December 2019

Machine Learning Innovation

A Patent Analytics Report

Prepared by IP Australia

Yohan Ramasundara
President, ACS

Andrew Johnson
Chief Executive Officer, ACS

Foreword

Twelve months ago, ACS released ***Artificial Intelligence - A Starter Guide To The Future Of Business***. We wanted to assist businesses to position themselves for a future that optimises machine-readable data patterns, automates processes wherever possible, and enables new digital product development that can enhance Australia's exports.

The use cases identified in the Starter Guide covered Agriculture (AI-controlled field robots, crop and soil monitoring, predictive analytics); Finance (algorithmic trading, Robo-advisors, deep data analytics); Health care (diagnostics, augmented decision making, AI-managed administration); Mining (mineral exploration, autonomous equipment, automated sorting and optimised metallurgy, predictive maintenance); Manufacturing (self-learning monitoring, predictive maintenance, intelligent supply chain optimisation, AI-driven product design); Retail (product recommendations, product identification and personalisation, demand prediction, customer engagement); Communications and Entertainment (automated content creation, personalised marketing and advertising, automated telemarketing and polling); and Transport (autonomous vehicles and driver assistance, engineer monitoring and predictive maintenance, safety and efficacy).

Since the time of the Starter Guide's release, we have seen revenue from the global AI market forecast to increase more than tenfold from US $9.5 billion in 2018 to an estimated US$118 billion by 2025[i]. We have witnessed 14 of the world's most advanced economies announce over AU$86 billion in focused AI programs and activities[ii].

As momentum builds in this fourth industrial revolution, ACS wanted to move beyond the hype and undertake an evidenced-based investigation into machine learning progress.

We considered there to be no better place to research than patents – both as a guide to see where the world's innovators believe there are gaps in the market that can be commercialised through machine learning technologies, as well as potentially indicate those industries and use cases that appear untapped, thus providing innovators with insight as to where early leader advantage may still be achieved.

Having worked with the Patent Analytics Hub of IP Australia previously on blockchain patent trends, it was once again our privilege to now explore machine learning and artificial intelligence.

[i] Tractica (30 April 2019), Artificial Intelligence Software Market to Reach $118.6 Billion in Annual Worldwide Revenue by 2025 https://www.tractica.com/newsroom/press-releases/artificial-intelligence-software-market-to-reach-118-6-billion-in-annual-worldwide-revenue-by-2025/

[ii] Artificial Intelligence Roadmap, Australia's artificial intelligence roadmap, developed by CSIROs Data61 for the Australian Government.

Contents

01

SUMMARY

Summary

Machine learning technologies underpin artificial intelligence (AI).

Machine learning algorithms are applied to make predictions and conclusions when provided with data and allow computer systems to make intelligent decisions.[1]

Recent patent data shows machine learning technology is an area of outstanding growth and investment across the world. Since 2012, patenting of AI and machine learning related technologies has increased significantly, with a four-fold global increase in patent filings over five years from 2012–16. Applicants are commonly large corporations or universities.

China is both the largest patent filing destination and the largest source of patent family filings in machine learning, with the State Grid Corporation of China playing the largest role in this sector.

The United States, South Korea and Japan are the other major players in this sector, with Australia ranking tenth in the world as a patent filing destination.

In real world applications of machine learning, the telecommunications sector had the most patent filings (17 per cent of patents filed), likely reflecting the growth and net worth of this sector in the global economy. This is mirrored by development of core capabilities in image and video analysis (36 per cent of patents filed).

Underlying these advances, the most significant developments in foundational technologies since 2012 were found in supervised or semi-supervised learning techniques (38 per cent of all patent filings).

This report outlines findings from analysis of patents filed since 2012, including trends, innovators, filing destinations and commercial players in this space.

[1] CSIRO Data61, Artificial Intelligence and Machine Learning, Focus Area (accessed 9 October 2019) https://www.data61.csiro.au/en/Our-Work/AI-and-Machine-Learning

KEY REPORT FINDINGS

36,740 patent families have been filed in machine learning technologies since 2012.

25,319 originate from China.

Machine learning patent filings have grown 27% or more every year since 2012.

91% of AI and machine learning related patent families are in an active state.

The State Grid Corporation of China is the top global innovator, with 1,018 active patent families (88% are active).

China is the largest filing destination in the world, while Australia is 10th.

Australian innovators rank 17th globally, with 59 patent families.

