



Munich Personal RePEc Archive

Artificial intelligence for science – adoption trends and future development pathways

Hajkowicz, Stefan and Naughtin, Claire and Sanderson,
Conrad and Schleiger, Emma and Karimi, Sarvnaz and
Bratanova, Alexandra and Bednarz, Tomasz

CSIRO

November 2022

Online at <https://mpra.ub.uni-muenchen.de/115464/>
MPRA Paper No. 115464, posted 01 Dec 2022 07:13 UTC

PERSONAL COPY OF THE REPORT

Artificial intelligence for science

Adoption trends and future development pathways

Citation

Hajkowicz S, Naughtin C, Sanderson C, Schleiger E, Karimi S, Bratanova A, Bednarz T (2022). Artificial intelligence for science – Adoption trends and future development pathways. CSIRO Data61, Brisbane, Australia.

Contents

1	Introduction.....	1
2	Artificial intelligence – Why now?.....	3
2.1	A turbulent history	3
2.2	Why is now different?.....	4
3	The science sector.....	7
3.1	What is the science sector?.....	7
3.2	The global science sector.....	7
3.3	The Australian science sector.....	10
3.4	Productivity decline.....	14
4	Artificial intelligence and knowledge discovery.....	17
4.1	Enablement or transformation?.....	17
4.2	Case studies – Artificial intelligence applications for science.....	18
5	Science domain adoption trends.....	23
5.1	Data sources and methods (bibliometric analysis).....	23
5.2	Publishing intensity and volumes – all research fields.....	26
5.3	Adoption trends in application domains	26
5.4	Artificial intelligence technology diffusion trends.....	32
5.5	Trends in artificial intelligence technologies.....	34
6	Future development pathways	37
6.1	Software, hardware and open access resources.....	37
6.2	The quest for better data.....	38
6.3	Education, training and capability uplift.....	39
6.4	Towards collaborative artificial intelligence.....	42
6.5	Artificial intelligence workforce diversity.....	44
6.6	The rise of ethical expectations and regulations.....	45
7	Conclusion.....	47
8	References.....	48
	Appendix A – Science occupations.....	56
	Appendix B – Artificial intelligence phrases.....	57

1 Introduction

We are currently amid the largest surge, arguably ‘boom’, in the application and development of artificial intelligence (AI) for scientific research in history. Scholarly publications, patents, education, training, salaries, research activity and investment are increasing at unprecedented rates. We may now be on the steepest part of the adoption and development curve. This is happening across the entire economy. Practically all industry sectors, advanced economies, professions and world regions are seeing rapid uptake of AI. The science sector is no exception. There’s a worldwide competitive race, and collaborative movement, to develop AI capability [8]. Many scientists, and science organisations, are aiming to uplift AI capability.

We are seeing steep rates of adoption in science because, in many well-publicised cases, AI is improving the speed, cost-effectiveness, quality and safety of scientific research. In some cases, the benefits from using AI are transformational; scientists have been able to solve problems hitherto beyond reach. There is a hope that AI can provide a much-needed productivity boost for science. AI may help scientists address humanity’s greatest challenges such as climate change, pollution, resource scarcity and infectious diseases. However, not all AI projects have met the expectations of scientists. Sometimes AI projects can be complex, costly, time-consuming and labour-intensive with limited results. The pathway to AI enablement, which most science organisations have embarked upon, is both rewarding and challenging.

This report has been prepared to help science managers, science organisations and investors understand plausible development pathways for AI. Our aim is to describe how AI has changed science and what the future may hold. We hope this will help science-sector workers make informed decisions about how they prepare for an AI-enabled future. Such decisions may be about investment, divestment, capability uplift, education, training and organisational design. We think most of the world’s science and research organisations are currently working through these, and related, issues as they seek to harness the opportunities and mitigate the risks of AI technology.

Other research institutes examining the impact of AI on science and exploring related issues include:

- The Alan Turing Institute. The institute was awarded £38.8 million (\$70.6 million) in 2018 for a 5-year research program on ‘AI and Data Science for Science, Engineering, Health and Government’. This applied-research program aims to understand and accelerate productive application of AI within these sectors [9].
- The ‘Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery’ (AI4SD) program aims to explore and demonstrate how AI technologies can boost discovery in all fields of research [10]. It is funded by the United Kingdom Engineering and Physical Sciences Research Council. A recent AI4SD conference was held at Chilworth Manor in the United Kingdom (and online) during 1–3 March 2022.
- The Organisation for Economic Co-operation and Development (OECD). Under the broader umbrella of the AI Policy Observatory, the OECD held a conference on ‘artificial intelligence and the future of science’ from 29 October to 5 November 2021. The conference examined the science productivity slump and the extent to which AI may provide a solution [11]. At this conference experts from across the globe presented data showing declining productivity in the science sector and discussed improvements possible via AI technology [12, 13].

- The Stanford University Human-Centered AI Institute. This institute prepares the AI index and associated reports, which are regularly updated to provide a comprehensive dataset on AI technology uptake and adoption across all fields of science, industry and society [14]. This institute also has a lead role in the 'One Hundred Year Study of Artificial Intelligence' (AI100), which provides an 'insider' perspective about the past, present and future trajectory of AI development. The most recent AI100 report was published in 2021 and is titled 'Gathering Strength, Gathering Storms' [15], indicating enhanced AI capability but coming with increased risks and complexity.
- The Argonne Laboratory AI for Science Project. The Argonne Laboratories at Oak Ridge and Berkeley hosted 'town hall' meetings for over 1,000 scientists during July to October 2019 about the use of AI, big data and high-performance computing. The findings were captured in the 'AI for Science' report written by over 80 authors. This report provides a detailed account about the state of the art, grand challenges, advances over the next decade, accelerating development and expected outcomes from AI application in nine major fields of scientific research [16].
- The University of Adelaide Australian Institute for Machine Learning (AIML). Along with the Australian Strategic Policy Institute, the AIML recently published a report titled 'Artificial intelligence: Your questions answered' [17], which examines issues of development, adoption and adaptation to AI technologies in Australia. It also examines issues of sovereign capability and why Australian industry often cannot buy AI 'off-the-shelf'. The AIML actively monitors and examines issues relating to AI in Australia.
- The 20-year community roadmap for AI research in the United States [18]. This document, and associated program of activity, is concerned with AI capability uplift in the United States out to the year 2040. It identifies research priorities in the areas of (a) integrated intelligence, (b) meaningful interaction, and (c) self-aware learning. The report makes recommendations about hardware and software resources, training and education, ethics, policy, workforce transitions and mission-led research for AI, amongst other matters [18].

This report contributes to the understanding about how AI will enable, and potentially transform, science from a global and Australian perspective. Our report opens with a brief history of AI and what makes now, the current boom cycle, different from the past. We then describe the global and Australian science sectors, highlighting science's productivity slump which AI can potentially help solve. We next present a bibliometric analysis of AI adoption across all science domains and patterns of AI science and technology development. Lastly, the report explores AI development pathways in science over the coming decade and examines the strategic implications for scientists and science organisations aiming to uplift capability for an AI-enabled future.

2 Artificial intelligence – Why now?

2.1 A turbulent history

There are many detailed accounts of AI's history. In this section we draw upon key publications [19-23] to provide a brief summary of AI's historic development to contextualise our subsequent analysis of its current status and future development pathways. We emphasise that AI isn't new to science; it has been on a long and turbulent journey with interest waxing and waning through history.

AI is a field of science that has been widely considered done-and-dusted a few times. However, AI has shown a strong ability to bounce back and re-establish scientific prominence. It is hard, if not impossible, to identify a start date for AI research. Scientists were publishing on concepts related to AI in the 1930s and 1940s. For example, Walter Pitts and Warren McCulloch [24] published a paper in 1943 about how artificial neurons can perform logical functions. One of their students, Marvin Minsky, later developed the 'stochastic neural analog reinforcement calculator' [25] which has evolved into AI neural networks

used widely in machine learning (deep learning) today.

One of the pivotal papers in AI was written by Alan Turing (Figure 1) and published in October 1950 [26]. This paper opens with the words 'I propose to consider the question, can machines think?'. It lays out the future challenges for AI to solve. Turing's question is still being asked [27] but remains unanswered. Turing died in 1954. The field of AI got its name at the Dartmouth Workshop of 1956. Organised by John McCarthy, Marvin Minsky, Nathaniel Rochester and Claude Shannon, this meeting brought together leading AI experts of the time [28]. Workshop attendees agreed to adopt 'artificial intelligence' as the name of their emerging research field. Naming AI helped connect a related set of technologies, concepts and theories. It helped formalise and establish an identity for a new field of science.

Investment and activity in AI escalated during the 1950s and, more so, in the 1960s. Significant advances occurred in the fields of natural language processing, automated reasoning, computational modelling, autonomous systems and robotics. The United States Defense Advanced Research Projects Agency (DARPA), the National Research Council and the United Kingdom Government were among the more notable investors in AI capability. The 1960s can be considered AI's first boom time. However, sentiment changed in the early 1970s. The first AI 'winter' lasted from 1974 to 1980. It was triggered by the Lighthill Report commissioned by the British Government and written



Figure 1. Alan Turing – Can machines think?

Data source: Sketch by Natata on Shutterstock.

by mathematician James Lighthill. The report was highly critical of AI's failure to achieve its 'grandiose objectives' [29]. The funding agencies mostly agreed with the view that AI had over-promised and under-delivered. The flow of resources for AI research was reduced to a trickle [23].

Despite the setbacks, the 1980s saw a return to boom times with the rise of expert systems and connectionism – an approach in the cognitive sciences which explains mental phenomena using artificial neural networks. The Japanese Government began aggressively funding AI through the fifth-generation computer project [30]. The United Kingdom and United States governments were soon to follow, again injecting substantial funds into a range of AI research initiatives in the early/mid 1980s. The business community became engaged as private companies boosted funding for AI research and development (R&D).

However, the boom times of the 1980s were followed by a second winter in 1987–1993. This was triggered by the business community which increasingly felt their investments in AI were failing to achieve commercial outcomes. Similar to the first winter, there was again a prevailing sentiment that AI had over-promised and under-delivered.

Expectations had risen higher than what was achievable. It was reported that by 1993 over 300 AI companies had shut down or gone bankrupt [23, 31]. This triggered a review of AI investment by governments; and again, AI R&D funds were suddenly and substantially reduced in the United States, the United Kingdom and across the globe.

Despite the two winters, the field of AI found a return to growth in the 1990s seeing the rise of new paradigms, tools, theories and applications. The 1990s saw the rapid growth of the internet, data and computing power. Since the late 1990s AI has remained on a strong growth trajectory. Research, investment, capability and adoption have continued to expand; there hasn't been a third winter. There are no apparent signs of a slowdown. In the current era AI is having a greater impact on scientific research than ever before.

2.2 Why is now different?

Many of the historic conditions which characterised the time periods leading up to previous AI winters, sometimes called AI 'springs', exist at the current time. There has been a huge and sudden boom in investment. There is much hype. Expectations are running high, and there is considerable mythology and confusion surrounding AI's capabilities and functions. If AI again fails to deliver on its perceived promises, it may enter another winter. However, there are reasons to believe the current era is different from the past. In this section we briefly explore what is different about the current epoch in the timeline of AI for science.

2.2.1 Greater depth and breadth of adoption

Compared to historic booms, today's AI surge has greater depth and breadth of technology penetration within diverse scientific fields, industry sectors, geographies, policy spheres and demographics. AI has gotten into practically everything everywhere. This is creating greater resilience for AI compared to historic boom–bust cycles. AI is too deeply embedded in too many places to suddenly lose relevance, as happened in the two winters of 1974–1980 and 1987–1993.

This embeddedness can be seen through publishing and patent trends. Our bibliometric analysis based on data from The Lens [32] reveals that 5.7% of all peer-reviewed research publications refer to AI in the title, abstract or keywords. This is up from 3.1% in 2017 and 1.2% in 2000.

In 2021 alone, 344,000 journal papers, books, book chapters and conference papers were published on the topic of AI [33]. In 2020 Google Scholar reported that AI attracted more citations than any other research field and five of the seven top-cited papers were on AI topics [1]. Patents for AI have also been increasing sharply. According to data from The Lens the number of published patents worldwide on the topic of 'artificial intelligence' rose from 11,000 in 2017 to 57,000 in 2021 representing an average year-on-year growth of 84% over the last 5 years [32]. This growth is happening in all world regions and most countries, with China being a standout: the number of peer-reviewed publications on AI from China now exceeds that from both the United States and Europe. Furthermore, the growth in AI publishing is happening in all industry sectors with sharp increases in the corporate, government, medical and other sectors [14].

In terms of expenditure, we are seeing sustained growth which is likely to continue over coming years. Most governments from advanced economies have announced and funded significant AI strategies, roadmaps, plans and policies. Canada was among the first of the OECD countries to commit to a national AI strategy in 2017 [34]. Since then over 700 AI policy and strategy initiatives have been developed across 60 countries and territorial jurisdictions [3, 4]. By late 2019, over \$86 billion in funding had been announced for AI initiatives [35]. Investment has continued to grow. In 2021 worldwide spending on AI products and services grew 15.2% year-on-year, reaching US\$341.8 billion. Growth of 18.8% is expected for 2022 with total spending forecast to exceed US\$500 billion per year by 2024 [36]. The share of this expenditure being invested in science and research is unknown, but due to the novelty and complexity of developing AI – e.g. training machine-learning algorithms – it is likely to be substantial.

It is also worth noting that AI has now found its way into people's day-to-day lives. Billions of people use, and increasingly depend upon, AI on a regular basis. Countless companies use AI technologies to provide goods and services to their customers. Before the turn of the century this was not the case. In the 1990s (and before) AI was a concept beyond the realm of most people's lived experience. In comparison, people today routinely interact with powerful AI through smart-phones, smart-cars and smart-speakers. This makes AI both tangible and practically useful. The contemporary widespread familiarity with AI makes it easier for today's research community to communicate its value proposition.

2.2.2 Hardware, software and data availability

Throughout history AI scientists have struggled to implement and test their ideas due to a lack of computational resources and/or lack of data to train machine-learning algorithms. Much was theorised but could not be proven nor developed. That is often why scientists hit a wall and could not realise their objectives. This boom is different. This boom comes with much better tools and much better data. Scientists can turn their ideas into technologies and technologies can be turned into consumer products with commercial value.

For example, in the last 10 years we have seen the rise of graphics processing units (GPUs) which are well suited to support parallel computing. The GPU has been transformative for AI. It has enabled low-cost high-power computing for a vast range of complex machine-learning challenges. Furthermore, cloud-based computing services are bringing these within reach. We are also seeing software tools and platforms, such as PyTorch, Tensorflow, Theano, MxNet, Microsoft Azure, and Amazon Web Services, make AI much more accessible to a broader cross-section of scientists; a subset of AI functionality is usable for those without highly developed and specialised AI skills. More recently we have seen code-free AI tools that allow users to perform AI functions via relatively simple graphic user interfaces (GUIs). These are likely to improve over time.

Another limiting factor for AI scientists in history has been the availability of data to train machine-learning algorithms. However, there's no shortage of data today. There are challenges about managing an overwhelming volume, variety and velocity of data. There are also challenges about verifying data, as well as challenges about handling private and confidential data. If these challenges can be addressed, today's scientists can have access to more data on every topic than ever before in history. The data comes from human internet usage, sensory systems and countless other rapidly expanding sources. These data are providing scientists with new opportunities to use AI to identify patterns, test hypotheses and make predictions.

2.2.3 Commercial drivers

While there have been surges in AI investment in the past, they do not come close to what is happening today. Private investment flows into AI have increased substantially over the past several years (Figure 2). Despite the pandemic, private investment in AI companies increased by a record high of 9.3% in 2020 year-on-year – which is above the 5.7% increase of 2019 – and exceeded US\$40 billion [14]. Venture capital investment in AI has also been growing compared to other areas of investment. According to the OECD, the share of venture capital investments in AI start-ups reached 20% of all venture capital investments in 2020, up from 3% in 2012 [37]. The number of venture capital deals in AI companies grew by 34% annually between 2012 and 2020 from 500 deals in 2012 to 3,900 deals in 2019 [37]. In 2020 Australia was ranked 11th in the world by the total amount of private investment in AI companies [14] with the United States, China, United Kingdom and Israel at the top of the list. With so much invested, AI activity by R&D providers and product developers is likely to be sustained for some time.




Figure 2. Private investment in artificial intelligence companies worldwide (billions of US dollars).

Data source: Stanford University Artificial Intelligence Index [14].

2.2.4 Improved scientific knowledge and technological capability

Lastly, today's surge in artificial intelligence comes with solutions and/or improvements to some – but by no means all – of the longstanding machine-learning problems, and with knowledge gaps filled which limited AI's development historically. For example, CSIRO re-implemented a classic machine-learning algorithm, Random Forest, enabling it to overcome the 'curse of dimensionality', which was brought on by today's larger and more detailed datasets [6]. Other algorithmic improvements include new regularisation techniques which modify the learning algorithm to improve the generalisability of the model plus improve its performance on unseen data [38].

These techniques can reduce the problem of 'overfitting' in machine learning which happens when the model is too closely matched to input data and is, therefore, unable to predict future observations accurately. We have also seen the emergence of robust optimisers such as Adam [39], RMSprop [40] and modification of the stochastic gradient descent (SDG) procedure [41]. These approaches speed up optimisation algorithms and generate higher quality solutions compared to earlier methods. The last 10 years have also seen the emergence of improved

backpropagation algorithms which improve the accuracy of artificial neural networks by finely adjusting mathematical weight functions. The recent transition to the Rectified Linear Unit (ReLU) activation function has substantially helped address the vanishing gradient problem; a longstanding challenge in the field of machine learning [14].

Solutions to these, and other, AI barriers have opened up entirely new avenues for continued problem solving and improvement of AI technologies. This means AI science has a greater chance of delivering on expectations. The future is likely to see continued discovery and innovation in the field of AI enabling the development of enhanced technological capabilities.

2.2.5 No slowdown in sight

There's so much momentum behind the current AI growth cycle it is hard to see it ending anytime soon. If AI were to experience another winter in one field of research – such as computer science – it is unlikely to be winter everywhere. The field of AI has become so large and diverse it is likely to be experiencing all four seasons in sub-fields, application domains, geographies and industry sectors at any one point in time.

3 The science sector

In this section we describe the size, structure and trends within the Australian and global science sectors. We examine how the sector has expanded and we examine the science productivity slump. It is likely that AI will play an important role in boosting science productivity and, in turn, economy-wide productivity.

3.1 What is the science sector?

The science sector is an interconnected, collaborative and dynamic global community with highly porous boundaries. It captures a diverse range of conceptual frameworks, paradigms, methodologies and cultural approaches to knowledge discovery. The Australian Academy of Science [42] says, ‘science can be thought of as both a body of knowledge (the things we have already discovered), and the process of acquiring new knowledge (through observation and experimentation—testing and hypothesising)’.

Science happens in households, communities, start-ups, large companies, government agencies, research institutions and universities. There are many ways of classifying scientific activity. The Australian Academy of Sciences identifies four broad and high-level categories of scientific research [42]:

- natural science – the study of living organisms and physical sciences which includes the study of the material universe
- social science – the study of human individuals, communities, societies and institutions and how they interact and behave
- formal science – the study of logic and mathematics
- applied science – the adaptation and use of existing scientific knowledge for industry and societal applications.

The United Nations recently conducted global consultation to define the science sector for statistical purposes [43]. This analysis identified the concept of ‘scientific and technological activities’ which includes three components of science activity: (a) research and experimental development, (b) scientific and technological education and training, and (c) scientific and technical services.

We have defined and conceptualised the science sector in the same manner in the subsequent analysis.

3.2 The global science sector

For some time, the world has been growing its science workforce and research spending (Figure 3). In 2018 the global research workforce was estimated by the United Nations at 9.33 million workers up from 8.01 million in 2014. This workforce is estimated to be growing over three times faster than population growth; increasing by 16% during the 5-year period 2014 to 2018 [44]. Data from the OECD show that the number of researchers per 1,000 employed persons has increased from 6.1 to 8.9 during 2000 to 2020 [2].

Expenditure patterns for R&D are indicative of aggregate global-level science spending. Spending on R&D is outpacing global economic growth, reaching 2.2% of global GDP by 2020 compared to 2.0% a decade ago. In the OECD it has grown faster, reaching 2.48% in 2019. Recent estimates suggest the world now spends over US\$1.7 trillion per year on research with 10 countries accounting for 80% of expenditure. The top five spenders in absolute terms are the United States, China, Japan, Germany and South Korea. The country spending the most on R&D relative to GDP is Israel, which invested 4.93% of GDP in R&D in 2019 [2].

