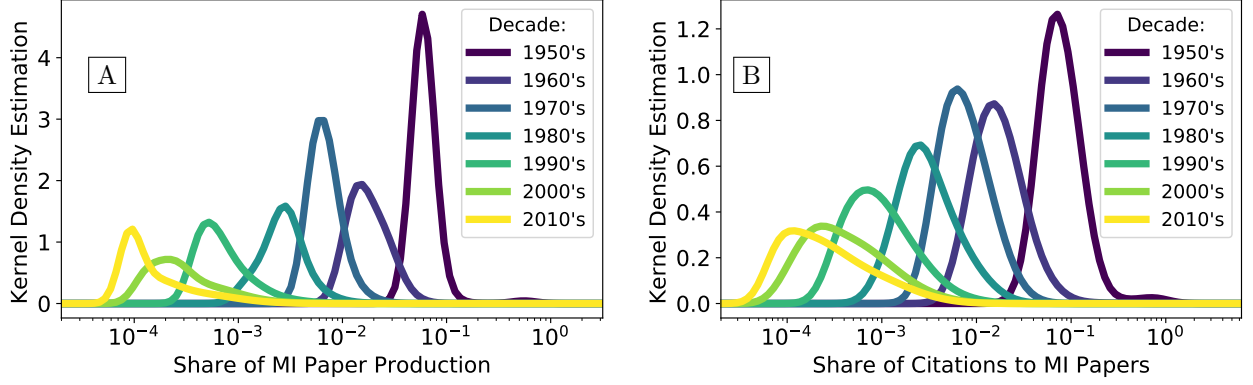


Figure 11: For each decade, we plot the distribution of **(A)** artificial intelligence paper production across each research institution producing at least one AI paper, and **(B)** the distribution of citations to AI research after 10 years across research institutions producing at least one AI publication. All curves are approximated using a Gaussian kernel density estimator.

4 The Preference of Academic Fields for Industry Artificial Intelligence Publications




Figure 12: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 13: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.

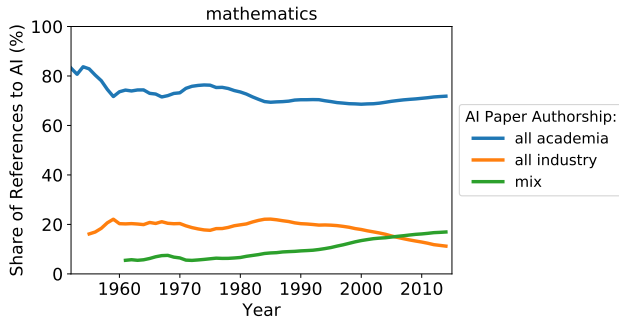


Figure 14: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.

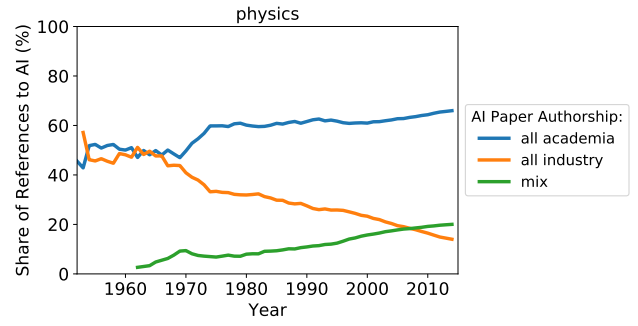


Figure 17: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 15: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 18: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 16: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 19: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 20: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 23: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 21: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 24: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 22: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 25: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 26: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 28: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 27: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.




Figure 29: The academic field's referencing behavior towards artificial intelligence papers according to the papers' authorship.

5 Bibliometric Diversity by Academic Field




Figure 30: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 31: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 33: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 32: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 34: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 35: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 37: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 36: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 38: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 39: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 41: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 40: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 42: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 43: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 45: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 44: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 46: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.




Figure 47: The diversity of the annual distribution of all papers (black), authors (red), and all citations to papers (green) across research institutions according to the Gini coefficient.

6 Authorship by Field of Study Over Time




Figure 48: For each Microsoft Academic Graph field of study, along with AI-related research as identified in the main text, we provide the average number of authors per paper published in each field in each year. The bump in AI research in the late 1950's is the result of a few publications with many co-authors combined with a reduced sample size in comparison to later years.

7 Authorship by AI Research Institution




Figure 49: In each year, we plot the average number of unique authors with at least one AI-related publication in that year across research institutions with at least one AI publication in that year. The grey area represents the 95% confidence interval for each year.

8 The Recent Rise of Chinese Institutions in AI Research




Figure 50: Ranking the prominence of AI research institutions in recent years. Similar to Figure 4a in the main text, we calculate the PageRank of AI research institutions from the references made to AI papers published by other AI research institutions. Here, we rank-order the AI research institutions in each year since 2000 and highlight the institutions exhibiting the greatest increase in rank. In addition to Google’s dramatic rise in PageRank rank, several academic institutions from around the world, but most notably in China, are rising in prominence within the AI research community. Gray lines represent AI research institutions that have fallen or remained constant in the AI prominence ranking in recent years.

References

- [1] Newman, M. E. Finding community structure in networks using the eigenvectors of matrices. *Physical review E* **74**, 036104 (2006).