The telecommunications sector leads real world applications of machine learning, with 17% of patent applications.

Image and video analysis is the largest core capability application, with 36% of patent applications.

02

Introduction

AI is the simulation of human intelligence processes by machines. It is the term given to a collection of interrelated technologies, modern algorithms and forms of artificial programs or systems capable of perception, cognition, decision making and implementation without explicit guidance from a human being.[2] Machine learning is a form of AI that enables a system – or machine – to learn from data using algorithms. These algorithms refine the system by processing new data. As a result, the system is trained to adapt to changing patterns and associations in the data, allowing intelligent predictions and decisions to be made.[3]

Revenue from the global AI market is forecast to increase more than tenfold from US $9.5 billion in 2018 to an estimated $118 billion by 2025.[4] Understanding how AI and machine learning technologies are commercialised, identifying potential partners and determining what capability currently exists in Australia will be important in assisting the Australian ICT sector to leverage opportunities both nationally and internationally as this technology matures.

[2] Dawson, D. et al. (2019) Artificial Intelligence: Australia's Ethics Framework (A Discussion Paper). Data61 CSIRO, Australia https://consult.industry.gov.au/strategic-policy/artificial-intelligence-ethics-framework/supporting_documents/ArtificialIntelligenceethicsframeworkdiscussionpaper.pdf

[3] SAS, Machine Learning Analysis Insight (accessed 9 October 2019) https://www.sas.com/en_au/insights/analytics/machine-learning.html

[4] Tractica (30 April 2019), Artificial Intelligence Software Market to Reach $118.6 Billion in Annual Worldwide Revenue by 2025 https://www.tractica.com/newsroom/press-releases/artificial-intelligence-software-market-to-reach-118-6-billion-in-annual-worldwide-revenue-by-2025/

Why patent data?

Patents can be used as indicators of innovative activity. By extracting and analysing data associated with patent documents, we can measure aspects of inventive activity such as scope, intensity, collaboration and impact. These metrics can be developed across technology sectors and by measures including individuals (inventors), institutions (applicants), countries and regions.

Patents are granted for devices, substances, methods or processes that are new, inventive and useful, giving exclusive commercial rights in exchange for full public disclosure of the invention. This means patents are a source of data on innovation trends in science and technology. More information about the patent system is given in Appendix A: Definitions.

Technology trends across the broad field of AI have been analysed using patent data in recent publications by the World Intellectual Property Office (WIPO)[5] and the United Kingdom Intellectual Property Office (UKIPO)[6]. These reports show that AI-related inventions are booming and shifting from theoretical to commercial application, with the field of machine learning dominating this growth.

The authors and purpose of this report

This report provides a specific analysis of emerging technologies in machine learning relating to AI.

IP Australia is dedicated to building prosperity for Australia and ensuring that Australians benefit from great ideas. Using patent data analytics to provide evidence of innovation trends, we leverage our unique access to IP data, knowledge and expertise to deliver value to the broader community.

This report was prepared in partnership with the Australian Computer Society (ACS) as part of IP Australia's commitment to support Australian research and innovation. ACS is the professional association and peak body representing Australia's ICT sector. It aims to deliver authoritative independent knowledge and technology insights, building relevant technology capacity and capability that catalyses Australian innovation, and speeds the adoption of technology for the benefit of commerce, government and society.

ACS is at the forefront of education and policy development in Australia's ICT sector, contributing to building capability in the development of, and response to, emerging technologies. The 2019 *ACS Australia's Digital Pulse* report examines digital workforce trends and takes stock of Australia's digital economy.[7]

In this patent analytics report we have analysed the development and uses of technology relating to machine learning. This report also provides insight into innovators in this field – based on patents filed – with potential connections for research, development and commercialisation.

[5] World Intellectual Property Organization (2019), WIPO Technology Trends 2019, Artificial Intelligence. WIPO, Geneva https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf

[6] Intellectual Property Office (2019) Artificial Intelligence: A worldwide overview of AI patents and patenting by the UK AI sector. IPO, Newport https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/808891/AI-a-worldwide-overview-of-AI-patents.pdf

[7] Australian Computer Society (2019) ACS Australia's Digital Pulse 2019, https://www.acs.org.au/insightsandpublications/reports-publications/digital-pulse-2019.html

03 Patenting trend overview

As a basis for this analysis, we searched the PATSTAT database (2018 Autumn edition) for inventions relating to machine learning. The search[8] found 36,740 unique DOCDB patent families (see Appendix A: Definitions) relating to machine learning filed worldwide since 2012. This year range allows a focus on the most recent patent filings and includes complete patent data over five years.