Global technology corporations are another key source of expanding funds, and overall activity, for scientific research. This has been a relatively recent phenomenon and is especially relevant to AI science. During 2005–2018, global private-sector R&D spending rose from \$523.8 billion to \$1.1 trillion [45]. A NASDAQ report [46] reveals the R&D budgets of the top spenders in 2020: Amazon (\$62.3 billion), Alphabet (\$40.2 billion), Huawei (\$32.1 billion), Microsoft (\$28.1 billion), and Apple (\$27.3 billion). These trends suggest that private corporations are funding or doing much of the world’s scientific research. Given the business objectives of these companies, much of this R&D is likely within the fields of data science and AI. The entry of the private/corporate sector into R&D is impacting scientific research.

EXPENDITURE ON R&D US DOLLARS, TRILLIONS•

PERCENT OF GDP SPENT ON R&D




Figure 3. Increasing spending on R&D within the OECD.

Data source: OECD Statistics [2].

The bulk of the global research effort by publishing volume lies in the physical sciences, which contained 44% of all publishing in the year 2021. This is followed by the health sciences, life sciences and the social sciences and humanities, which contained 22%, 20% and 14%, respectively, of publishing in the same year. The relative level of publishing within these domains has remained relatively stable over history. When we look beneath the top-level research domains, we see that the fields of medicine, biochemistry genetics and molecular biology, engineering, social sciences and computer science account for over half of all research publishing (Figure 4). Medicine is by far the largest research field, accounting for one-fifth of all research publishing. Again, this has remained relatively stable over history.

The volume of research published across all fields of science has been increasing consistently for over 60 years (Figure 5). However, during the peak COVID-19 pandemic years of 2020 to 2022, global combined scientific publishing via journal papers, conference papers, books, book chapters and dissertations declined by 12%. This may have multiple causes, including: (a) the redirection of research effort towards urgent COVID-19 issues; (b) a switch to alternative non-traditional publishing venues; (c) decreased productivity of researchers working in lockdowns; and (d) reduced research funds in the university sector associated with reduced revenue due to COVID-19 disruptions.

The patterns and consequences of pandemic-related research activity contraction are explored in the research literature in greater depth [48, 49]. These studies indicate the possibility of long-term effects and the disproportionate impact on female researchers with young children.




Figure 4. Share of global peer-reviewed publishing by research field in 2022.

Data source: The Lens [32] and Scopus, Elsevier All Science Journal Classification [47]. Data sourced for 1 Jan 2022 to 20 Sept 2022.

NUMBER OF PUBLICATIONS THOUSANDS

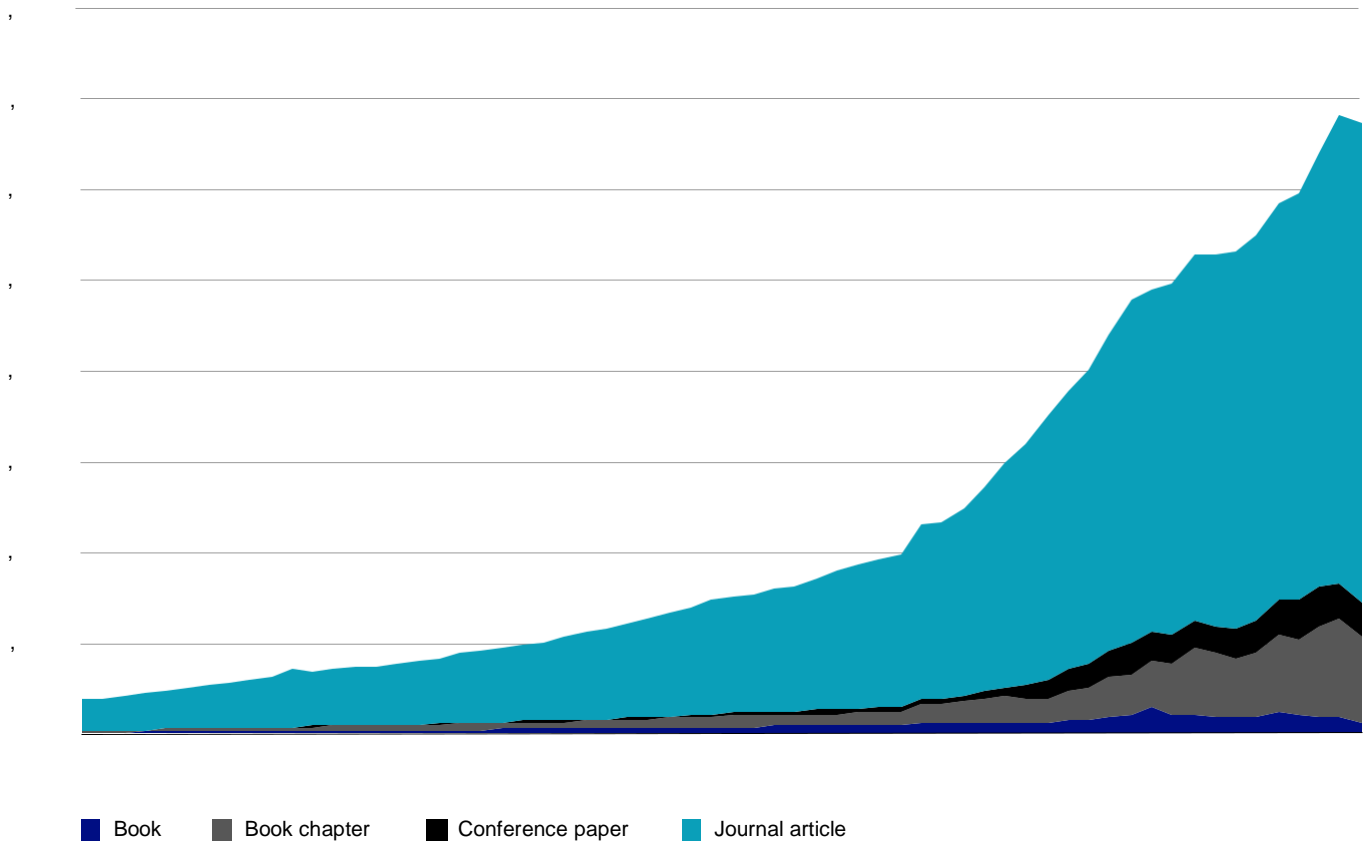


Figure 5. Worldwide peer-reviewed scientific publishing in all fields for 1960–2021.

Data source: The Lens [32] database of peer-reviewed scholarly publications.

Publishing venues may also be changing. During recent times, new open-access research-sharing platforms have risen quickly. For example, during 2016–2021 submissions of research papers to the arXiv pre-print server increased 60% from 113,380 to 181,630 [50]. The material is not subject to peer review and, therefore, comes with a significant quality and reliability caveat. However, it does allow for rapid sharing of research results and is being used extensively. In another example, the ‘Papers with Code’ platform now reports 62,856 research papers (as of 8 Jan 2022) [51]. The extent to which these alternative venues displace and/or complement traditional peer-reviewed publishing, and/or morph into new publishing models, remains to be seen. Peer review is still needed to separate knowledge from opinions. The question is whether new models can uphold this critical function whilst speeding up the process of knowledge sharing [52]. What’s evident at the current time is sizeable and rapidly increasing usage of the alternative venues by the science community.

3.3 The Australian science sector

In the Australian and New Zealand Industry Classification (ANZIC) the science sector fits under the industry sub-grouping ‘scientific research services’ [53]. In 2016, this industry contained 28,850 workers, roughly 0.3% of the Australian workforce. However, most scientists work within, and most science happens within, other industry sectors (Figure 6). The vast bulk of Australian scientists do not work in the science industry. They use science to problem-solve in other industries and societal spheres. The ‘science sector’ is spread widely across the entire economy.

Using the most recent population census for 2016, we identify 109,890 natural, physical and social scientists in Australia representing 1% of the total workforce [53]. We note a degree of uncertainty in this estimate due to the category groupings used by the Australian Bureau of Statistics. It is likely that we are including some non-scientists and excluding some scientists because there is no category for ‘scientist’ in our statistics; it requires aggregating and disaggregating other categories.

Nevertheless, we believe a reasonable approximation has been achieved using the occupation information in the census (Appendix A). From this we identify 11 sub-categories of the scientist profession (Figure 7). The two largest groupings are medical laboratory scientists and environmental scientists, which together account for about 30% of the scientific workforce [53].

At 22% of the total science workforce, Australia’s professional services sector contains the largest share of scientists [53]. It includes sub-industry groupings such

as management consulting, architectural, engineering, law, accounting, market research, veterinary and other services. This analysis suggests that much of AI for science (the focus of this report) is likely to occur within Australian industry as opposed to dedicated research organisations and universities. This is because most scientists work in industry. However, research organisations may be more focused on theoretical and early-stage developmental aspects of AI compared to industry scientists who may be focused on adoption and application.




Figure 6. The number of Australian scientists by industry grouping.
Data source: Australian Bureau of Statistics, 2016 Population Census [53].

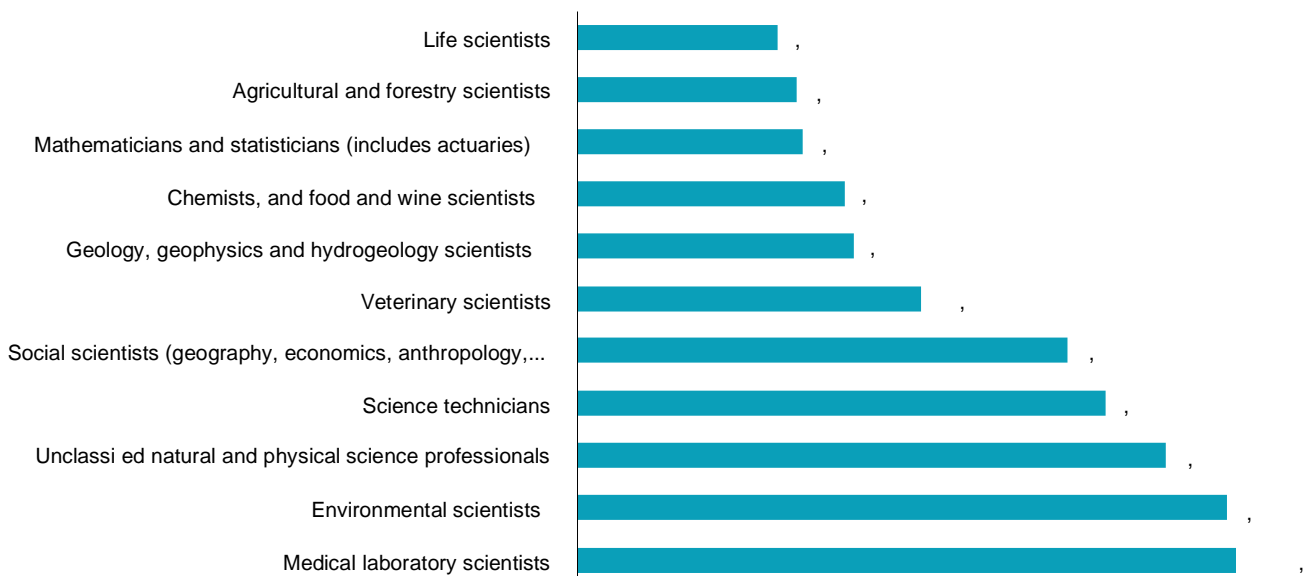


Figure 7. The number of Australian scientists by field of science.
Data source: Australian Bureau of Statistics 2016 Census (see Appendix A for methods and assumptions) [53].
Note: The social science field contained a small number of non-science occupations that could not be disaggregated.

R&D spending patterns recorded by the Australian Bureau of Statistics [54] provide insight into science investment patterns. In 2019–2020, total expenditure on R&D by business, government, higher education and not-for-profit sectors in Australia amounted to \$35.602 billion, representing an increase of \$2.54 billion (8%) over the preceding 2 years. Much of the growth in R&D spending comes from the higher education sector, which increased from \$11.24 billion in 2017–2018 to \$12.71 in 2019–2020, an increase of \$1.47 billion (13%). However, over the longer-term Australian R&D spending declined as a percentage of GDP from 2.1% in 2012 to 1.8% in 2020.

At the higher levels of the research field taxonomy, the publishing profile for Australian research has similar expression to global research. However, at a more granular level we can see that Australia has comparative specialisation in certain fields. To examine this, we calculate specialisation quotients at the second-level of the All Science Journal Classification (ASJC) by Scopus, Elsevier [47] which contains 26 unique research fields. The specialisation quotient is calculated as follows:

A specialisation quotient above 1 implies the research field is associated with greater output in Australia compared to the average for all research fields. This analysis reveals that over the 10-year period 2012–2021, the fields of psychology, health professions, earth and planetary science, nursing, agricultural and biological science and environmental science have higher levels of comparative specialisation in Australia (Figure 8). By comparison, Australia has lower levels of comparative specialisation in dentistry, physics and astronomy, mathematics, chemistry and materials science. At the even more granular (third) level of the research field hierarchy, the top 10 fields of research by specialisation for Australia include: emergency medical services; research and theory; community and home care; tourism, leisure and hospitality management; economic geology; chiropractics; occupational therapy; pharmacy; physical therapy, sports therapy and rehabilitation; and ecological modelling.

$$\begin{array}{l}
 \text{Specialisation quotient} \\
 \text{for a given research} \\
 \text{field in Australia}
 \end{array}
 =
 \frac{\text{Number of publications} \\
 \text{in research field by} \\
 \text{Australian authors}}{\text{Number of publications} \\
 \text{in research field by} \\
 \text{worldwide authors}}
 \div
 \frac{\text{Number of publications} \\
 \text{in all research fields by} \\
 \text{Australian authors}}{\text{Number of publications} \\
 \text{in all research fields by} \\
 \text{worldwide authors}}$$




Figure 8. Comparative levels of specialisation in Australian science and research during 2012–2021.

Data source: The Lens [32] and Scopus, Elsevier All Science Journal Classification [47]. A higher score indicates greater specialisation within the given field of research compared to the average comparison of Australia and the world.

3.4 Productivity decline

Productivity is the efficiency via which inputs are converted to outputs. Productivity is an important determinant of short-run economic growth and the primary determinant of long-run economic growth (and wealth generation). Ideas are the fuel source for productivity. When people discover how to produce something or deliver a service more efficiently it leads to productivity growth. Science is one of the most important pathways to discovery and ideation. The positive associations between science, ideas and productivity are well established

and accepted, although the magnitude of benefit is not easily quantified [55]. However, over the past decade or so, the world's advanced economies, and the science sector itself, have been caught within a productivity slump (Figure 9). This is harming economic growth and limiting long-term improvement of living standards.





Figure 9. Average (mean) annual OECD-country multi-factor productivity growth.

Data source: OECD Statistics [2].

A comprehensive study published in the American Economic Review in 2020 finds that within the science sector ‘research effort is rising substantially while research productivity is declining sharply’ [56]. The analysis examines agricultural crop yields, semi-conductors, cancer treatments, heart disease treatments, intellectual property patents and overall economic productivity. The analysis uses outcome metrics related to benefits such as changes in crop yields (e.g. wheat, corn) per unit of area resulting from agricultural R&D. Similar outcome metrics were used for the other categories. The researchers show that the cost of developing new pharmaceutical products to treat illness doubles every 9 years. They find that while the research workforce has grown, productivity (output per researcher) has decreased [56]. We are getting fewer ground-breaking ideas for each dollar invested.

For example, the cost of developing a new antibiotic was estimated at US\$1.581 billion in 2017 [57]. This far exceeds the costs of antibiotic discovery compared to the ‘golden era’ of the 1970s and 1980s. For antibiotic discovery, and many other types of scientific discovery, the next wave of discovery appears harder to achieve [58].

The authors conclude that the United States needs to double research effort, and double its research workforce, every 13 years to maintain science output. Without this investment, they argue, the United States economy will experience productivity decline and declining rates of GDP growth [56]. The United States economy, and other advanced economies worldwide, increasingly depend on science and technology improvements to sustain growth. Therefore, it is critical that the science sector – the engine room for the creation of ideas – keeps operating at full pace. If science productivity is declining, the only way to achieve this is via investing in more scientists and more science resources. The economists who did this study indicate that in advanced economies, income growth and improved living standards depend on research productivity and research effort [56]:



The centrality of innovation, science and research to productivity uplift and economic growth is well accepted and demonstrated within the field of economics. It means that if research productivity is declining, the only way of ensuring income growth is via increasing research effort. This is the approach followed by most advanced economies, as shown in the R&D expenditure data and statistics on the growing R&D workforce presented earlier in this chapter. Most OECD economies are growing the share of GDP spent on R&D to offset declines in research productivity and achieve overall economic growth.

The United States economic study finding productivity decline isn't a standalone. A study from the Research Institute of Economy, Trade and Industry in Japan applied the same techniques and found 'significant decline of R&D efficiency in the Japanese information service industry' [12]. The researchers recommend that the Japanese government implement R&D policies that address the decline. Another team of economists from the Leibniz Centre for European Economic Research and the Copenhagen Business School replicated the United States study for China and Germany using firm-level data over three decades [13]. They find evidence of productivity decline in both countries and 'strong decline' in Chinese R&D productivity. The authors conclude that 'diminishing returns in idea production are a global phenomenon, not just confined to the United States' [13]. These studies were presented at a recent OECD workshop examining science productivity and AI [11].

There will be many policy interventions needed to solve the productivity slump in science. However, the recent surge in AI capability and adoption is likely to play an important role. Recent years have seen AI substantially improve the speed, quality, safety and cost-effectiveness of scientific research. AI is already enabling discoveries which were hitherto beyond reach. Although AI has been used by scientists since the 1960s, it hasn't been mainstreamed until the last several years. The last few years have seen a huge increase in AI development and application in all scientific fields. AI is likely to be one of the most important mechanisms for boosting science productivity and escaping the slump. The need for science to reinvent itself, and problem-solve for industry and society, is likely to be a driving consideration for AI development and adoption into the future.

4 Artificial intelligence and knowledge discovery

4.1 Enablement or transformation?

To demonstrate the potential of AI to impact knowledge discovery, let us consider the case of electricity generation via nuclear fusion. According to the International Thermonuclear Experimental Reactor organisation in France, nuclear fusion creates 4 million times more energy than chemical reactions such as burning coal, oil or gas [59]. Electricity generation via nuclear fusion represents one of the most important (future) scientific discoveries for humanity. If discovered, nuclear fusion would provide an abundant and practically inexhaustible source of clean energy. Nuclear fusion does not produce the high-activity and long-lived nuclear waste associated with nuclear fission. Nuclear fusion provides greatly enhanced safety, substantially reduced financial costs and reduced risk of weaponisation [59]. Nuclear fusion can help solve climate change while supplying abundant energy for ample food, water and mineral (via mining and recycling) production. It is a game changer for humanity.

The main problem with electricity generation via nuclear fusion is that it currently cannot be accomplished in a practical and industrial way. The scientific community has been trying for decades. However, it appears this capability is getting closer. In early 2022 British scientists reported the production of 59 megajoules of energy sustained for 5 seconds from a nuclear fusion reaction; while the duration is minuscule by industrial electricity generation standards, this is nevertheless a major improvement upon previous records [60]. Another significant breakthrough came from the field of AI at about the same time. In February 2022 the results of a collaboration between DeepMind and the Swiss Plasma Center were published in Nature [61]. In this project reinforcement learning, a type of machine learning, was used to control the super-heated suspended plasma needed for the nuclear fusion reaction within a device called a tokamak. Reinforcement learning was used to control voltage in the tokamak and, thereby, the shape of the suspended plasma, ensuring it met experimental requirements while not touching the walls of the tokamak.

This is a well-understood but extremely difficult-to-solve optimisation problem for nuclear physicists. Reinforcement learning was able to identify plasma configurations not previously known. As reported in Wired magazine [62], Ambrogio Fasoli (fusion and plasma physicist and director of the Swiss Plasma Center) says this represents a 'significant step' on the pathway to nuclear fusion and that AI enables 'us to explore things that we wouldn't explore otherwise, because we can take risks with this kind of control system we wouldn't dare take otherwise' and that 'if we are sure that we have a control system that can take us close to the limit but not beyond the limit, we can actually explore possibilities that wouldn't otherwise be there for exploring' [62].

For scientists working on nuclear fusion, AI has provided a big boost. It has removed one of the critical barriers on the pathway to discovery; the ability to control the plasma within the tokamak. The future impact of AI on science and knowledge discovery can be viewed as a continuum of possibility. At one end of the continuum is enablement: the useful application of AI tools to help scientists do what they are already doing faster, cheaper, safer and better. At the other end of continuum is transformation: the use of AI to remove major barriers to scientific progress leading to paradigmatic shifts, new approaches to knowledge discovery and new possibilities for problem solving.

The case for enablement is well demonstrated through thousands of published AI studies within practically all science fields over recent decades. Our observations in this report, that the share of global scholarly publishing on AI has risen since 2020 from 1.2% to 5.7% and AI is now applied in virtually all disciplines, provide evidence of the usefulness of AI to scientists and researchers. The extent to which AI will (in the future) be transformative and associated with paradigmatic shifts in approaches to knowledge discovery and major leaps in problem-solving capability is less clear. However, some AI scientists see this as a distinct possibility.

Hiroaki Kitano is a Japanese AI scientist and is the director of both the Systems Biology Institute and Sony Computer Science Laboratories in Tokyo. Writing in Nature [63], Kitano proposes the 'Nobel Turing Challenge', which 'aims to develop a highly autonomous AI system that can perform top-level science, indistinguishable from the quality of that performed by the best human scientists, where some of the discoveries may be worthy of Nobel Prize level recognition and beyond'. In this paper Kitano also argues that future AI science 'may be an alternative form of science that will break the limitation of current scientific practice largely hampered by human cognitive limitation and sociological constraints' and that such approaches 'could give rise to a human-AI hybrid form of science that shall bring systems biology and other sciences into the next stage' [63]. Kitano is connected with a team at the Alan Turing Institute in the United Kingdom which is delivering a project ('The Turing AI scientist grand challenge') which tackles similar and related objectives [64]. As stated on the Alan Turing website, this ongoing project started in January 2021 and is, among other things [64, 65]:

- reviewing current autonomous systems capable of performing scientific research with a focus on AI approaches capable of pushing disciplinary science beyond the current cutting-edge
- developing a multi-year roadmap charting a scientific and technical pathway for AI for science with milestones identified in materials, biomedicine and environmental sciences.

A related line of inquiry pursued by researchers working on the philosophy of science is about the possibility of theory-free, data-intensive science [66-68]. Starting in the early 2000s 'this approach is supposed to be data driven, strongly inductive, and relatively theory independent' [69]. Data-intensive science can be considered transformative as it represents a paradigm shift challenging existing approaches to knowledge discovery. The idea comes from the successful application of data science for forecasting in fields like meteorology [70], economics, energy, and demographics [71].

The idea has credence because sometimes data can work better than theory for modelling and predicting system behaviour [71]. However, science philosophers have criticised the idea of theory-free data-intensive science [72]. Some have argued that researchers using 'big data' approaches (akin to data-intensive science) may sidestep the critical hurdle of causality and rely upon statistical correlation to explain and predict system behaviour [73]. The concern about sole reliance on

data-intensive approaches is they (a) fail to draw upon existing theory, and (b) fail to establish and understand causality [73]. This can lead to errors and accidents.

Recent perspectives suggest that 'data versus theory' is a false dichotomy and that there is no competition between the two [69]. Instead, data science approaches are inextricably linked to theory and have unique application in practically every field of science [74]. Techniques such as linear regression analysis have long been used by scientists to understand and explain real-world phenomena. When linked with other tools and ideas, these techniques can lead to confirmation of existing theories or the development of new theories. Scientists have always used data, and as the tools of data science get better, they are getting better at using data. Theory and causality haven't vanished; they remain critically important. The observations about data-intensive science are likely to apply to AI-intensive science. Data science often uses AI, and AI almost always uses data science; the two fields are increasingly inseparable.

4.2 Case studies – Artificial intelligence applications for science

In this section we describe case studies where AI has improved the efficiency and effectiveness (the productivity) of scientific research. A more detailed, comprehensive and up-to-date repository of case studies is available on the CSIRO website (www.csiro.au/en/research/technology-space/ai).

The case studies also illustrate how AI is enabling science and research in diverse fields of study. This includes examples of enablement and transformation, where AI has enabled scientists to solve complex problems and has created an elevated platform of capability and knowledge discovery. Overall, AI and its constituent technologies (such as pattern recognition and machine learning) appear to considerably enhance and accelerate the scientific process by allowing: (a) faster processing of data, (b) handling of very large and datasets, (c) handling of disparate datasets, (d) offloading of menial tasks, (e) deeper and wider exploration of the experimental space, (f) more accurate predictions due to better models, and (g) faster and more reliable detection of salient and/or anomalous patterns or events. However, the enhancement and acceleration of science via AI is largely predicated on relevant data being available in digital format, thereby necessitating that any physical experiments need to be designed and executed such that the primary objectives include data acquisition in digital format.

4.2.1 Predicting the 3D structure of proteins

Proteins are essential for the growth and maintenance of all cells and tissues [75, 76]. Understanding their structure, or the way they fold, is key to identifying their function – a time-consuming and challenging problem that scientists have spent decades trying to solve [75, 76]. AlphaFold is a neural network system developed by Google's DeepMind that can map the 3D structure of proteins with significantly greater accuracy than conventional methods [75, 77]. AlphaFold has been applied to map the human proteome (the entire set of proteins that make up the human body) and was able to predict the structure of 98.5% of human proteins [77]. Together with the European Molecular Biology Laboratory's European Bioinformatics Institute, DeepMind has developed the AlphaFold Protein Structure Database to make these predictions available to the scientific community [78]. This technology can potentially support future advances in biological research and drug development.

4.2.2 Accelerating solar panel research

Researchers at CSIRO have developed a research robot that can autonomously test flexible solar panel samples [5]. These researchers developed the autonomous system during the 'second wave' of COVID-19 in Melbourne in 2020. Before this, researchers could manually test up to 20 solar cells per day and had to be physically present in the lab [5]. The new automated research system is controlled remotely and could test 12,000 cells in 24 hours, which represents a 600-times improvement in productivity [5]. AI and machine learning is also being applied to efficiently analyse and predict parameters for solar cell manufacturing of organic solar cells [77]. These applications illustrate how autonomous testing, combined with machine learning, can increase the efficiency of scientific research and accelerate the development of new technologies, even while scientists are working from home [5].

4.2.3 Enhancing the reach of citizen science

AI is being applied in citizen science, assisting civic educators and scientists in engaging the community in scientific endeavours and collecting large datasets on rare or difficult to access phenomena. Examples include iNaturalist, a platform run by the California Academy of Sciences and National Geographic where members of the public can submit photos of the natural world, including animals and plants [79]. This platform uses computer vision and a machine-learning model previously trained on an existing research-grade dataset of images [80, 81]. Citizen science can enhance the spatial and temporal resolution of data in ecological monitoring projects relative to traditional methods [82]. Using AI systems can improve the cost efficiency of collecting, processing and analysing data generated by the public [82] enabling more researchers to leverage the benefits of citizen science.

4.2.4 Predicting the replicability of scientific studies

The replicability of scientific findings was brought into question with a series of publications demonstrating that a large share of studies in psychology, economics, and medicine could not be replicated [83-85]. Non-replicability can impede scientific progress, hinder public support and trust in science, and waste finite funding resources [86]. Researchers from Northwestern University used machine learning to accurately estimate the replicability of a study; meaning the extent to which it is possible to replicate the methodology but not necessarily whether (or not) the results hold-up. The machine-learning approach performed as well as expert survey predictions, which is the current gold-standard, but resource-intensive, method of assessing replicability [86]. While this research is preliminary, it suggests that AI can potentially be used to test the replicability of scientific findings without imposing additional time and resource requirements on scientists.

4.2.5 Discovering and developing new materials

AI is being widely applied across materials science to accelerate the rate at which scientists can discover and develop new materials [87]. An early example is the Autonomous Research System developed by researchers from the Air Force Research Laboratory, UES Inc. and Lockheed Martin Advanced Technology Laboratories [88]. This system combines robotics, AI, data science, and in situ technologies to design, execute, and analyse experiments faster than traditional human-driven approaches [88]. This approach has been used to explore the synthesis of carbon nanotubes – a well-suited material for electronics applications that scientists have spent decades trying to understand [88]. Through autonomous experimentation, AI can increase the speed and cost of materials science research, increase the productivity of scientists, and maximise the value that can be derived from complex multi-dimensional datasets [87].

4.2.6 Untangling mathematical relationships

Advances in mathematics depend on the ability of mathematicians to discover new patterns and formulate statements around the potential relationship between objects (referred to as a conjecture) [89]. These insights are then used to develop new mathematical proofs [89]. AI can assist mathematicians in the initial step of detecting patterns between objects, which can help guide them in developing mathematical formulae and theorems [89].

A team of researchers from DeepMind, the United Kingdom and Australia have applied AI to study the algebraic and geometric structure of knots – a longstanding mathematical challenge [90]. The machine-learning approach enabled the researchers to discover novel and surprising patterns, and develop new conjectures [89].

4.2.7 Improving the efficiency in conservation science

Like many aspects of science, conservation research operates in a resource-constrained environment. Conservation managers need to determine the most effective way of managing these finite resources and identify when to stop efforts to manage and survey an endangered species population. Still, they often have insufficient information for making these decisions [91].

To provide better intelligence, deep-learning techniques have been used to count the number of endangered species animals from aerial survey images [92, 93]. Using images collected from motion-sensor cameras placed in natural habitats, researchers found that a deep-learning system can identify animals as accurately as a human observer [93]. The use of AI here to classify around 5.5 million images saved over 8.4 years in human labour [93]. This demonstrates the significant cost and time savings that AI can provide in conducting conservation research.

4.2.8 Predicting high-impact research

The impact of scientific outputs is typically measured through citation metrics, such as h-indices and journal impact factors. These metrics can be discipline-specific, biased, or reflect lag quality indicators [94, 95]. With trillions of dollars invested in research globally each year, having reliable predictors of impactful research is critical. Researchers from the Massachusetts Institute of Technology have developed a DELPHI framework (Dynamic Early-warning by Learning to Predict High Impact), which uses machine learning to predict the likely impact of scientific publications [96]. It draws upon a rich collection of publication, journal and citation data [96]. The model was able to predict high-impact research the year it was published with 77% accuracy, and it was a better predictor than citation metrics [96]. This work demonstrates how AI can potentially be used to inform funding decisions to maximise the return on investment and impact of scientific research.

4.2.9 Decoding the human brain

The human brain is arguably the most complex system known to humankind, but AI is helping scientists to crack the neural code. AI provides scientists with opportunities to directly explore the functioning of healthy, neurotypical human brains, which has previously been limited, if not impossible, due to practical or ethical considerations [97]. For example, researchers from the Massachusetts Institute of Technology have used deep neural networks to demonstrate the hierarchical structure of the human visual system [98]. The neural-network model can identify objects as well as a human and exhibits a similar pattern of neural activity to a monkey brain performing a similar task [98]. Similar applications have been used to show the hierarchical organisation of the human auditory cortex [99]. These neural-process models can help scientists generate hypotheses around certain brain functions and inform their experimental design [97].

4.2.10 Identifying drugs for antibiotic-resistant bacteria

AI has been applied in a range of medical contexts, and increasingly, it is being used to accelerate the discovery of new drugs [100]. Improving the efficiency of drug discovery research is particularly critical in the face of the increasing prevalence of antibiotic-resistant bacteria and the diminishing returns on investment and high risk associated with drug discovery [101-103]. Researchers have used deep learning to develop a model that can identify candidate molecules capable of inhibiting the growth of *Escherichia coli* (*E. coli*) – a common bacteria associated with antimicrobial resistance [101]. Similarly, scientists have also used AI to rapidly identify and test novel compounds that can inhibit discoidin domain receptor 1 (a common receptor implicated in fibrosis and other diseases) [104]. This emerging work demonstrates how deep-learning models can lead to faster and more cost-effective experiments than traditional approaches [100].

4.2.11 Representing spatial phenomena

Geospatial AI is an emerging field that combines spatial science with AI methods to derive rich insights from big spatial data [105]. These approaches can be used to understand the environmental factors that people may be exposed to at a given geographical location and time and how this exposure may impact their health [105].

A group of researchers from the University of Southern California used this approach to develop a model that can predict air quality (i.e. particulate matter air pollution <2.5 µm in diameter, or PM2.5) [106]. The model accurately predicted PM2.5 concentration levels without relying on prior domain knowledge and quantified the impact of various geographic features (e.g. parking lots, commercial buildings) on air quality [106]. The fine spatiotemporal resolution of this model provided insights into the impacts of air pollution (e.g. health or environmental outcomes) on specific populations [106].

4.2.12 Automating literature reviews and bibliometric analyses

Scientists are under increasing strain to keep up with the ever-growing number of scientific publications, as are editors and peer-reviewers [107]. Moreover, scholarly works are usually written for expert audiences in specific academic fields, limiting their use beyond academia, or even other academic fields. Emerging AI tools can assist with reviewing, appraising and summarising scientific publications. Examples include the Artificial Intelligence

Review Assistant, launched by open-access publisher Frontiers in June 2020, which can screen the language quality, integrity of figures, instances of plagiarism and potential conflicts of interest in submissions [108]. Other AI-enabled systems use natural language processing tools to synthesise academic papers into language that a 7-year-old child can understand [109]. These tools can potentially improve the productivity of academics and publishers and assist them in identifying and evaluating relevant and impactful research [107, 108].

4.2.13 Identifying archaeological samples

In archaeology, deep-learning approaches have been increasingly applied to make sense of often unstructured and disparate datasets, and to derive new insights from archaeological records [110]. Traditional sampling methods can be time and resource-intensive [111]. Existing applications have used pattern recognition and AI to identify patterns in pottery and engraved wooden artifacts [111, 112] as well as to sort and filter images and identify objects in images (e.g. rock art, tools, shell or animal bone) [110]. With advances in large-scale lidar, satellite and aerial imagery, archaeologists have access to richer geospatial data for archaeological mapping of sites [110]. As a result, using machine learning to analyse geospatial data can help archaeologists profile the landscape characteristics without physically accessing archaeological sites [110, 113, 114].

4.2.14 Faster chip design

When designing computer chips (such as CPUs and GPUs), various components must be placed in a floorplan that satisfies many operational requirements, including metrics such as power consumption, performance, and chip area. This typically requires considerable manual effort over many months to generate manufacturable layouts. To address this problem, researchers at Google designed an AI system for automatic floorplan generation based on reinforcement learning (a type of machine learning which learns from its past mistakes) that required a training dataset of 10,000 chip layouts of varying quality. The trained system was then able to automatically generate floorplans in only 6 hours, with the resulting floorplans having comparable or better metrics than human-designed floorplans [115]. Similar reinforcement-learning approaches may be applicable in science domains where experimental designs require time-consuming trials to explore the space of many possible outcomes. The potential time savings may allow researchers to focus on higher level tasks and hence be more productive; substantial time savings may also lead to discoveries that hitherto were too time-consuming to pursue.

5 Science domain adoption trends

In this section we explore temporal patterns of AI application and development during 63 years from 1960 to 2022 in science application domains. This shows how the field of AI has moved beyond computer science into other scientific and academic research disciplines. Overall, AI continues to be increasingly adopted across all areas of science, as evidenced by the increasing share of AI publications relative to total publications. This trend is likely to continue for some time until AI usage is normalised within science domains and as researchers grapple with the ever-increasing volume of scientific data.

5.1 Data sources and methods (bibliometric analysis)

We applied bibliometric analysis to explore trends in AI adoption [116]. Bibliometric analysis involves the examination of terms and phrases in research literature to understand important trends or patterns. Bibliometric analysis is becoming increasingly popular as the body of published research continues to expand [117, 118].

The volume and rate of publishing in some fields makes comprehensive literature reviews by human researchers difficult or infeasible. Semi-automated literature searches which augment human researchers are increasingly needed to achieve up-to-date coverage of all relevant publications.

Previous studies have used bibliometric analysis to explore AI research patterns. One such study used Microsoft Academic Graph [119] to examine the extent to which social science research fields are cited within AI publications [120]. The authors found social science fields such as geography, art and philosophy were under-represented in AI research. They conclude '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' [120]. In another bibliometric study, web-of-science data was used to examine AI research efforts across countries, sponsors, institutions and disciplines [121]. This study found that AI technology development has arisen from high levels of interdisciplinary research. The Stanford University AI index report also uses bibliometric analysis to measure AI publishing intensity [14].

Our work contributes to this body of knowledge by using a novel, large, comprehensive and up-to-date dataset of scholarly publishing. We use a broader definition of AI with a larger and more diverse set of search phrases developed by the OECD [122] via expert consultation. Our bibliometric analysis is focused on AI application within other research fields; not the mirror (opposite) issue covered in earlier work [120] about how other fields have been used within AI research. Our analysis uses a formal, comprehensive and granular classification of research covering all major fields of physical, natural and social sciences and arts and humanities. We take a historical perspective examining AI publication trends from 1960 to 2022. We also examine patent citations relating to AI technologies and how different types of AI technology have evolved over time.

Our bibliometric analysis of AI research trends across various scientific application domains is based on The Lens database [123]. The Lens is a product of a collaboration between the Queensland University of Technology and a not-for-profit Brisbane-based firm Cambia. It received funding from the Bill and Melinda Gates Foundation, the Rockefeller Foundation and other organisations. As of 20 September 2022, The Lens database contained 249 million scholarly publications and over 143 million patent records.

The Lens has previously been used to analyse science trends relating to genetics [33] and COVID-19 [124]. The Lens draws upon data from the following databases:

- CrossRef
- The Open Researcher and Contributor Identifier (ORCID)
- PubMed
- Impactstory
- COncnecting REpositories (CORE)
- Microsoft Academic (which ceased operations on 31 December 2021)
- European Patent Office (EPO)
- United States Patent and Trademark Office (USPTO)
- Intellectual Property (IP) Australia
- World Intellectual Property Organization (WIPO).

We extracted bibliometric data from The Lens using their Application Programming Interface (API) with Python scripts during 18–29 April 2022 (to capture the 1960–2021 time period) and on 20 September 2022 (for the last, partial, year of data of 2022). The last year, therefore, contains data for 72% of the year. Whilst the absolute numbers of AI publications are likely to rise and change over the remaining days of the year, the relative shares (percentages) are likely to remain stable. The Lens database contains records on all scholarly publications in all fields of research by the whole world over all history. As such it is a large, complex and continually evolving database. The content,

structures and definitions/rules of the database are likely to be changing. This means that future extractions of the data could yield different results but are unlikely to change the main implications/results from our study.

To identify AI-related publications, we used a set of 214 AI search phrases developed by the OECD via expert consultation (Appendix B). AI-related publications had to contain one or more of the 214 search phrases in the publication title, abstract or keywords. We limited our analysis to scholarly publications that were journal papers, books, book chapters, conference proceedings and conference proceedings articles; all of which are peer-reviewed. This search strategy returned 3.35 million AI-related scholarly works published between 1 January 1960 and 20 September 2022 (Figure 10).

Each publication was classified by field of science using the All Science Journal Classification (ASJC) system. The ASJC is a three-level hierarchical taxonomy that represents a comprehensive classification of global research covering all fields of study and is maintained by Elsevier (Table 1). The top level has four categories (health sciences, life sciences, physical sciences, and social sciences and humanities), the second level has 26 categories, and the third level has 333 categories. The ASJC codes are assigned by Elsevier’s team of in-house experts at the time of publication. The assignment of the code is based on the aims, title and content of the publication [47]. A single publication can be assigned multiple ASJC codes.

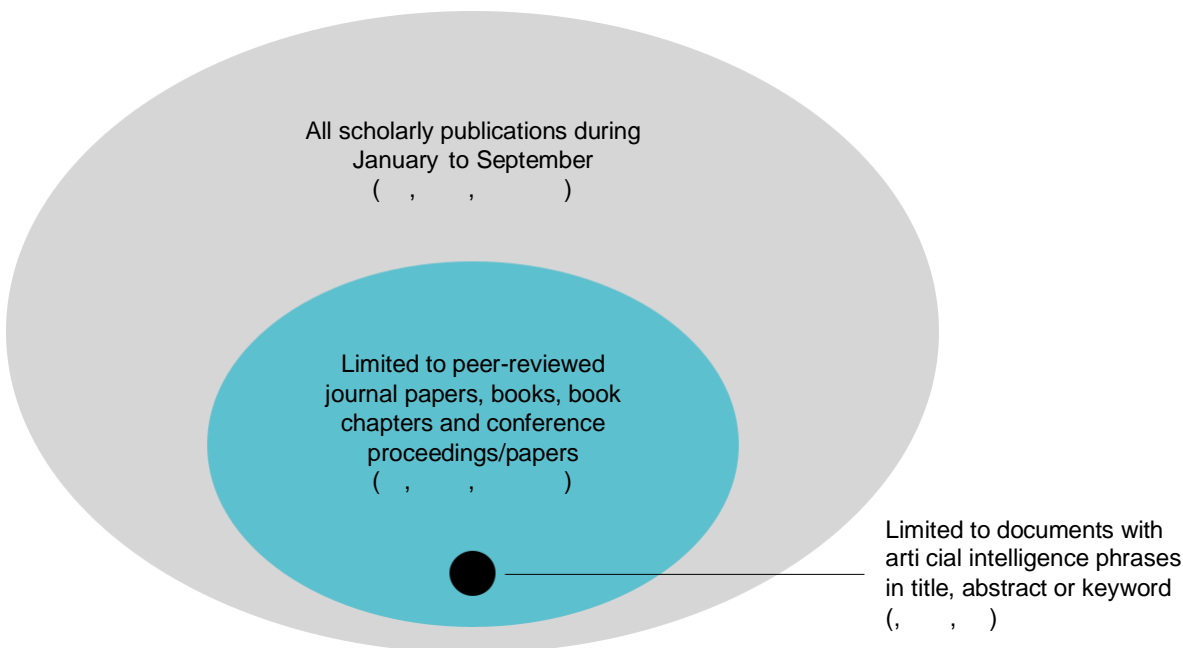


Figure 10. Identifying artificial intelligence scholarly works from 1960 to 2022.

Data source: The Lens [32].

Table 1. All Science Journal Classification (ASJC) categories [47].

First-level research field	Second-level research field	Number of third-level research fields
Health sciences	Dentistry	7
	Health professions	17
	Medicine	49
	Nursing	24
	Veterinary	5
Life sciences	Agricultural and biological sciences	12
	Biochemistry genetics and molecular biology	16
	Immunology and microbiology	7
	Neuroscience	10
	Pharmacology toxicology and pharmaceuticals	6
Physical sciences	Chemical engineering	9
	Chemistry	8
	Computer science	13
	Earth and planetary sciences	14
	Energy	6
	Engineering	17
	Environmental science	13
	Materials science	9
	Mathematics	15
	Physics and astronomy	11
Social sciences and humanities	Arts and humanities	14
	Business management and accounting	11
	Decision sciences	5
	Economics econometrics and finance	4
	Psychology	8
	Social sciences	23

5.2 Publishing intensity and volumes – all research fields

In this section we examine temporal patterns in AI publishing across all ASJC fields of research. We found the volume of AI publishing continually increased over history, both in terms of the total number of AI publications and the relative share of total publications (Figure 11). The only exceptions were 1964 and 1971, where the number of publications contracted, before returning to growth in the following year. Interestingly, we did not find convincing evidence of a decline in AI publications associated with 1974–1980 and 1987–1993, the periods corresponding to the first and second AI winters, respectively (Figure 12). However, our analysis is not exhaustive and there may be time lags, spatial patterns and impacts on specialised AI fields worthy of further consideration.

Growth in AI publishing has been greatest in the past 5–6 years, with the relative share of AI publishing rising from 2.9% of all publications in 2016 to 5.7% of all publications in 2022. The total number of AI publications rose from 159,426 to 344,265 over this time period which equates to a 2.2-times increase. In most fields of research, the amount of AI adoption in the past several years roughly equals what happened over all preceding history. Across all fields of research, the volume of peer-reviewed publications on AI in the past 7 years (1.6 million documents) exceeds all prior AI publishing over the preceding 55 years (1.5 million documents).

Assuming these trends continue, it is likely that a much greater share of publishing will be on the topic of AI by 2030. Our analysis of AI-related publications suggests that we are currently on the steepest part of the adoption curve and there are no signs of a slowdown. As such, the full potential impact of AI on science and research domains lies ahead. AI will become increasingly integrated into routine research practices. As AI becomes normalised researchers may apply AI tools and concepts without using AI phrases in the title, abstract or keywords.

5.3 Adoption trends in application domains

In this section we explore AI publishing intensity within the four first-level (Figure 13) and 26 second-level (Table 2) research fields. We report AI publishing intensity as the percentage of AI-related publications out of the total number of publications. We also report the AI publication counts in absolute terms. Physical sciences account for the bulk of AI publishing in relative and absolute terms. In 2021, there were 461,000 AI-related publications in the physical sciences, accounting for 9.4% of total publication output. AI-related publishing was roughly evenly distributed across other first-level fields of science, with the social sciences and humanities, life science and health sciences making up 3.9%, 3.4% and 2.6%, respectively.

We examined trends in AI publishing intensity over time to find evidence of AI winters across specific fields of research. There is evidence of an AI winter from 1974 to 1980 in the field of computer science (the first winter), with publishing intensity dropping and then plateauing during this period. Across all other fields of research, there is no clear evidence of a similar pattern during either the first or second AI winter. This suggests the first AI winter may have been isolated to the computer science domain. Beyond this we did not find evidence of domain-specific slow-downs in AI publishing associated with either of the two AI winters.

Looking across the second-level research fields, we found that computer science dominates AI publishing, with one-quarter of all publications in the field on AI. Mathematics (14%), engineering (11%) and decision sciences (11%) also have a high AI publishing intensity. A similar pattern was also observed when looking at the number of AI publications across fields of research. The lowest level of AI penetration was observed within dentistry, nursing, veterinary science, pharmacology toxicology and pharmaceuticals, where publishing intensity ranged from 1% to 2%. However, this appears to be changing, as AI publishing has increased in these fields (and most other fields of research) in recent years.

SHARE OF PUBLISHING ON ARTIFICIAL INTELLIGENCE %




Figure 11. Peer-reviewed research publications on artificial intelligence.

Data source: The Lens [32]. Date range is from 1 January 1960 to 20 September 2022.

YEAR ON YEAR CHANGE IN AI PUBLICATIONS %




Figure 12. Annual change in peer-reviewed AI-related publications (%) over time.

Data source: The Lens [32].

Table 2. Artificial intelligence publishing intensity by research field (percentage).

Fields of Research (Second-Level ASJC)	1970	1980	1990	2000	2010	2015	2020	2021	2022*
Agricultural and biological sciences	0.1	0.1	0.3	0.7	1.2	1.5	2.5	2.8	3.5
Arts and humanities	0.0	0.1	0.3	0.4	0.6	0.7	2.3	3.2	2.7
Biochemistry genetics and molecular biology	0.1	0.1	0.2	0.4	1.3	1.9	3.2	3.8	4.8
Business management and accounting	0.5	0.5	0.9	1.3	2.2	2.6	4.8	5.0	6.3
Chemical engineering	0.1	0.0	0.2	0.7	1.0	1.1	4.1	4.8	5.4
Chemistry	0.0	0.1	0.2	0.4	0.6	1.0	2.7	3.2	3.6
Computer science	3.7	1.9	6.9	12.4	16.0	17.1	22.7	25.7	29.5
Decision sciences	2.3	1.4	2.1	4.5	7.1	8.5	9.8	11.3	14.9
Dentistry	0.0	0.0	0.1	0.3	0.3	0.3	0.9	1.7	1.6
Earth and planetary sciences	0.1	0.2	0.5	0.9	1.7	2.5	4.4	5.5	7.3
Economics econometrics and finance	0.0	0.1	0.3	0.8	0.9	1.1	2.7	3.5	3.9
Energy	0.1	0.2	0.3	0.8	1.5	2.1	4.5	5.2	6.0
Engineering	0.3	0.4	1.6	3.0	4.4	5.2	10.1	11.3	12.4
Environmental science	0.1	0.2	0.3	0.7	1.3	1.7	2.9	3.3	4.0
Health professions	0.1	0.2	0.4	1.1	1.4	2.2	3.2	4.1	7.3
Immunology and microbiology	0.1	0.1	0.1	0.3	0.7	1.4	1.9	2.3	3.1
Materials science	0.1	0.1	0.3	0.5	0.7	0.9	4.2	4.1	4.2
Mathematics	0.6	0.8	1.9	4.9	7.9	9.0	12.7	14.1	15.3
Medicine	0.0	0.1	0.2	0.3	0.8	1.1	2.2	2.7	3.5
Neuroscience	0.1	0.1	0.5	1.1	2.3	3.5	5.1	6.1	8.3
Nursing	0.0	0.0	0.1	0.2	0.3	0.5	1.1	1.2	1.9
Pharmacology toxicology and pharmaceuticals	0.0	0.0	0.1	0.3	0.6	0.9	1.7	2.0	2.2
Physics and astronomy	0.1	0.2	0.6	0.8	1.2	1.7	5.6	7.0	7.2
Psychology	0.2	0.4	0.7	1.2	1.7	2.2	2.7	2.9	3.6
Social sciences	0.1	0.1	0.3	0.4	0.8	1.2	2.8	3.6	4.1
Veterinary	0.0	0.0	0.1	0.1	0.2	0.4	1.0	1.1	1.4

Level of AI publishing intensity	Low	Medium	High
----------------------------------	-----	--------	------

* Final year of data captures publishing activity within 1 January to 20 September 2022 only.

SHARE OF TOTAL PUBLISHING ON ARTIFICIAL INTELLIGENCE

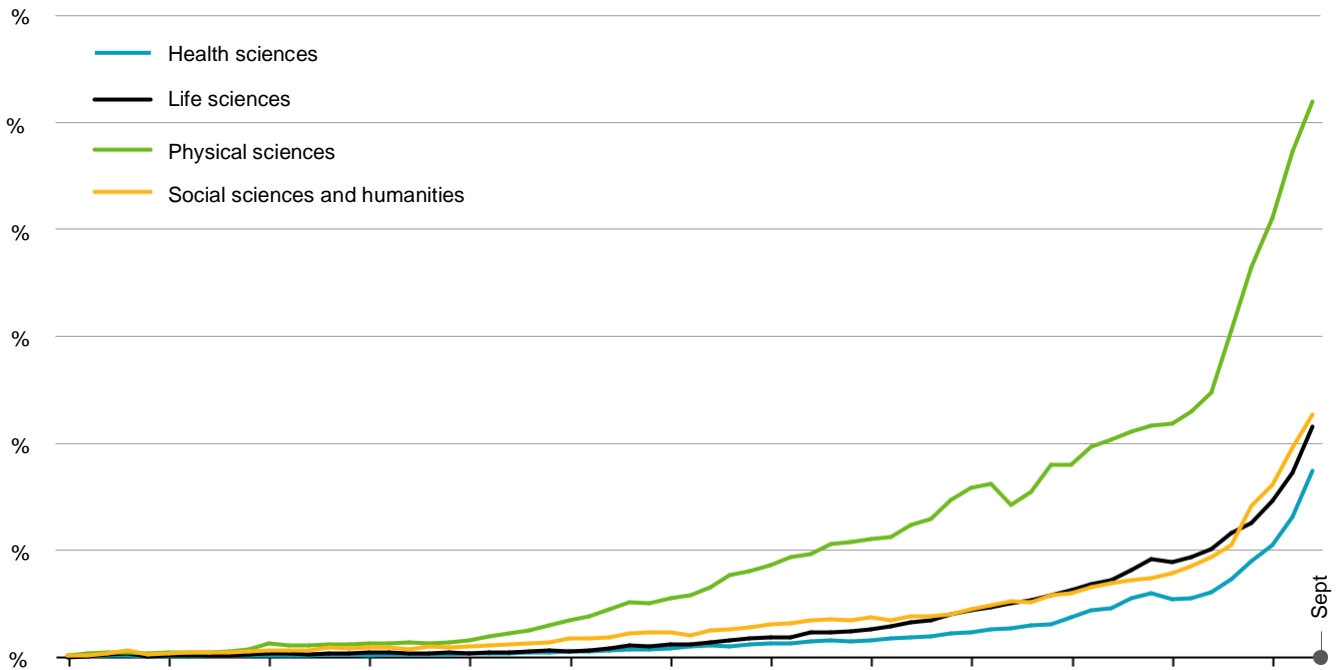


Figure 13. Artificial intelligence publishing intensity by main research domains.

Data source: The Lens [32]. Data sourced from January 1960 to September 2022.

The AI publishing intensity within a field of science is influenced by the overall publishing volume; both metrics are important. For example, while medicine has comparatively lower AI publishing intensity relative to other fields (2.6% of publications were AI-related in 2021), it accounts for a substantial share of AI publication volume. In 2021 there were 55,374 AI-related publications within the field of medicine, which accounts for 8.4% of all AI-related publications. This makes medicine the third-largest field by publication volume after computer science (25.5% of all AI-related publications) and engineering (18.1%). As such, medicine represents an important area for AI science and technology development. Within the overall field of medicine, the usage and development of AI technologies is most pronounced within the third-level fields of health informatics, as well as radiology, nuclear medicine and imaging, where access to digital data is readily available.

The development and adoption of AI technology started within the areas of computer science, mathematics, engineering and decision sciences, with the AI publishing intensity picking up in these areas from the early 1980s or before. Most other fields do not show substantial uptake until the 2000s or 2010s. What is consistent across all fields of research, however, is a sudden and substantial surge in AI publishing intensity from 2017 to 2021 (Figure 14 to Figure 18). Except for materials science, publishing intensity has continued to rise, or accelerate, in all fields of research during the COVID-19 pandemic. Overall, AI is having a greater impact in the current era, compared to all history, for the physical sciences, life sciences, health sciences, social sciences and humanities.

SHARE OF TOTAL PUBLISHING ON ARTIFICIAL INTELLIGENCE

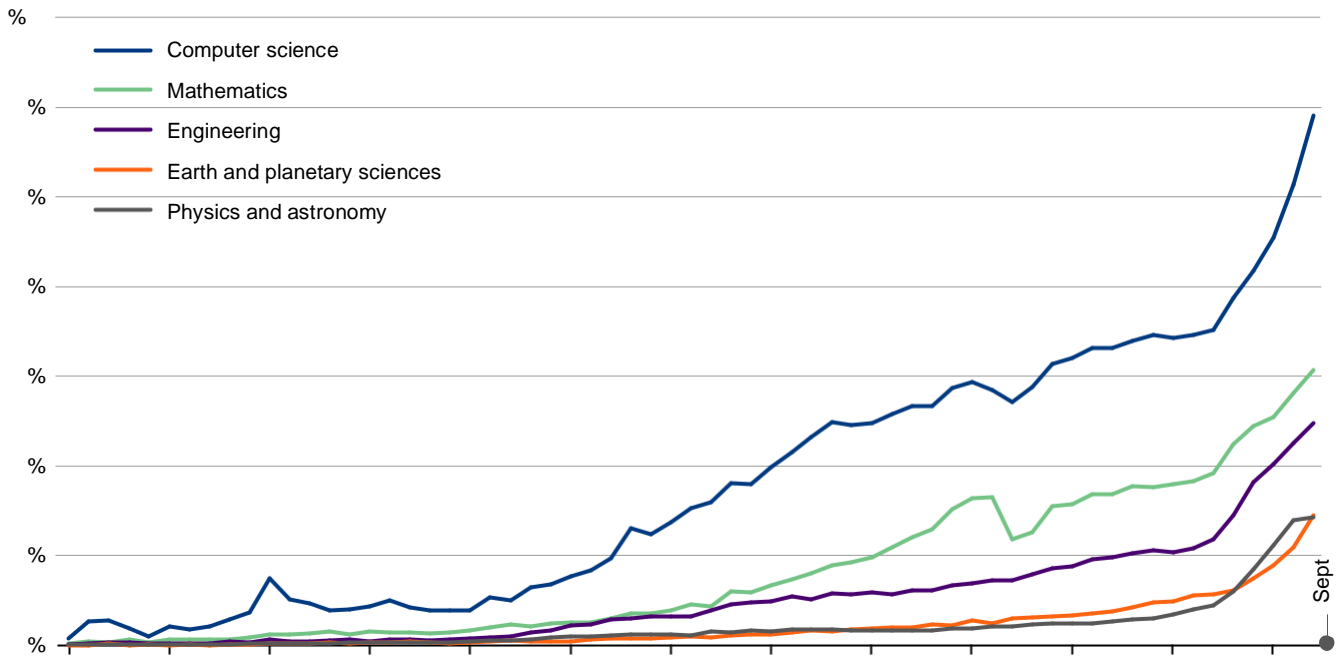


Figure 14. Artificial intelligence publishing intensity in the physical sciences (1–5 fields).

Data source: The Lens [32]. Data is sourced from January 1960 to September 2022.

SHARE OF TOTAL PUBLISHING ON ARTIFICIAL INTELLIGENCE




Figure 15. Artificial intelligence publishing intensity in the physical sciences (6–10 fields).

Data source: The Lens [32]. Data is sourced from January 1960 to September 2022.

SHARE OF TOTAL PUBLISHING ON ARTIFICIAL INTELLIGENCE




Figure 16. Artificial intelligence publishing intensity in the health sciences.

Data source: The Lens [32]. Data is sourced from January 1960 to September 2022.

SHARE OF TOTAL PUBLISHING ON ARTIFICIAL INTELLIGENCE

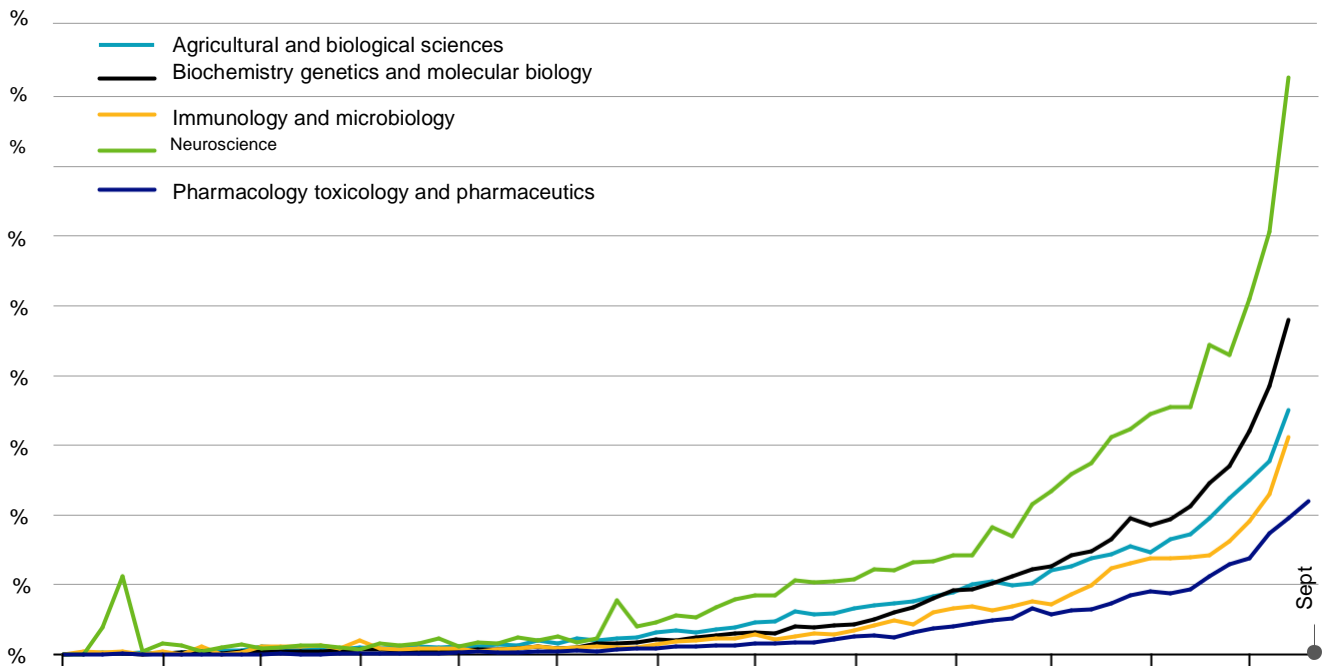


Figure 17. Artificial intelligence publishing intensity in the life sciences.

Data source: The Lens [32]. Data is sourced from January 1960 to September 2022.

SHARE OF TOTAL PUBLISHING ON ARTIFICIAL INTELLIGENCE

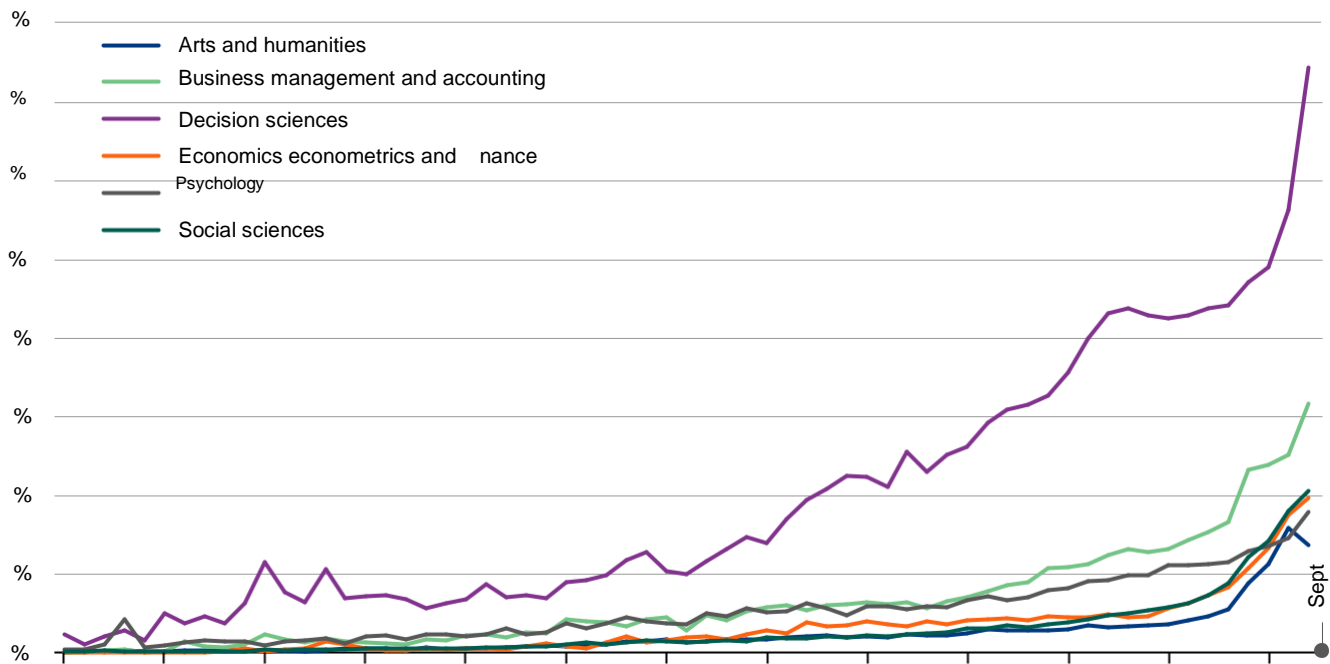


Figure 18. Artificial intelligence publishing intensity in the social sciences and humanities.

Data source: The Lens [32]. Data is sourced from January 1960 to September 2022.

5.4 Artificial intelligence technology diffusion trends

In the early stages of its development AI research mostly occurred within the computer sciences field and, to a lesser extent, engineering, mathematics and decision sciences. However, the usefulness of AI was soon discovered by scientists in many other fields of research. The pattern via which AI technologies developed by computer scientists were adopted in other fields can be considered technology diffusion, a sub-field of research in the discipline of technology economics [125]. In this section we explore the diffusion of AI technology. We do this by analysing time-series data on the concentration of AI activity and temporal patterns of AI adoption. We analyse the most granular level of the ASJC classification, the third-level, which contains 333 unique fields of research.

Our analysis of the concentration of AI publishing over time uses the Gini Coefficient (GC). The GC is a measure of concentration for any variable across multiple categories, such as how evenly wealth is distributed among individuals within society. We used the GC to measure the extent to

which AI is concentrated in a few fields of research versus evenly distributed across all fields, an approach taken in earlier AI technology diffusion analyses [120]. The GC ranges from 0 to 1, where 1 implies all AI publications are in one field of research and 0 implies a perfectly equal distribution of AI publishing across all fields.

We found that from 1960 to 1980 the GC fell from 0.9 to 0.7 as AI diffused beyond the foundation disciplines of computer science, mathematics and engineering into a much broader range of application domains (Figure 19). This level of AI diffusion has been sustained over the following four decades. The main reason the GC hasn't fallen further, indicating a more even distribution of AI activity, is that computer science has increased AI publishing faster than other research fields.

The number of fields of research using AI has increased over the same time period, with close to all fields publishing on AI by 2021; up from 70% in 1980 (Figure 20). It took roughly 25 years from 1960 to 1985 for AI technologies to be represented in over 80% of all research fields. We can see from these data that AI has been embedded and applied in most fields of research for decades.

LEVEL OF CONCENTRATION OF AI PUBLISHING GINI COEFFICIENT




Figure 19. Concentration of artificial intelligence publishing across research fields.

Note: The Gini Coefficient is a statistical measure of concentration. It is used here to measure the level of concentration of AI publishing across fields of research. A higher value represents increased concentration.

NUMBER OF RESEARCH FIELDS WITH AI PUBLISHING TOTAL = ... FIELDS•




Figure 20. Diffusion of artificial intelligence technology into research fields.

Note: There are a total of 333 third-level All Science Journal Classification research fields. The graph shows the number of research fields with artificial intelligence publishing from 1960 to 2021.

5.5 Trends in artificial intelligence technologies

In this section we explore publishing trends across various subfields of AI. We analysed publication counts for the 214 OECD AI phrases (Appendix B) from 1960 to 2021.

The majority of AI technologies saw publishing volumes increase over time, with some developing faster than others. The top six AI technologies with strongest growth over the past 20 years were convolution neural networks, followed by deep learning, random forest, generative adversarial networks (defined as ‘adversarial network’), sentiment analysis and transfer learning (Figure 21). Deep learning accounts for the greatest share of AI publications in 2021 (7.6% of total AI publications) and has shown the strongest 20-year increase, growing 25.4 times from 2002 to 2021. Looking at the past 5 years alone, the strongest growth has been observed across ‘generative adversarial networks’ and ‘transfer learning’ (growing by 2.3 and 5.3 times during 2016–2021, respectively).

We also examined the extent to which various AI technologies have been translated into commercially valuable and socially useful products. We did this by examining the number of patent citations attributed to research publications containing an AI-related search term in the title, abstract or keyword. Patent citations were taken as a proxy measure of commercial (and societal) impact, where a greater number of patent citations is assumed to indicate greater impact. The analysis revealed that scholarly publications which included reference to ‘neural network’, ‘convolutional neural network’, ‘deep learning’, ‘machine learning’ and ‘computer vision’ generated the greatest number of patents (Figure 22). These AI subfields are associated with the most significant commercial value and impact.

THOUSANDS OF PUBLICATIONS

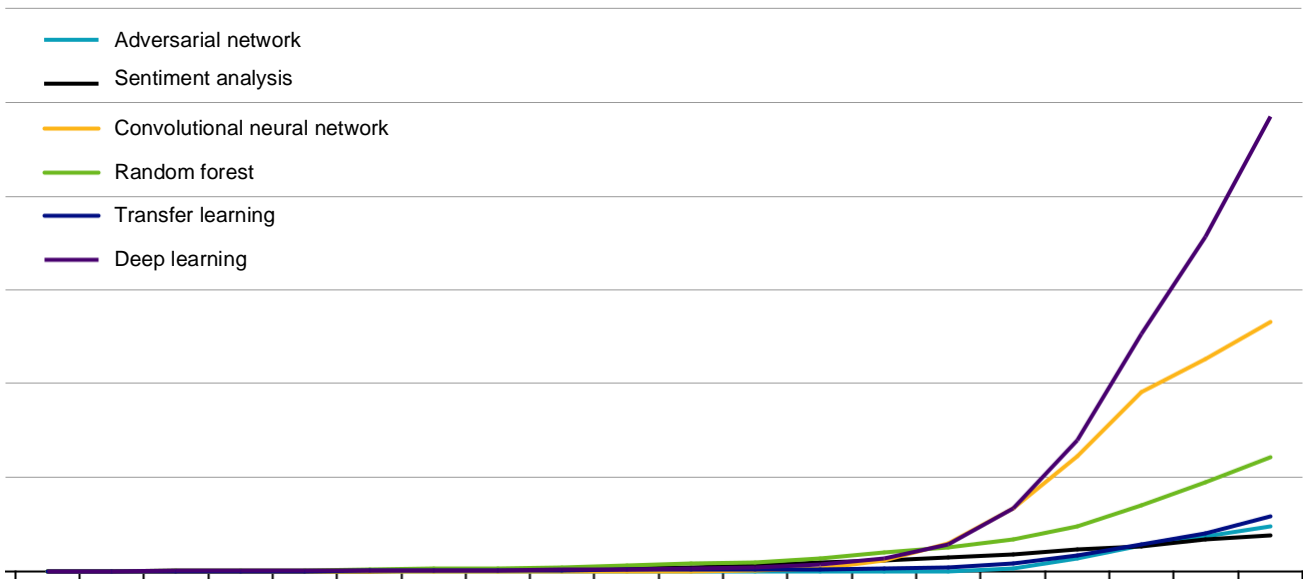


Figure 21. The six fastest-growth artificial intelligence technologies over past 20 years.

Data source: The Lens [32], OECD [122]. Note: The graph shows the number of peer-reviewed publications with the search term appearing in the title, abstract or keywords for journal papers, conference papers, books and book chapters.




Figure 22. Patent citation counts for the top 20 artificial intelligence phrases for 2017–2021.

Data source: The Lens [32].

6 Future development pathways

6.1 Software, hardware and open access resources

Hardware and software toolkits available to scientists and researchers wanting to apply AI in their work are continually improving. This will boost the productivity of advanced AI developers and researchers with limited AI knowledge seeking to perform a particular AI function or procedure. The technology upgrades are having the impact of democratising and industrialising AI. They will play an important role in the diffusion of AI technology across all fields of science and research.

AI accelerators are computing processors specifically designed to handle matrix algebra operations used in machine learning. AI accelerators improve the speed and reduce the latency of AI computations. They enable more time-efficient and cost-efficient development of AI systems. Over the past few years, the number of AI accelerators available on the market has increased in quantity and diversity. They are used in applications such as autonomous vehicles, speech recognition, natural language processing and video object detection [126].

A recent review [126] of AI hardware accelerators by the Massachusetts Institute of Technology describes over 70 AI accelerators under several categories: (a) research chips, (b) very low power chips, (c) embedded chips and systems, (d) autonomous systems, (e) data centre systems, (f) data centre chips and cards. The last category is further broken up into CPUs (with new low-level instructions targeting AI/ML workloads), FPGA-based accelerators, GPU-based accelerators, and dataflow chips, such as Google's Tensor Processing Units (TPUs). In addition to those already on the market, the study reviews 12 AI accelerators that have been announced for future release in the short term, including accelerators from Qualcomm, a large company in the mobile/cellular phone space. The future is likely to see continued improvement in computational power and efficiency of AI systems with an increased diversity of specialised processors suited for various applications.

A related development in computing hardware is the rise of quantum computing. Quantum computers use theories of quantum physics to store and analyse data. Unlike conventional computers, which use a binary (on/ off) system to represent data, quantum computers use qubits which can be in many states at any given point in time. Writing in Nature in 2019, Google scientists reported that their quantum processor called 'sycamore' can solve a problem in 200 seconds that would take a

state-of-the-art classical supercomputer 10,000 years to solve. They refer to this as an 'experimental realisation of quantum supremacy' [127]. The researchers conclude that 'as a result of these developments, quantum computing is transitioning from a research topic to a technology that unlocks new computational capabilities' and that 'we are only one creative algorithm away from valuable near-term applications' [127]. Quantum computers may eventually have the capacity to solve AI problems beyond the reach of conventional computers. This can potentially lead to a paradigm shift and step change in AI capability.

In addition to improved hardware, the AI field is seeing the rapid growth of software frameworks to support AI operations. Examples of popular frameworks include PyTorch, Tensorflow, Keras and Caffe. Using these frameworks in environments such as Python and R, researchers can design and/or adapt machine-learning algorithms relatively quickly, often without the need to delve into the low-level details of the algorithms. This is how many researchers are likely to develop and apply AI within their fields of expertise. These AI frameworks have played, and will continue to play, an important role in facilitating AI technology diffusion across all fields of physical, natural and social sciences. We are also seeing the emergence of code-free AI software tools delivered through graphic user interfaces (GUIs). A research team from Moorfields Eye Hospital in the United Kingdom recently evaluated code-free AI tools for training machine-learning algorithms from corporations such as Amazon, Apple, Clarifai, MedicMind and Microsoft [128]. The code-free deep learning (CFDL) software tools were used for the classification of medical imagery. They conclude 'that CFDL platforms have the potential to improve access to deep learning for both clinicians and biomedical researchers, and represent another step towards the democratization and industrialization of AI' [128]. Furthermore, mass-market software tools such as Microsoft Excel and Microsoft Power BI are increasingly making code-free machine-learning functions available to users to perform common tasks [129, 130].

The open-access frameworks for AI computations are supported by a large and growing number of platforms which enable knowledge sharing. Examples include GitHub, Bitbucket, SourceForge, Gogs, Gitbucket, AWS CodeCommit, Beanstalk, Phabricator, Gitea, Allura, Rhodocode, CodeGiant, Cloud Source Repositories (by Google), Azure DevOps Services, Google Developers and Trac [131]. There are many other such platforms.

These are powerful information resources which speed up and assist scientists developing software code for AI. The volume of material and continuously improving search tools allow a software developer to find a code snippet, library or dataset to quickly solve a problem they're working on. These platforms also facilitate Q&A style discussions where software developers can turn to their community for help. Furthermore, there are platforms such as Kaggle and ImageNet which host competitions for AI experts to access datasets and solve problems. Competitions on these platforms have fast tracked AI problem solving in many areas [132, 133]. Collectively these open-access resources will provide a big boost to AI application in diverse fields of science.

Lastly, the rise of accessible cloud-based computing services is also facilitating the adoption of AI across all fields of research and world regions. A recent report by information technology (IT) consulting firm Gartner finds that the global cloud computing market grew from US\$270 billion in 2020 to a (forecast) US\$397 billion in 2022 with 23% growth during 2021 [134]. The analysts observed that the pandemic fuelled the growth of cloud computing with many business operations moving online. A similar pattern is likely to have occurred within the global science and research community as remote work was needed due to movement and/or quarantine restrictions. Market research by ReportLinker forecasts the continued growth of the cloud rising at an annual compound annual growth rate of 16.3% during 2021 to 2026 and reaching US\$948 billion per year [135]. The science sector is an enthusiastic adopter of cloud computing [136]. The field of genetics, for example, depends on cloud computing for storage, sharing and analysis of vast quantities of data for cross-organisational and international science teams [137].

6.1.1 Implications for science and research organisations

Science organisations seeking to uplift AI capability will need to make decisions about hardware, software and computational infrastructure upgrades, including the access to cloud computing services. These tools have improved substantially over recent years. They are likely to follow a pathway of ongoing improvement in the future. There are many unknowns about quantum computing; it could potentially lead to a step change and paradigm shift in AI resulting in substantially elevated capability. Quantum computing services are already available to AI developers and science organisations will need to factor this into their longer-term AI capability development strategies.

6.2 The quest for better data

We are living in the era of 'big data', where the volume, variety and velocity of data inflows continue to expand. Big data have supported the training of machine-learning algorithms. For example, vast image datasets (labelled by users through search terms) supported the development of image recognition systems able to accurately identify dogs, cats, birds or practically any object within an image. Speech recognition, face recognition and emotion/expression recognition systems have benefited similarly from vast volumes of labelled data.

However, big data can be problematic. Big data contain considerable noise in addition to the signal. Big data can contain spurious entries which are camouflaged and hard to identify amid the other entries. This can degrade the accuracy and reliability of machine-learning models. For example, a recent analysis of 62 published scientific studies using machine learning on chest radiographs and CT scans to detect and prognosticate COVID-19 found that 'none of the models identified are of potential clinical use due to methodological flaws and/or underlying biases' [7]. Most of the problems related to duplication and quality issues in the datasets used for machine learning:

- Incorrectly sourced datasets – In these cases data were incorrectly sourced from demographic age groups that led to biased and inaccurate results when applied at the population level.
- Frankenstein datasets – In these cases public datasets were assembled from numerous other datasets and then redistributed under a new name. This meant algorithms were being trained on multiple identical or overlapping datasets with significant duplication.
- Biased datasets – In these cases images shared publicly and/or contained within published documents often have a form of selection bias. For example, people with certain conditions and/or disease severity may be more/less likely to share their images. This leads to bias in the data.

The masses of data, which on the surface looked like a powerful resource for training machine-learning algorithms, had serious limitations which in every case made the model unusable in clinical settings. Similar problems have been observed in earlier reviews of AI-based models for COVID-19 diagnosis and prognosis. For example, another review of 169 studies containing 232 machine-learning models found 'all models were rated at high or unclear risk of bias, mostly because of non-representative selection of control patients, exclusion of patients who

had not experienced the event of interest by the end of the study, high risk of model overfitting, and unclear reporting' [138]. Most of these problems stemmed from the incorrect use of datasets that were not fit-for-purpose.

Given such findings the AI science community in healthcare, and other domains, is likely to invest greater effort in developing higher quality and fit-for-purpose datasets. Professor Luciano Floridi from the University of Oxford and Alan Turing Institute recently published a paper [139] on the near-term future for AI. One of the trends he identifies is the move from 'big data' to 'small data' – and he defines small data as being higher quality, well curated and provenance assured. He gives the example of an AI-based system developed by Google's DeepMind in partnership with Moorfields Eye Hospital in London. Historically, medical imaging diagnostics based on AI typically relied on 'databases of millions of annotated images'. However, the system was successfully trained using only 14,884 eye scans for early detection of sight-threatening eye diseases and was found to reach or exceed diagnostic accuracy by experts and is considered clinically applicable [140]. The dataset of 14,884 scans represents a considerably smaller-than-usual dataset. The dataset wasn't just smaller; it was also well curated, labelled by experts, reviewed/examined by experts, provenance assured, and fit-for-purpose [140].

6.2.1 Implications for science and research organisations

The ability of a research organisation to achieve competitive differentiation and problem-solve with AI will, in large part, be determined by the quality of its datasets. Vast volumes of publicly available data have been, and will remain, important for the development of AI. For example, a CSIRO team used machine learning to identify COVID-19 virus mutations with likely impact on disease severity [141]. This was done on only 0.3% of the available viral data points due to a lack of patient information for the rest [142]. This illustrates the rate-changing potential for high-quality datasets. Specifically, future AI capability uplift will require investments in high-quality data which is fit-for-purpose, provenance assured, validated, up-to-date, and ethically obtained. As such, it is likely to be critical for an organisation to know what datasets it owns, their provenance, reliability and suitable uses, as well as detailed metadata of those datasets.

This means science and research organisations seeking to upgrade AI capability will need to become adept at acquiring, storing, protecting and recording metadata

on the right (top-priority) datasets. Some of the datasets needed will be novel and will require new investment/capture. Other datasets will be historic and may require formatting, validating and fixing. The data imperative means that research organisations, like any organisation, need to move towards becoming increasingly data-driven. This implies changes to business processes, infrastructure, skills and organisational culture.

Strategies about how to become a data-driven organisation are well covered in the management sciences literature [143, 144]. Data-driven organisations need strong capacity to acquire, analyse, interpret protect, store, share and communicate data. Furthermore, data-driven organisations demonstrably use data in decision making to achieve organisational objectives. They also know the value of their current and future-planned data assets. Research organisations need to acquire these traits to achieve aspirations for AI capability upgrades.

6.3 Education, training and capability uplift

The surge in AI development and application is being accompanied by a surge in AI training and education. An analysis of 18 universities across 9 countries found that the number of undergraduate courses teaching students skills necessary to build and deploy AI models doubled from 2016 to 2020, and increased by 42% for postgraduate courses [14]. Similarly, enrolments in introductory courses for AI or machine learning have grown by close to 60% over the same time [14]. Data from the OECD AI Policy Observatory shows the number of AI courses (delivered in English) worldwide increased 80.1% during 2018–2021, and AI now comprises 27.3% of all computer science and IT courses [145].

In Australia the number of AI courses offered by universities has grown 1.2 times over the past 4 years, with 235 courses on offer in 2021 [145]. The University of Queensland and the University of Sydney are ranked among the world's top 100 academic institutions in AI according to the Nature Artificial Intelligence Index. These institutes are placed in 55th and 76th positions, respectively [146]. The universities ranked in the top 10 positions are all in the United States, Germany and the United Kingdom [146]. Australian institutions such as the Australian Institute for Machine Learning at the University of Adelaide are also expanding, with the number of staff increasing from 80 in 2017 to 140 in 2021 [147]. Tertiary educational and vocational training institutes in Australia are offering a growing range of studying opportunities for people seeking to gain AI skills.

In Australia the fastest growth in AI course offerings was observed for master's degrees [148]. The science and research sectors are in direct competition with industry for AI skills. An analysis of job postings in the United States in 2019 found 9.1% of postings were for AI-related positions [149]. These AI-related jobs tended to be higher skilled positions with 80% requiring a 4-year bachelor degree as a minimum requirement [149]. In Australia, 37,587 AI-related job advertisements were posted in 2015–2019 by Adzuna – a leading job advertisement search engine [150]. These jobs tend to be concentrated in Australian states with the largest population, with the notable exception of the Northern Territory, which had a higher rate of AI jobs relative to its population. AI-related positions make up a fraction of the total job advertisements in Australia, accounting for 0.5% of all postings made between 2015 and 2019 [150].

An analysis conducted by the OECD examined the prevalence of AI skills across occupations using LinkedIn member profiles in 2015–2020 [151]. This analysis found Australia ranked in the middle of the list (13th place out of 26 OECD countries) with the highest penetration of AI skills in the United States, followed by Germany and Israel [151]. When it comes to the AI scientific workforce, as measured through the Global AI Talent Tracker by MacroPolo, the majority of AI scientists are currently based in the United States (59%), followed by China (11%) and Europe (10%) [152]. Most AI scientists completed their undergraduate degrees in China (29%), followed by the United States (20%) and Europe (18%). This reveals a strong net-inward movement of AI talent into the United States [152].

A growing number of students in Australia are graduating with degrees in IT or computer science and have skills suitable for AI-related occupations. The number of university graduates with degrees in IT, as a field of education, has been on a steep rise since 2017 after over a decade of slump [153]. Between 2003 and 2013 the number of IT graduates was mainly in decline, which was a warning trend for Australia's transition to the digital economy [154]. In 2013, the trend reversed to a slow growth (2013–2017) and moved into a steep rise since 2017 [153]. By 2020, the number of IT graduates in Australia almost tripled compared to 2013, exceeding 31,700 people [153]. In 2020, IT became the fourth-largest field of education by the number of graduates after management and commerce, society and culture and health [153]. The growth was the highest among the postgraduate overseas students – the number

of IT graduates in this sector grew over 5-times between 2013 and 2020 to over 15,600 students [153]. This can be partly attributed to the COVID-19 pandemic as IT degrees may have been more easily shifted to online education.

Research has shown that workers who have a background in computer science and programming are more supportive of the development of AI [155], suggesting that AI literacy can be an additional important factor in encouraging adoption of AI across the science sector. While there is limited understanding around the level of AI awareness and understanding across the science sector in Australia, the majority of the Australian public report low (62%) or moderate (26%) subjective knowledge of AI [156]. Examining and improving the AI literacy of the workforce will help science and research organisations identify gaps in the current understanding of AI in the scientific workforce and future capability areas that require attention [155, 156].

Another dimension to AI upskilling is the importance of interdisciplinarity [157]. Research projects using AI typically require high levels of interdisciplinarity involving expertise in the science application domain along with specialised expertise in areas such as machine learning, natural language processing, computer vision, robotics and other sub-fields of AI. A study into the interdisciplinary nature of AI finds that 'the relationship between AI and interdisciplinary research must be considered as a two-way street' [157]. The authors of this study note that more effort is going in one direction (applying AI to other research fields) than the other (applying other research fields to AI). Other researchers using bibliometric analyses have observed similar patterns [120]. Both these studies identify a need for improved two-way interdisciplinarity collaboration to achieve improved outcomes from AI for science.

Lastly, it is worth noting that the future AI talent pipeline for science and research organisations does not start at university. Education researchers have found an interest, motivation and capability for science, technology, engineering and mathematics (STEM) expertise – including mathematics and the foundational skills needed for AI – is typically acquired in early childhood learning [158, 159], primary school and high school [160, 161]. Therefore, a longer-term view of the AI talent pipeline requires investment in all lifelong stages of learning.

6.3.1 Implications for science and research organisations

Research organisations will need strategies for talent acquisition and retention, given the strong demand for AI skills. Industry employers are often able to lure skilled AI workers with high salaries. Research organisations will need to offer competitive remuneration packages to attract and retain skilled AI workers. The feasibility of talent acquisition and retention will need consideration by research organisations while making decisions about whether to grow certain types of AI capability.

At some point the supply of AI skills is likely to adjust to meet demand. Many industries are on the same steep AI adoption curves as we have shown for the science sector; they are hungry for the same AI skills.

The number of, and variety of, education and training courses available to scientists seeking to upgrade AI capability is continuing to expand. There is a wide range of course formats: from micro-credentialling or flash courses to acquire specific skills for specific tasks and timeframes, through to longer in-depth courses designed to develop deeper skills and knowledge. Research organisations, and researchers, can take advantage of these educational offerings to upgrade AI capabilities. Furthermore, research organisations may already have staff with professional backgrounds which make them well-suited for a career transition into AI-focused roles. In addition to training

and education, research organisations will need strategies to bolster two-way interdisciplinary collaboration in AI projects. This involves a flow of expertise from AI specialists into science application domains along with the flow of science domain expertise back into the field of AI.

Some of the upskilling required will exist beyond the organisation's immediate sphere of influence. Longer term and broad-based development of the AI talent pipeline will involve developing foundational skills in mathematics and computational logic for children and teenagers in early learning, primary school and high school contexts. That's when an interest, motivation and capability for advanced STEM skills begins to develop. Research organisations can work with schools and learning organisations to promote the foundational education needed for the future AI workforce. In the same way that science and research organisations have developed school engagement programs to raise awareness around STEM careers (e.g. CSIRO's STEM Professionals in Schools program), similar efforts can be used to strengthen the knowledge and understanding around AI-related career opportunities.

There is also value in uplifting societal awareness and understanding of AI. Improving the general knowledge of AI will help create informed users who can better manage the risks, and harness the opportunities, associated with AI technology. It will also help society work towards effective policies, laws and regulations.

6.4 Towards collaborative artificial intelligence

Initial predictions about the impact of technology on the workforce focused on the areas where technologies like AI would substitute and replace humans [162, 163]. However, increasingly the focus has shifted to how humans and AI can work together (augmentation). Human-AI collaboration is a field of research about how humans and AI can meaningfully interact and cooperate to carry out tasks to higher standards than either can achieve alone [164]. Human-AI collaboration can lead to significant productivity gains and expand the bounds of human capacity. IT company Accenture estimated that organisations that invested in human-AI collaboration would increase their revenues by 38% and employment by 10% between 2018–2022 [165]. Human-AI collaboration is important within the science sector as per other industry sectors. The productive use of AI by scientists depends in large part upon the quality of human-AI collaboration and individual, team and organisational levels.

How does human-AI collaboration happen? Partnership on AI, a global multistakeholder organisation, has developed a standardised framework to investigate and characterise human-AI collaboration [166]. Using seven case studies, the framework covers the following: the nature of the collaboration (e.g. the goals of the interaction, how the human and AI are engaged and their level of agency); the situational context of the collaboration (e.g. whether the human and AI are physically co-located, AI awareness, trust in AI system and potential consequences); the AI system characteristics (e.g. whether the AI system is interactive, adaptable, predictable, explainable and human-like); and the characteristics of the human collaborator [166]. Understanding and examining the nature of human-AI collaboration is critical to progressing the responsible design and governance required to ensure safe, reliable and productive design and development of collaborative AI.

FastMRI is an example of a human-AI collaborative system which aims to accelerate the rate at which doctors can acquire brain scans using magnetic resonance imaging (MRI) without compromising on the image quality [166].

This AI system, developed by Facebook and NYU School of Medicine's Department of Radiology, interprets lower quality image data, which has been rapidly acquired, and predicts the missing data to create a higher quality image. This higher resolution image can then be interpreted by the doctor to determine whether an abnormality is present, increasing their productivity, reducing patient time in the MRI scanner and potentially increasing diagnostic accuracy.

Another emerging area of research related to human-AI collaboration looks at how workers perceive future AI developments. Initial surveys and interviews of healthcare professionals, librarians and qualitative researchers, data scientists and the public have revealed a number of common themes around how workers view human-AI collaboration [155, 156, 167-178]. First, there is a generally positive view towards the value that AI can provide. A 2019 survey of the American public found that 79% of participants were either supportive or neutral towards future AI developments, with support for AI strongest in higher educated and higher income cohorts, or those that have a computer science or programming background [155]. This share is even higher in similar surveys that have been conducted in Australia (85% supportive or neutral towards AI) [178].

The positive sentiment is driven by the perceived benefits of AI. A survey of clinicians in South Korea found that 83.4% felt AI would be useful in medicine, particularly for diagnostic purposes [179]. The European Society of Radiology also found that radiologists felt that AI can potentially result in higher productivity, resulting in more available time to spend with patients [168]. AI also opens opportunities for researchers to access and derive greater value from existing large datasets that would otherwise be prohibitively resource-intensive to manually analyse [170-172]. Scientists traditionally use meta-analyses to synthesise findings from a large collection of scientific studies, a highly time and labour-intensive process. AI could provide a means to automate the process [171], helping researchers stay across the fast-moving research landscape.

Studies exploring perceptions of AI in the research community have found that most scientists and researchers do not think AI can, nor should, replicate the research process [169, 172]. Instead, AI systems and human scientists could operate in a 'synergistic partnership' [171-173, 180].

This synergistic relationship acknowledges the unique and complementary strengths of humans and AI systems and provides opportunities to collaborate to address the limitations of humans and AI on their own [167, 169-172, 175, 179-181]. Under a human-AI collaboration scenario, the human scientist would delegate tasks that can be completed more efficiently by an AI system, leaving the human scientist to invest their time and resources into tasks that rely on uniquely human cognitive abilities.

AI systems have an advantage over human workers when it comes to rapidly processing and analysing large masses of information and identifying patterns or relationships [171, 180]. By automating these manual tasks, scientists can have more time for creative and higher order tasks [170, 171]. Conversely, humans perform better in ambiguous or uncertain decision-making contexts [180]. Researchers are sceptical whether an AI system will be able to replicate a human's ability to make complex associations or complete tasks that require subjective judgement or specialised knowledge [172]. This includes tasks that require prior knowledge around the data or domain [169] or those that require human judgement, such as reviewing scientific publications [169-172] or interpreting ambiguous medical results [179, 182].

Surveys of scientists have found that transparency and explainability are critical factors in determining trust in an AI system [169]. Without this transparency, researchers are concerned that AI could produce biased results or exacerbate existing societal inequalities [169]. A group of researchers from IBM and Rensselaer Polytechnic Institute have shown that 'transparency features' are critical in human-AI systems as they help build trust between the user and the machine [183]. These researchers used visualisations to enhance transparency, and in turn, user trust, around the data that goes into the model and the process through which the AI generates a predictive model.

6.4.1 Implications for science and research organisations

The ability to capture value of human-AI collaboration rests upon buy-in from the scientific community and there are several factors that influence this. One of these is trust, particularly in non-data science domains, which relates to AI literacy. Research has shown that workers who have a background in computer science and programming are more supportive of the development of AI [155]. While there is limited understanding around the level of AI awareness and understanding across the science sector in Australia, the majority of the Australian public report low (62%) or moderate (26%) subjective knowledge of AI [156]. Examining and improving the AI literacy of the workforce will help science and research organisations build understanding and acceptance of AI [155, 156]. In general, a greater share of Australians places their trust in AI systems (41%) than other countries (e.g. 35% in the United Kingdom and 33% in the United States) [156]. To maximise the complementary strengths of human scientists and AI systems, AI applications need to be designed and evaluated with the human scientist workflow in mind. In certain cases, AI systems might be used to automate a discrete part of the workflow (e.g. pre-processing large datasets), or work in concert with the human scientists (e.g. validating previous analyses).

Addressing the productivity slump is an area of concern in the contemporary research sector. While AI presents an opportunity to do more with fewer resources and free scientists up for higher value tasks, it is important that this increase in productivity does not coincide with a decline in research quality or impact [171]. For example, using AI might speed up the rate at which scientists can collect and analyse data, but this may not contribute to meaningful advancements in the field if such tools are implemented in the absence of sufficient human oversight and specialised domain input. Multidisciplinary research teams involving data science and domain knowledge specialists will likely feature more heavily to support quality outputs from human-AI collaborations.

6.5 Artificial intelligence workforce diversity

There is a lack of gender and cultural diversity in the AI research workforce. This is evident in both technology corporations and academic/research organisations.

For example, a recent study of gender diversity in AI research workforces based on analysis of arXiv publications by the National Endowment for Science, Technology and the Arts in the United Kingdom found [184]:

- Worldwide, 13.8% of authors on AI research papers are female and the portion of papers written by at least one female author has not increased since 1990.
- Less than 25% of AI researchers are female in most academic institutions with a few exceptions.
- There are relatively few female authors of AI research papers from technology corporations such as Google (11.3%), Microsoft (11.95%) and IBM (15.7%).
- AI research papers with at least one female author tended to be more applied and were more likely to use human terms such as 'fairness, human mobility, mental, health, gender and personality'.

Gender disparity in the AI workforce is similar to the gender disparity in the broader field of computer science and within STEM disciplines. According to the 2021 STEM equity monitor by the Australian Government [185], 28% of the STEM workforce are women, and males in STEM professions earn (on average) \$28,994 per year more than females compared to a pay gap of \$25,534 across all industries. The STEM equity monitor also records changes to gender disparity over time:

- During 2015 to 2019 the proportion of women enrolled in STEM courses at Australian universities increased from 34% to 36%.
- During 2016 to 2020 the proportion of women working across all STEM-qualified industries increased from 24% to 28%.
- During 2016 to 2020 the proportion of managers and senior managers in STEM roles who are female increased from 18% to 23%.

In addition to gender issues the STEM workforce, likely to reflect the AI workforce, has a lack of cultural and ethnic diversity. For example, professional body Science and Technology Australia finds that 'one in 200 Aboriginal or Torres Strait Islander people of working age have a STEM degree – while one in 20 non-Indigenous working age people have a STEM degree' [186]. Australian mathematician Rowena Ball writing on the topic in Australian Quarterly in 2015 says that 'unless this percentage of Indigenous enrolments in STEM is increased what does follow is that Indigenous people are being systematically locked out of high paying jobs in science related -fields' [187]. It also means that Australian science and technology is not capturing the full benefits of Indigenous science and knowledge. This knowledge can help us understand the world and problem-solve in many contexts.

6.5.1 Implications for science and research organisations

Many research organisations acknowledge the lack of gender and ethnic/cultural diversity in AI and STEM workforces as a challenge they are working to address. While there has been some progress, much remains to be done. Improving workforce diversity will be an important developmental pathway for AI capability uplift within research organisations over the coming decade. A review of Australia's strategies to achieve gender equality in STEM was recently published by authors from several universities and non-government organisations (NGOs) [188]. This provides details on the outcomes associated with various strategies and priorities for the journey ahead. There are also initiatives to promote Indigenous science in Australia. For example, the national science agency of Australia, CSIRO, has an Indigenous science program which aims 'to create Indigenous-driven science solutions that support sustainable futures for Indigenous peoples, cultures and Country' [189].

In one project under this program Microsoft, CSIRO and Kakadu National Park rangers are combining AI, science and Indigenous knowledge for environmental management and biodiversity protection [190].

6.6 The rise of ethical expectations and regulations

A recent review of AI ethics policies was published in January 2021 by researchers at the School of Public Policy at the Georgia Institute of Technology [191]. They identified 112 documents prescribing AI ethics principles, frameworks, policies and strategies from 25 countries produced during 2016–2019. The documents were published by governments, companies and NGOs. The top five (of 25) ethics topics covered in these documents were: (a) social responsibility; (b) transparency; (c) bias and fairness; (d) privacy; and (e) safety and reliability. The authors found that ‘public and NGO documents are more participatory in their creation and more engaged with the law’ and that ‘private-sector documents appear to be more concerned with client and customer-related ethical issues that may lend themselves to a technical fix’.

Overall, the study points towards a substantial expansion in AI ethics expectations across all sectors with the public and NGO sectors leaning towards future legislative implications. It complements several earlier studies examining the development of AI ethics policies, laws, guidelines and frameworks across the globe [192-194]. One study found convergence around five ethical principles across the globe: (a) transparency; (b) justice and fairness; (c) non-maleficence; (d) responsibility; and (e) privacy [194]. All these topics feature in the AI ethics principles of the Australian Government, quoted as follows [195]:

- ‘Human, societal and environmental wellbeing: AI systems should benefit individuals, society and the environment.
- Human-centred values: AI systems should respect human rights, diversity, and the autonomy of individuals.
- Fairness: AI systems should be inclusive and accessible, and should not involve or result in unfair discrimination against individuals, communities or groups.
- Privacy protection and security: AI systems should respect and uphold privacy rights and data protection, and ensure the security of data.
- Reliability and safety: AI systems should reliably operate in accordance with their intended purpose.
- Transparency and explainability: There should be transparency and responsible disclosure so people can understand when they are being significantly impacted by AI, and can find out when an AI system is engaging with them.

- Contestability: When an AI system significantly impacts a person, community, group or environment, there should be a timely process to allow people to challenge the use or outcomes of the AI system.
- Accountability: People responsible for the various phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled.’

These principles are identified by the Australian Government as voluntary and intended to be ‘aspirational and to complement – not substitute – existing AI regulations and practices’ [195]. A recent review of the application of these principles was done by researchers at CSIRO [196].

As we look into the future and across the globe, it is possible that the currently voluntary and aspirational AI principles may become regulations and laws. A recent April 2021 paper in the Harvard Business Review [197] explores this issue and opens with the statement ‘Over the last few weeks, regulators and lawmakers around the world have made one thing clear: New laws will soon shape how companies use artificial intelligence’. Examples of recent developments include:

On 31 March 2021 the five main financial regulators in the United States (including the Federal Treasury) issued an information request to financial institutions to provide detailed information on their use of AI and machine learning. They indicated the information provided is to help ensure ‘compliance with applicable laws and regulations’ [198].

On 21 April 2021 the European Union proposed the first legal framework on AI which includes fines of up to 6% of company revenue for non-compliance [197, 199]. Furthermore, the European Union’s general data protection regulation (GDPR) includes articles limiting the use of automated decision systems including requirements related to explainability and contestability.

Clearly there is a considerable pathway ahead before AI-specific laws are enacted across the globe. Sectors such as finance and retail may be at the forefront of these regulations due to their extensive and routine handling of confidential customer data. However, over the coming years and decades AI policies, regulations and laws are likely to increase. The science sector will be impacted along with other sectors; research organisations will need to ensure they are compliant. Furthermore, AI ethics go beyond compliance. There are also rising expectations for ethical AI from society, investors and AI researchers and developers themselves.

However, ensuring that the development and application of AI is both compliant and meets (and exceeds) the ethical expectations of society can be challenging. The AI research community is still working to resolve the operational meaning of concepts such as explainability, transparency, repeatability and interpretability as they apply to machine-learning systems. There is also considerable work underway to develop software and systems to deliver on AI ethics. For example, a recent study [200] identifies and reviews state-of-the-art technologies to enable explainable AI (XAI), including: (a) features-oriented methods, (b) global methods, (c) concept models, (d) surrogate models, (e) local pixel-based methods, and (f) human-centric methods. Another area of technological innovation to achieve improved AI ethics is privacy-preserving analytics. Recent review papers have been published on this rapidly emerging field [201, 202]. There is also a growing body of work and technology development on improved ways to identify and manage bias in machine-learning projects; a recent review paper describes 25 bias mitigation methods [203].

6.6.1 Implications for science and research organisations

The implication arising from this AI development pathway is that the AI ethics performance bar is likely to be higher and more tightly regulated into the future. What are currently voluntary principles and guidelines could become laws

in the future. Societal awareness about the issues and expectations for ethical AI is likely to rise. Over the last several years, governments, companies and not-for-profits have identified principles and expectations for ethical AI. There are high levels of agreement in these principles about transparency, fairness, explainability and privacy.

However, merely signalling an intention to deliver ethical AI may not be sufficient. Delivering on complex ethical requirements will require improved scientific knowledge and technological capability. It will require skills and capability uplift within the AI workforce. Early investment in ethical capability – including technology, skills and cultures – will help research organisations stay ahead of the regulations.

Lastly, there's a complex balance between efforts to ensure the ethics of AI and the development of novel technologies which improve (or save) people's lives. Effective approaches to AI ethics will ensure principles are upheld without limiting the pace or quality of innovation and discovery. Furthermore, many of today's innovative technologies and approaches enabling improved ethical performance – as discussed above – have grown organically within the AI community. This has mostly happened in the absence of laws and regulations. There's much evidence of a strong drive, coming from within the AI research and development community itself, to achieve improved ethical performance.

7 Conclusion

This report has shown how AI application and development within computer science, and all other major fields of science and research, has increased substantially. The growth has been strongest over the past several years. The coming decade is likely to see the growth continue, and AI become more deeply and broadly adopted in most scientific research domains. References to AI techniques in the titles, abstracts and keywords of research papers are likely to (at some point) decline as the technology becomes commonplace. At this point AI technology will be subsumed into application domains as 'business as usual'.

The current surge in AI activity is not without historical precedent. Twice before in history AI research, investment and activity has surged. The peaks were followed by troughs; two AI winters are generally considered to have occurred during 1974–1980 and 1987–1993. Many of the conditions leading up to these two winters are present today; however, there are significant differences. The sheer size and momentum of the current AI boom is unlikely to end anytime soon. There's so much AI-related investment, upskilling, organisational change and policy development that it will be some time before it levels-off.

The uptake of AI technologies within science and research domains holds the promise of productivity improvement. This is much needed as the global science sector is amid an ongoing productivity slump where more research effort is being invested to achieve the same (or fewer) outcomes. The science productivity slump is causing a broader productivity slump across most industries and the entire economy. As a general-purpose technology, AI can improve productivity in all domains of science and research and, therefore, all industries. However, at this stage we still refer to this as 'a promise' for productivity uplift. There is much evidence from case studies that AI is improving the efficiency and effectiveness of science, enabling discoveries to happen faster, safer and at lower cost. However, this empirical evidence is currently not sufficient as incontrovertible proof of the productivity gains of AI.

The implications for science organisations arising from this report are captured under the future development pathways of AI for science. Overall, science organisations have an imperative to upgrade AI capability to remain competitive and capable for the future. This will require education, training, hardware and software upgrades. It will require the development of data assets and changed ways of working to become a more data-driven organisation. It will also require ensuring the development and application of AI is ethically sound and responds to societal expectations, regulations and legislation.

The notion that AI will be doing research by itself seems unlikely. Scientific research requires creativity, judgement, logic and communication skills that lie beyond the reach of current and foreseeable future AI capability. However, human scientists working in harmony with powerful AI technologies (where AI augments human capabilities), are likely to achieve better outcomes, such as a higher rate of scientific discovery.

The economic depression of 1920–1921 resulting from the 'Spanish Flu' of 1918–1920 was followed by the 'Roaring 20s', a decade of unprecedented economic growth. Economic historians [204] studying the Roaring 20s identify the general-purpose technology of electricity as the primary driver of productivity growth in manufacturing. This productivity growth in manufacturing stimulated overall economic growth with spectacular results. It is possible that AI is the general-purpose technology of our time which leads to improved productivity in science which, in turn, improves productivity and growth in the whole economy.

8 References

1. Crew, B. (2020) Google Scholar reveals its most influential papers for 2020 - Artificial intelligence papers amass citations more than any other research topic. *Nature News* (13 July). London.
2. OECD (2022) OECD Statistics. Organisation for Economic Cooperation and Development, Statistics Website (<https://stats.oecd.org/>) accessed 11 March 2022. Paris.
3. OECD (2021) Database of national AI policies. Organisation for Economic Cooperation and Development, Artificial Intelligence Policy Observatory (website accessed on 13/01/2022). Paris.
4. Van Roy, V., et al. (2021) AI Watch - National strategies on Artificial Intelligence: A European perspective, 2021 edition. European Commission. Publications Office of the European Union, Luxembourg.
5. Malcolm, E. (2021) Flexible solar panels: new stretch of the imagination. *CSIROscope* (18 May 2021). Canberra.
6. Bayat, A., et al. (2020) VariantSpark: Cloud-based machine learning for association study of complex phenotype and large-scale genomic data. *GigaScience* 9(8): 1-12.
7. Roberts, M., et al. (2021) Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. *Nature Machine Intelligence* 3(3): 199-217.
8. Nicholson, K. and A. Slonim (2022) Artificial intelligence: your questions answered. Australian Strategic Policy Institute. Australia.
9. UKRI (2022) Strategic Priorities Fund - Artificial intelligence - AI and Data Science for Science, Engineering, Health and Government. United Kingdom Research and Innovation, United Kingdom Government (website accessed 4 March 2022, <https://www.ukri.org/our-work/our-main-funds/strategic-priorities-fund/>). London.
10. AI4SD (2021) Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery (AI4SD). University of Southampton (<https://www.ai3sd.org>, website accessed 7 March 2022). United Kingdom.
11. Nolan, A. (2021) Artificial intelligence and the future of science. Organisation for Economic Cooperation and Development, Artificial Intelligence Policy Observatory (25 October). Paris.
12. Miyagawa, T. and T. Ishikawa (2019) On the Decline of R&D Efficiency (RIETI Discussion Paper Series 19-E-052). Research Institute of Economy, Trade and Industry. Japan.
13. Boeing, P. and P. Hünermund (2020) A global decline in research productivity? Evidence from China and Germany. *Economics Letters* 197: 1-4.
14. Zhang, D., et al. (2021) The AI Index 2021 Annual Report. AI Index Steering Committee, Human-Centered AI Institute, Stanford University. Stanford CA, United States.
15. Littman, M., et al. (2021) Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) Stanford University (<http://ai100.stanford.edu/2021-report>, website accessed 4 March 2022). California, United States.
16. Nichols, J., et al. (2019) Artificial intelligence for science. Argonne National Laboratory. Lemont, Illinois, United States.
17. Nicholson, K. and A. Slonim (2022) Special Report - Artificial intelligence: Your questions answered. The University of Adelaide Australian Institute for Machine Learning, Australian Strategic Policy Research Institute. Adelaide.
18. Gil, Y. and A. Selman (2019) A 20-Year Community Roadmap for Artificial Intelligence Research in the US. Computing Community Consortium (CCC) and Association for the Advancement of Artificial Intelligence (AAAI). Washington DC, United States.
19. Kaul, V., S. Enslin, and S.A. Gross (2020) History of artificial intelligence in medicine. *Gastrointestinal Endoscopy* 92(4): 807-812.
20. Haenlein, M. and A. Kaplan (2019) A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence. *California Management Review* 61(4): 5-14.
21. Buchanan, B.G. (2005) A (Very) Brief History of Artificial Intelligence. *AI Magazine* 26(4): 53.
22. Flasiński, M. (2016) History of Artificial Intelligence, in *Introduction to Artificial Intelligence*, M. Flasiński, Editor, Springer International Publishing: London.
23. Wikipedia (2022) History of artificial intelligence. Wikipedia - The free encyclopedia (website accessed 22 Jan 2022). San Francisco.
24. McCulloch, W.S. and W. Pitts (1943) A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* 5(4): 115-133.
25. Minsky, M. and D. Edmonds (1954) *Stochastic Neural Analogue Reinforcement Calculator*. Princeton University. United States.

26. Turing, A. (1950) Computing machinery and intelligence. *Mind - A quarterly review of psychology and philosophy* 59(236): 433-460.
27. Warwick, K. and H. Shah (2016) Can machines think? A report on Turing test experiments at the Royal Society. *Journal of Experimental & Theoretical Artificial Intelligence* 28(6): 989-1007.
28. McCarthy, J. (2006) The Dartmouth Workshop-- as planned and as it happened. Stanford University (website accessed 25 Jan 2022). United States.
29. Agar, J.O.N. (2020) What is science for? The Lighthill report on artificial intelligence reinterpreted. *The British Journal for the History of Science* 53(3): 289-310.
30. Odagiri, H., Y. Nakamura, and M. Shibuya (1997) Research consortia as a vehicle for basic research: The case of a fifth generation computer project in Japan. *Research Policy* 26(2): 191-207.
31. Newquist, H. (2020) *The Brain Makers: The History of Artificial Intelligence – Genius, Ego, And Greed In The Quest For Machines That Think*. The Relayer Group. New York.
32. lens.org (2022) Search, analyse and manage patent and scholarly data - Lens serves global patent and scholarly knowledge as a public good to inform science and technology enabled problem solving. The Lens Website by Cambia (www.lens.org) accessed on 8 March 2022. Brisbane, Australia.
33. Jefferson, O.A., et al. (2021) Mapping CRISPR-Cas9 public and commercial innovation using The Lens institutional toolkit. *Transgenic Research* 30(4): 585-599.
34. Dutton, T., B. Barron, and G. Boskovic (2018) *Building an AI world. Report on national and regional AI strategies*. Canadian Institute for Advanced Research. Toronto, Canada.
35. Hajkowicz, S., et al. (2019) *Artificial Intelligence: Solving problems, growing the economy and improving our quality of life*. CSIRO. Brisbane, Australia.
36. IDC (2021) *IDC Forecasts Companies to Spend Almost \$342 Billion on AI Solutions in 2021 (Media Release 4 August 2021)*. International Data Corporation. Needham, MA, United States.
37. OECD (2021) *OECD Science, Technology and R&D Statistics: Revealed technology advantage in selected fields*. OECD Science, Technology and R&D Statistics (database). Paris.
38. Jain, S. (2018) *An Overview of Regularization Techniques in Deep Learning (with Python code)*. Analytics Vidhya. Gurugram, India.
39. Zhang, Z. (2018) Improved Adam Optimizer for Deep Neural Networks. 2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS). DOI: 10.1109/IWQoS.2018.8624183.
40. Mukkamala, M.C. and M. Hein, Variants of RMSProp and Adagrad with Logarithmic Regret Bounds, in *Proceedings of the 34th International Conference on Machine Learning*, P. Doina and T. Yee Whye, Editors. 2017, PMLR: Proceedings of Machine Learning Research. p. 2545--2553.
41. Bottou, L. (2012) Stochastic Gradient Descent Tricks, in *Neural Networks: Tricks of the Trade: Second Edition*, G. Montavon, G.B. Orr, and K.-R. Müller, Editors, Springer Berlin Heidelberg: Berlin, Heidelberg.
42. AAS (2022) What is science? (Reviewed by Pauline Ladiges and Oliver Mayo). The Australian Academy of Science (website accessed 4 February 2022). Canberra.
43. UNESCO (2017) *Measuring Scientific and Technological Services (STS): Draft Paper for Consultation*. United Nations Educational, Scientific and Cultural Organisation, Institute for Statistics. Montreal, Quebec, Canada.
44. UIS (2022) *The UNESCO Institute for Statistics UIS.Stat*. United Nations Economic and Social Council website. Paris, France.
45. PwC (2019) *The 2018 global innovation 1000 study*. PwC. London, UK.
46. Bajpai, P. (2021) Which companies spend the most in research and development? *NASDAQ News*. New York City.
47. Scopus (2021) What are Scopus subject area categories and ASJC codes? Elsevier Website (accessed on 3 March 2022, last updated 29 July 2021). Amsterdam, Netherlands.
48. Gao, J., et al. (2021) Potentially long-lasting effects of the pandemic on scientists. *Nature Communications* 12(1): 6188.
49. Riccaboni, M. and L. Verginer (2022) The impact of the COVID-19 pandemic on scientific research in the life sciences. *PLOS ONE* 17(2): e0263001.
50. arXiv (2022) arXiv usage statistics. arXiv website (accessed 1 July 2022), Cornell University. United States.

51. PWC (2022) Trends - Paper Implementations grouped by framework. Papers With Code Website (<https://paperswithcode.com/trends>). United States.
52. Schwartz, S.J. and B.L. Zamboanga (2009) The Peer-Review and Editorial System: Ways to Fix Something That Might Be Broken. *Perspectives on Psychological Science* 4(1): 54-61.
53. ABS (2017) The Australian National Census for 2016 - Data accessed through the ABS table builder professional. Australian Bureau of Statistics (ABS). Canberra.
54. ABS (2021) Research and experimental development, businesses, Australia. Australian Bureau of Statistics. Canberra, Australia.
55. Parham, D. (2009) Empirical Analysis of the Effects of R&D on Productivity - Implications for productivity measurement. Organisation for Economic Cooperation and Development. Paris.
56. Bloom, N., et al. (2020) Are Ideas Getting Harder to Find? *American Economic Review* 110(4): 1104-44.
57. Towse, A., et al. (2017) Time for a change in how new antibiotics are reimbursed: Development of an insurance framework for funding new antibiotics based on a policy of risk mitigation. *Health Policy* 121(10): 1025-1030.
58. Kraus, C.N. (2008) Low hanging fruit in infectious disease drug development. *Current Opinion in Microbiology* 11(5): 434-438.
59. ITER (2022) Advantages of fusion - The next decades are crucially important to putting the world on a path of reduced greenhouse gas emissions. International Thermonuclear Experimental Reactor (website accessed 6 April 2022). France.
60. CCFE (2022) Fusion energy record demonstrates powerplant future (9 February). Culham Centre for Fusion Energy (CCFE), United Kingdom Energy Authority (website accessed 6 April 2022). Oxfordshire, United Kingdom.
61. Degraeve, J., et al. (2022) Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature* 602(7897): 414-419.
62. Katwala, A. (2022) DeepMind Has Trained an AI to Control Nuclear Fusion - The Google-backed firm taught a reinforcement learning algorithm to control the fiery plasma inside a tokamak nuclear fusion reactor. *Wired Magazine* (16 Feb 2022). San Francisco, United States.
63. Kitano, H. (2021) Nobel Turing Challenge: creating the engine for scientific discovery. *npj Systems Biology and Applications* 7(1): 29.
64. Gil, Y., R. King, and H. Kitano (2020) Posing an AI Scientist Grand Challenge: Artificial Intelligence Systems Capable of Nobel-Quality Discoveries (Workshop Summary developed by Workshop Attendees). The Alan Turing Institute (February 2020). London.
65. ATI (2022) The Turing AI scientist grand challenge - Developing AI systems capable of making Nobel quality scientific discoveries highly autonomously at a level comparable, and possibly superior, to the best human scientists by 2050. The Alan Turing Institute, Research Projects (website accessed 6 April 2022). London.
66. Kelling, S., et al. (2009) Data-intensive Science: A New Paradigm for Biodiversity Studies. *BioScience* 59(7): 613-620.
67. Leonelli, S. (2012) Classificatory Theory in Data-intensive Science: The Case of Open Biomedical Ontologies. *International Studies in the Philosophy of Science* 26(1): 47-65.
68. Ahrens, J., et al. (2011) Data-Intensive Science in the US DOE: Case Studies and Future Challenges. *Computing in Science & Engineering* 13(6): 14-24.
69. Pietsch, W. (2015) Aspects of Theory-Ladenness in Data-Intensive Science. *Philosophy of Science* 82(5): 905-916.
70. Nikam, V.B. and B.B. Meshram (2013) Modeling Rainfall Prediction Using Data Mining Method: A Bayesian Approach. Fifth International Conference on Computational Intelligence, Modelling and Simulation. DOI: 10.1109/CIMSim.2013.29.
71. Hassani, H. and E.S. Silva (2015) Forecasting with Big Data: A Review. *Annals of Data Science* 2(1): 5-19.
72. Callebaut, W. (2012) Scientific perspectivism: A philosopher of science's response to the challenge of big data biology. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 43(1): 69-80.
73. Cowls, J. and R. Schroeder (2015) Causation, Correlation, and Big Data in Social Science Research. *Policy & Internet* 7(4): 447-472.
74. Elragal, A. and R. Klischewski (2017) Theory-driven or process-driven prediction? Epistemological challenges of big data analytics. *Journal of Big Data* 4(1): 19.

75. Senior, A.W., et al. (2020) Improved protein structure prediction using potentials from deep learning. *Nature* 577(7792): 706-710.
76. DeepMind (2020) AlphaFold: Using AI for scientific discovery. Deepmind (<https://deepmind.com>) website accessed 27 April 2022. London.
77. Jumper, J., et al. (2021) Highly accurate protein structure prediction with AlphaFold. *Nature* 596(7873): 583-589.
78. AlphaFold and EMBL-EBI, AlphaFold Protein Structure Database, A.a. EMBL-EBI, Editor. 2022: Cambridgeshire, United Kingdom.
79. iNaturalist (2021) A Community for Naturalists. iNaturalist Web page accessed 24 January 2022. Australia.
80. Van Horn, G., et al. (2018) The inaturalist species classification and detection dataset. *Proceedings of the IEEE conference on computer vision and pattern recognition*.
81. Ceccaroni, L., et al. (2019) Opportunities and risks for citizen science in the age of artificial intelligence. *Citizen Science: Theory and Practice* 4(1).
82. McClure, E.C., et al. (2020) Artificial Intelligence Meets Citizen Science to Supercharge Ecological Monitoring. *Patterns* 1(7): <https://doi.org/10.1016/j.patter.2020.100109>.
83. Open Science Collaboration (2015) Estimating the reproducibility of psychological science. *Science* 349(6251): doi:10.1126/science.aac4716.
84. Camerer, C.F., et al. (2016) Evaluating replicability of laboratory experiments in economics. *Science* 351(6280): 1433-1436.
85. Prinz, F., T. Schlange, and K. Asadullah (2011) Believe it or not: how much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery* 10(9): 712-712.
86. Yang, Y., W. Youyou, and B. Uzzi (2020) Estimating the deep replicability of scientific findings using human and artificial intelligence. *Proceedings of the National Academy of Sciences* 117(20): 10762-10768.
87. Stach, E., et al. (2021) Autonomous experimentation systems for materials development: A community perspective. *Matter* 4(9): 2702-2726.
88. Nikolaev, P., et al. (2016) Autonomy in materials research: a case study in carbon nanotube growth. *npj Computational Materials* 2(1): 16031.
89. Castelveccchi, D. (2021) DeepMind's AI helps untangle the mathematics of knots. *Nature* 600(7888): 202-202.
90. Davies, A., et al. (2021) Advancing mathematics by guiding human intuition with AI. *Nature* 600(7887): 70-74.
91. Pascal, L., et al. (2020) A Shiny r app to solve the problem of when to stop managing or surveying species under imperfect detection. *Methods in Ecology and Evolution* 11(12): 1707-1715.
92. Torney, C.J., et al. (2019) A comparison of deep learning and citizen science techniques for counting wildlife in aerial survey images. *Methods in Ecology and Evolution* 10(6): 779-787.
93. Norouzzadeh, M.S., et al. (2018) Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences* 115(25): E5716-E5725.
94. Wilhite, A.W. and E.A. Fong (2012) Coercive Citation in Academic Publishing. *Science* 335(6068): 542-543.
95. Seglen, P.O. (1997) Why the impact factor of journals should not be used for evaluating research. *BMJ (Clinical research ed.)* 314(7079): 498-502.
96. Weis, J.W. and J.M. Jacobson (2021) Learning on knowledge graph dynamics provides an early warning of impactful research. *Nature Biotechnology* 39(10): 1300-1307.
97. Savage, N. (2019) How AI and neuroscience drive each other forwards. *Nature* 571(7766): S15-S15.
98. Yamins, D.L.K., et al. (2014) Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences* 111(23): 8619-8624.
99. Kell, A.J.E., et al. (2018) A Task-Optimized Neural Network Replicates Human Auditory Behavior, Predicts Brain Responses, and Reveals a Cortical Processing Hierarchy. *Neuron* 98(3): 630-644.
100. Rajpurkar, P., et al. (2022) AI in health and medicine. *Nature Medicine* (28): 31-38.
101. Stokes, J.M., et al. (2020) A Deep Learning Approach to Antibiotic Discovery. *Cell* 180(4): 688-702.
102. RAR (2014) Antimicrobial Resistance: Tackling a Crisis for the Health and Wealth of Nations. Review on Antimicrobial Resistance, United Kingdom Government. London.

103. Brown, E.D. and G.D. Wright (2016) Antibacterial drug discovery in the resistance era. *Nature* 529(7586): 336-343.
104. Zhavoronkov, A., et al. (2019) Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology* 37(9): 1038-1040.
105. VoPham, T., et al. (2018) Emerging trends in geospatial artificial intelligence (geoAI): potential applications for environmental epidemiology. *Environmental Health* 17(1): 40.
106. Lin, Y., et al. (2017) Mining public datasets for modeling intra-city PM_{2.5} concentrations at a fine spatial resolution. *Proceedings of the 25th ACM SIGSPATIAL international conference on advances in geographic information systems*.
107. Checco, A., et al. (2021) AI-assisted peer review. *Humanities and Social Sciences Communications* 8(1): 25.
108. Frontiers (2020) Artificial Intelligence to help meet global demand for high-quality, objective peer-review in publishing. *Frontiers Science News*. Lausanne, Switzerland.
109. Vincent, J. (2022) A new use for AI: Summarizing scientific research for seven-year-olds. *The Verge* (18 January 2022). Washington, District of Columbia, United States.
110. Bickler, S.H. (2021) Machine Learning Arrives in Archaeology. *Advances in Archaeological Practice* 9(2): 186-191.
111. Chetouani, A., et al. (2020) Classification of engraved pottery sherds mixing deep-learning features by compact bilinear pooling. *Pattern Recognition Letters* 131: 1-7.
112. Romanengo, C., S. Biasotti, and B. Falcidieno (2020) Recognising decorations in archaeological finds through the analysis of characteristic curves on 3D models. *Pattern Recognition Letters* 131: 405-412.
113. Bonhage, A., et al. (2021) A modified Mask region-based convolutional neural network approach for the automated detection of archaeological sites on high-resolution light detection and ranging-derived digital elevation models in the North German Lowland. *Archaeological Prospection* 28(2): 177-186.
114. Davis, D.S., R.J. DiNapoli, and K. Douglass (2020) Integrating Point Process Models, Evolutionary Ecology and Traditional Knowledge Improves Landscape Archaeology—A Case from Southwest Madagascar. *Geosciences* 10(287): 1-25.
115. Mirhoseini, A., et al. (2021) A graph placement methodology for fast chip design. *Nature* 594(7862): 207-212.
116. Ellegaard, O. and J.A. Wallin (2015) The bibliometric analysis of scholarly production: How great is the impact? *Scientometrics* 105(3): 1809-1831.
117. Hu, Q., S. Khosa, and N. Kapucu (2015) The Intellectual Structure of Empirical Network Research in Public Administration. *Journal of Public Administration Research and Theory* 26(4): 593-612.
118. Donthu, N., et al. (2021) How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research* 133: 285-296.
119. Wang, K., et al. (2020) Microsoft Academic Graph: When experts are not enough. *Quantitative Science Studies* 1(1): 396-413.
120. Frank, M.R., et al. (2019) The evolution of citation graphs in artificial intelligence research. *Nature Machine Intelligence* 1(2): 79-85.
121. Liu, N., P. Shapira, and X. Yue (2021) Tracking developments in artificial intelligence research: constructing and applying a new search strategy. *Scientometrics* 126(4): 3153-3192.
122. Baruffaldi, S., et al. (2020) Identifying and measuring developments in artificial intelligence. *Organisation for Economic Cooperation and Development*. Paris.
123. Penfold, R. (2020) Using the Lens Database for Staff Publications. *Journal of the Medical Library Association* : JMLA 108(2): 341-344.
124. Jefferson, O.A., et al. (2021) Mapping innovation trajectories on SARS-CoV-2 and its variants. *Nature Biotechnology* 39(4): 401-403.
125. Geroski, P.A. (2000) Models of technology diffusion. *Research Policy* 29(4): 603-625.
126. Reuther, A., et al. (2020) Survey of Machine Learning Accelerators. 2020 IEEE High Performance Extreme Computing Conference (HPEC). DOI: 10.1109/HPEC43674.2020.9286149.
127. Arute, F., et al. (2019) Quantum supremacy using a programmable superconducting processor. *Nature* 574(7779): 505-510.
128. Korot, E., et al. (2021) Code-free deep learning for multi-modality medical image classification. *Nature Machine Intelligence* 3(4): 288-298.
129. Microsoft (2022) Artificial Intelligence sample for Power BI: Take a tour (2 April 2022). Microsoft Corporation. Seattle, United States.

130. Microsoft (2018) Bringing AI to Excel—4 new features announced today at Ignite. Microsoft Corporation. Seattle, United States.
131. Martin, M. (2022) Fifteen best Github alternatives in 2022. Guru99 Blog. Wilmington, Delaware, United States of America.
132. Bojer, C.S. and J.P. Meldgaard (2021) Kaggle forecasting competitions: An overlooked learning opportunity. *International Journal of Forecasting* 37(2): 587-603.
133. Mishkin, D., N. Sergievskiy, and J. Matas (2017) Systematic evaluation of convolution neural network advances on the Imagenet. *Computer Vision and Image Understanding* 161: 11-19.
134. Gartner (2021) Gartner Forecasts Worldwide Public Cloud End-User Spending to Grow 23% in 2021. Gartner Newsroom - Press Release 21 April 2021. Stamford, Connecticut, United States.
135. ReportLinker (2021) Cloud Computing Market by Service, Deployment Model, Organization Size, Vertical And Region - Global Forecast to 2026. ReportLinker- Market Intelligence Platform (<https://ai.reportlinker.com/>) Lyon, France.
136. Drake, N. (2014) Cloud computing beckons scientists. *Nature* 509(7502): 543-544.
137. Langmead, B. and A. Nellore (2018) Cloud computing for genomic data analysis and collaboration. *Nature Reviews Genetics* 19(4): 208-219.
138. Wynants, L., et al. (2020) Prediction models for diagnosis and prognosis of covid-19: systematic review and critical appraisal. 369: m1328.
139. Floridi, L. (2019) What the Near Future of Artificial Intelligence Could Be. *Philosophy & Technology* 32(1): 1-15.
140. De Fauw, J., et al. (2018) Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine* 24(9): 1342-1350.
141. Ramarao-Milne, P., et al. (2022) Data-driven platform for identifying variants of interest in COVID-19 virus. *Computational and Structural Biotechnology Journal* 20: 2942-2950.
142. Bauer, D.C., et al. (2021) Interoperable medical data: The missing link for understanding COVID-19. *Transboundary and Emerging Diseases* 68(4): 1753-1760.
143. Kiron, D. (2016) Lessons from Becoming a Data-Driven Organization. MIT Sloan Management Review (18 October 2016). Massachusetts Institute of Technology, Cambridge, MA.
144. Anderson, C. (2015) Creating a data-driven organisation. O'Reilly Media Inc. Sebastopol, California.
145. OECD (2022) OECD.AI Policy Observatory: AI courses in English by educational level. OECD. AI Policy Observatory. Paris, France.
146. Nature Index (2020) Top 100 academic institutions in artificial intelligence. *Nature Index 2020*. London.
147. University of Adelaide (2020) Australian Institute for Machine Learning. Annual report 2020. University of Adelaide. Adelaide.
148. DandoloPartners (2020) Evaluation of early learning and schools initiatives in the National Innovation and Science Agenda. Report to the Department of Education. Department of Education, Australian Government. Canberra.
149. Toney, A. and M. Flagg (2020) U.S. Demand for AI-Related Talent. Center for Security and Emerging Technology (CSET). Washington D.C.
150. Mason, C., et al. (2022) Australian Skills Dashboard. Commonwealth Scientific and Industrial Research Organisation. Brisbane.
151. OECD (2022) Live data: Cross-country AI skills penetration (accessed on 27/01/2022). Organisation for Economic Cooperation and Development, AI Policy Observatory. Paris.
152. MacroPolo (2022) The Global AI Talent Tracker. MacroPolo (<https://macropolo.org/digital-projects/the-global-ai-talent-tracker>). Chicago.
153. DESE (2022) Award Course Completions Time Series. Department of Education, Skills and Employment, Australian Government. Canberra.
154. Hajkowicz, S., et al. (2016) Tomorrow's Digitally Enabled Workforce: Megatrends and scenarios for jobs and employment in Australia over the coming twenty years. CSIRO. Brisbane, Australia.
155. Zhang, B. and A. Dafoe, Artificial intelligence: American attitudes and trends. 2019, University of Oxford: Oxford, United Kingdom.
156. Gillespie, N., S. Lockey, and C. Curtis (2021) Trust in artificial Intelligence: a five country study. University of Queensland and KPMG. Brisbane, Australia.

157. Kusters, R., et al. (2020) Interdisciplinary Research in Artificial Intelligence: Challenges and Opportunities. *Frontiers in Big Data 3*: DOI: <https://doi.org/10.3389/fdata.2020.577974>.
158. Wang, M.-T. and J.L. Degol (2017) Gender Gap in Science, Technology, Engineering, and Mathematics (STEM): Current Knowledge, Implications for Practice, Policy, and Future Directions. *Educational Psychology Review 29*(1): 119-140.
159. Simoncini, K. and M. Lasen (2018) Ideas About STEM Among Australian Early Childhood Professionals: How Important is STEM in Early Childhood Education? *International Journal of Early Childhood 50*(3): 353-369.
160. Tai, R.H., et al. (2006) Planning early for careers in science. *Science 312*(5777): 1143-1144.
161. Maltese, A.V. and R.H. Tai (2011) Pipeline persistence: Examining the association of educational experiences with earned degrees in STEM among US students. *Science education 95*(5): 877-907.
162. McKinsey & Company (2017) Jobs lost, jobs gained: Workforce transitions in a time of automation. McKinsey & Company. New York, United States.
163. Frey, C.B. and M.A. Osborne (2017) The future of employment: how susceptible are jobs to computerisation? *Technological Forecasting and Social Change 114*(1): 254-280.
164. Dellermann, D., et al. (2019) The Future of Human-AI Collaboration: A Taxonomy of Design Knowledge for Hybrid Intelligence Systems. arXiv (arXiv:2105.03354). Cornell University.
165. Shook, E. and M. Knickrehm (2018) Reworking the revolution. Accenture. Sydney.
166. Partnership on AI (2019) Collaborations Between People and AI Systems (CPAIS): Human - AI Collaboration Framework and Case Studies. Partnership on AI (<https://partnershiponai.org/>). San Francisco, California, United States.
167. Castagno, S. and M. Khalifa (2020) Perceptions of Artificial Intelligence Among Healthcare Staff: A Qualitative Survey Study. *Frontiers in Artificial Intelligence 3*(84): DOI: <https://doi.org/10.3389/frai.2020.578983>.
168. European Society of Radiology (2019) Impact of artificial intelligence on radiology: a EuroAIM survey among members of the European Society of Radiology. *Insights Imaging 10*(105): <https://doi.org/10.1186/s13244-019-0798-3>.
169. Wang, D., et al. (2019) Human-AI collaboration in data science: Exploring data scientists' perceptions of automated AI. *Proceedings of the ACM on Human-Computer Interaction 3*(CSCW): 1-24.
170. Procter, R., B. Glover, and E. Jones (2020) Research 4.0: Research in the age of automation. Demos. London, United Kingdom.
171. Chubb, J., P. Cowling, and D. Reed (2021) Speeding up to keep up: exploring the use of AI in the research process. *AI & Society* DOI: <https://doi.org/10.1007/s00146-021-01259-0>.
172. Feuston, J.L. and J.R. Brubaker (2021) Putting Tools in Their Place: The Role of Time and Perspective in Human-AI Collaboration for Qualitative Analysis. *Proceedings of the ACM on Human Computer Interaction 5*(CSCW2): 1-25.
173. Jiang, J.A., et al. (2021) Supporting Serendipity: Opportunities and Challenges for Human-AI Collaboration in Qualitative Analysis. *Proceedings of the ACM on Human-Computer Interaction 5*(CSCW1): 1-23.
174. Hervieux, S. and A. Wheatley (2021) Perceptions of artificial intelligence: A survey of academic librarians in Canada and the United States. *The Journal of Academic Librarianship 47*(1): 1-11.
175. Lund, B.D., et al. (2020) Perceptions toward Artificial Intelligence among Academic Library Employees and Alignment with the Diffusion of Innovations' Adopter Categories. *College & Research Libraries 81*(5): DOI: <https://doi.org/10.5860/crl.81.5.865>.
176. Wood, B.A. and D. Evans (2018) Librarians' perceptions of artificial intelligence and its potential impact on the profession. Kennesaw State University. Kennesaw, GA.
177. Cox, A.M., S. Pinfield, and S. Rutter (2019) The intelligent library: Thought leaders' views on the likely impact of artificial intelligence on academic libraries. *Library Hi Tech 37*(3): 418-435.
178. Selwyn, N., et al. (2020) AI for Social Good-Australian Attitudes Toward AI and Society Report. pdf. Monash Data Futures Institute, Monash University. Victoria, Australia.
179. Oh, S., et al. (2019) Physician Confidence in Artificial Intelligence: An Online Mobile Survey. *Journal of Medical Internet Research 21*(3): DOI: 10.2196/12422.
180. Jarrahi, M.H. (2018) Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons 61*(4): 577-586.

181. Lubars, B. and C. Tan (2019) Ask not what AI can do, but what AI should do: Towards a framework of task delegability. arXiv preprint arXiv:1902.03245.
182. Krittanawong, C. (2018) The rise of artificial intelligence and the uncertain future for physicians. *Eur J Intern Med* 48: e13-e14.
183. Drozdal, J., et al. (2020) Trust in AutoML: Exploring Information Needs for Establishing Trust in Automated Machine Learning Systems.
184. Stathoulopoulos, K. and J. Mateos-Garcia (2019) Gender Diversity in AI Research. National Endowment for Science, Technology and the Arts. London.
185. DISER (2021) STEM Equity Monitor 2021, a national data report on girls' and women's participation in science, technology, engineering and mathematics (STEM). Australian Government Department of Industry Science Energy and Resources. Canberra.
186. STA (2020) Diversity gains in Australia's STEM workforce - But more needed. Science and Technology Australia (website accessed 8 February 2022). Canberra.
187. Ball, R. (2015) STEM the gap: Science belongs to us mob too. *Australian Quarterly* January-March: 13-19.
188. Latimer, J., et al. (2019) Australia's strategy to achieve gender equality in STEM. *The Lancet* 393(10171): 524-526.
189. CSIRO (2022) Indigenous science (web page). CSIRO (Website accessed 8 February 2022, <https://www.csiro.au/en/research/indigenous-science>). Canberra.
190. CSIRO (2019) AI transforms Kakadu management - CSIRO, Kakadu rangers and Microsoft meld science, Indigenous knowledge and technology in pioneering program. CSIRO Media Release (20 November). Canberra.
191. Schiff, D., et al. (2021) AI Ethics in the Public, Private, and NGO Sectors: A Review of a Global Document Collection. *IEEE Transactions on Technology and Society* 2(1): 31-42.
192. Schiff, D., et al. (2020) What's Next for AI Ethics, Policy, and Governance? A Global Overview. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society. New York. Association for Computing Machinery. DOI: 10.1145/3375627.3375804.
193. Hagendorff, T. (2020) The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds and Machines* 30(1): 99-120.
194. Jobin, A., M. Ienca, and E. Vayena (2019) The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1(9): 389-399.
195. DISER (2022) Australia's Artificial Intelligence Ethics Framework. Department of Industry, Science and Energy Resources, Australian Government. Canberra.
196. Sanderson, C., et al. (2022) AI Ethics Principles in Practice: Perspectives of Designers and Developers. arXiv:2112.07467v2 (<https://doi.org/10.48550/arXiv.2112.07467>).
197. Burt, A. (2021) New AI Regulations Are Coming. Is Your Organization Ready? *Harvard Business Review* (April 30, 2021).
198. Treasury (2021) Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, Including Machine Learning. Department of the Treasury, United States Government. Washington DC, United States.
199. EC (2021) Proposal for a Regulation laying down harmonised rules on artificial intelligence. European Commission (Policy and Legislation, Publication 21 April 2021). Brussels.
200. Angelov, P.P., et al. (2021) Explainable artificial intelligence: an analytical review. 11(5): e1424.
201. Briggs, C., Z. Fan, and P. Andras (2021) A Review of Privacy-Preserving Federated Learning for the Internet-of-Things, in *Federated Learning Systems: Towards Next-Generation AI*, M.H. Rehman and M.M. Gaber, Editors, Springer International Publishing: London.
202. Boulemtafes, A., A. Derhab, and Y. Challal (2020) A review of privacy-preserving techniques for deep learning. *Neurocomputing* 384: 21-45.
203. Fahse, T., V. Huber, and B. van Giffen (2021) Managing Bias in Machine Learning Projects. *Innovation Through Information Systems*. Cham. Springer International Publishing.
204. Harrison, S. and M. Weder (2009) Technological change and the roaring twenties: A neoclassical perspective. *Journal of Macroeconomics* 31(3): 363-375.

Appendix A – Science occupations

Considering the four types of science (natural science, social science, formal science and applied science) we identified corresponding occupations from the Australian and New Zealand Standard Classification of Occupations (ANZSCO). We sought occupations where the description involved research activity and was consistent with tasks comprising the scientific process. The occupations we identified (with ANZSCO codes in brackets) included:

- science technicians (3114)
- natural and physical science professionals (234)
- economists (224311)
- social professionals (272499)
- mathematicians (224112)
- statisticians (224113).

The category of natural and physical science professionals contains sub-categories of

- agricultural, fisheries and forestry scientists (2341)
- chemists, and food and wine scientists (2342)
- environmental scientists (2343)
- geologists, geophysicists and hydrogeologists (2344)
- life scientists (2345)
- medical scientists (2346)
- veterinarians (2347)
- and other natural and physical science professionals (2349).

The 'social professional' category captures social scientists conducting research in diverse fields: anthropologists, criminologists, geographers, political scientists, sociologists and others. We note that some of these categories may contain relatively small numbers of non-science occupations that we could not separate out.

Appendix B – Artificial intelligence phrases

During 2018–2019 the Organisation for Economic Co-operation and Development (OECD) held several expert working groups to identify, review and refine a set of phrases to search for artificial intelligence (AI) patents and scholarly publications. In the final meeting in 2019 the patent examiners and AI experts convened by the OECD went through the final draft list of phrases to validate and challenge them. The final set of phrases was published by the OECD [122]. We have also used these phrases to search for AI publications in this report. The phrases are listed below.

action recognition	cognitive computing	feature engineering
activity recognition	cognitive insight system	feature extraction
adaboost	cognitive modelling	feature learning
adaptive boosting	collaborative filtering	feature selection
adversarial network	collision avoidance	firefly algorithm
ambient intelligence	community detection	fuzzy c
ant colony	computational intelligence	fuzzy environment
ant colony optimisation	computational pathology	fuzzy logic
artificial bee colony algorithm	computer vision	fuzzy number
artificial intelligence	convolutional neural network	fuzzy set
artificial neural network	cyber physical system	fuzzy system
association rule	data mining	gaussian mixture model
autoencoder	decision tree	gaussian process
autonomic computing	deep belief network	generative adversarial network
autonomous vehicle	deep convolutional neural network	genetic algorithm
autonomous weapon	deep learning	genetic programming
backpropagation	deep neural network	gesture recognition
Bayesian learning	dictionary learning	gradient boosting
bayesian network	differential evolution algorithm	gradient tree boosting
bee colony	dimensionality reduction	graphical model
biped robot	dynamic time warping	gravitational search algorithm
blind signal separation	emotion recognition	hebbian learning
bootstrap aggregation	ensemble learning	hidden Markov model
brain computer interface	evolutionary algorithm	hierarchical clustering
brownboost	evolutionary computation	high-dimensional data
chatbot	extreme machine learning	high-dimensional feature
classification tree	face recognition	high-dimensional input
cluster analysis	facial expression recognition	high-dimensional model
cognitive automation	factorisation machine	high-dimensional space

high-dimensional system	lpboost	particle swarm optimisation
human action recognition	machine intelligence	pattern recognition
human activity recognition	machine learning	pedestrian detection
human aware artificial intelligence	machine translation	policy gradient methods
humanoid robot	machine vision	Q-learning
human-robot interaction	madaboost	quadruped robot
image classification	MapReduce	random field
image processing	Markovian	random forest
image recognition	memetic algorithm	rankboost
image retrieval	meta learning	recommender system
image segmentation	motion planning	recurrent neural network
independent component analysis	multi task learning	regression tree
inductive monitoring	multi-agent system	reinforcement learning
industrial robot	multi-label classification	relational learning
instance-based learning	multi-layer perceptron	robot
intelligence augmentation	multinomial naive Bayes	rough set
intelligent agent	multi-objective evolutionary algorithm	rule learning
intelligent classifier	multi-objective optimisation	rule-based learning
intelligent geometric computing	multi-sensor fusion	self-organising map
intelligent infrastructure	naive Bayes classifier	self-organising structure
intelligent software agent	natural gradient	semantic web
intuitionistic fuzzy set	natural language generation	semi-supervised learning
Kernel learning	natural language processing	sensor data fusion
K-means	natural language understanding	sensor fusion
latent dirichlet allocation	nearest neighbour algorithm	sentiment analysis
latent semantic analysis	neural network	service robot
latent variable	neural turing	similarity learning
layered control system	neural turing machine	simultaneous localisation mapping
learning automata	neuromorphic computing	single-linkage clustering
legged robot	non negative matrix factorisation	social robot
link prediction	object detection	sparse representation
logitboost	object recognition	spectral clustering
long short term memory (LSTM)	obstacle avoidance	speech recognition

speech to text
stacked generalisation
statistical relational learning
stochastic gradient
supervised learning
support vector machine
support vector regression
swarm intelligence
swarm optimisation
t s fuzzy system
Takagi-Sugeno fuzzy systems
temporal difference learning
text mining
text to speech
topic model
totalboost
trajectory planning
trajectory tracking
transfer learning
trust region policy optimisation
unmanned aerial vehicle
unsupervised learning
variational inference
vector machine
virtual assistant
visual servoing
wheeled mobile robot
xgboost