Patent filings

Analysing patent family filings across time indicates growth or decline in innovation. Additional in-depth analysis can help understand underlying factors and their correlation to these trends.

Figure 1 shows the number of patent families by their earliest priority year and patent status. More details on patent status are given in Appendix A: Patent status determination.

Global patenting activity relating to machine learning has increased dramatically from 2012–16, with more than a four-fold increase. The dip in patent filings from 2017–18 reflects incomplete data rather than a declining trend.

Of the 36,740 patent families filed since 2012, 33,505 (91 per cent) are in an active state with patent protection being sought or in force. The low levels of patent families that have lapsed, expired or withdrawn may be reflective of both recent patenting activity in this technology area, and a high commercial interest in AI and machine learning.

FIGURE 1 PATENT FAMILIES RELATING TO MACHINE LEARNING BY EARLIEST PRIORITY YEAR AND PATENT STATUS
(Note: data for 2017–18 is incomplete due to a lag in publication of patents)

Source: PATSTAT 2018 Autumn Edition

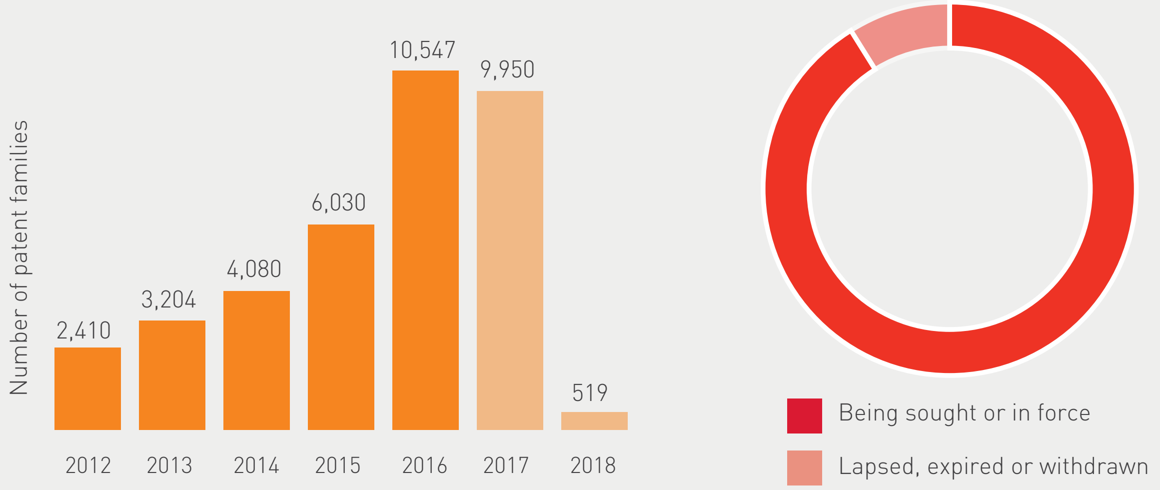

[8] The search used a combination of keywords, International Patent Classification (IPC) symbols and Cooperative Patent Classification (CPC) symbols. Detail of the search methodology is provided in Appendix B: Search strategy

Patent destination

Applicants must file patent applications in each country or patent jurisdiction where they wish to have enforceable patent protection. Reasons for patent protection include the country being a target for commercialisation, further research, or manufacturing.

Figure 2 shows the number of patent families filed in different jurisdictions. Patent Cooperation Treaty (PCT) applications have been excluded from this figure because these do not represent an enforceable right in any jurisdiction.

China is the largest patent destination, with 73 per cent of total machine learning related patent families filed here.

The United States is second with 22 per cent of the total, representing less than a third of patent filings in China. In third, fourth and fifth place are South Korea, Japan and Europe with only 4–5 per cent of total patent families filed in each.

Australia ranks tenth in patent destinations for machine learning related patents, with 295 patent families filed in Australia.

FIGURE 2: PATENT FAMILIES RELATING TO MACHINE LEARNING BY PATENT DESTINATION
(Note: only countries with more than 100 patent family filings are listed)

Source: PATSTAT 2018 Autumn Edition

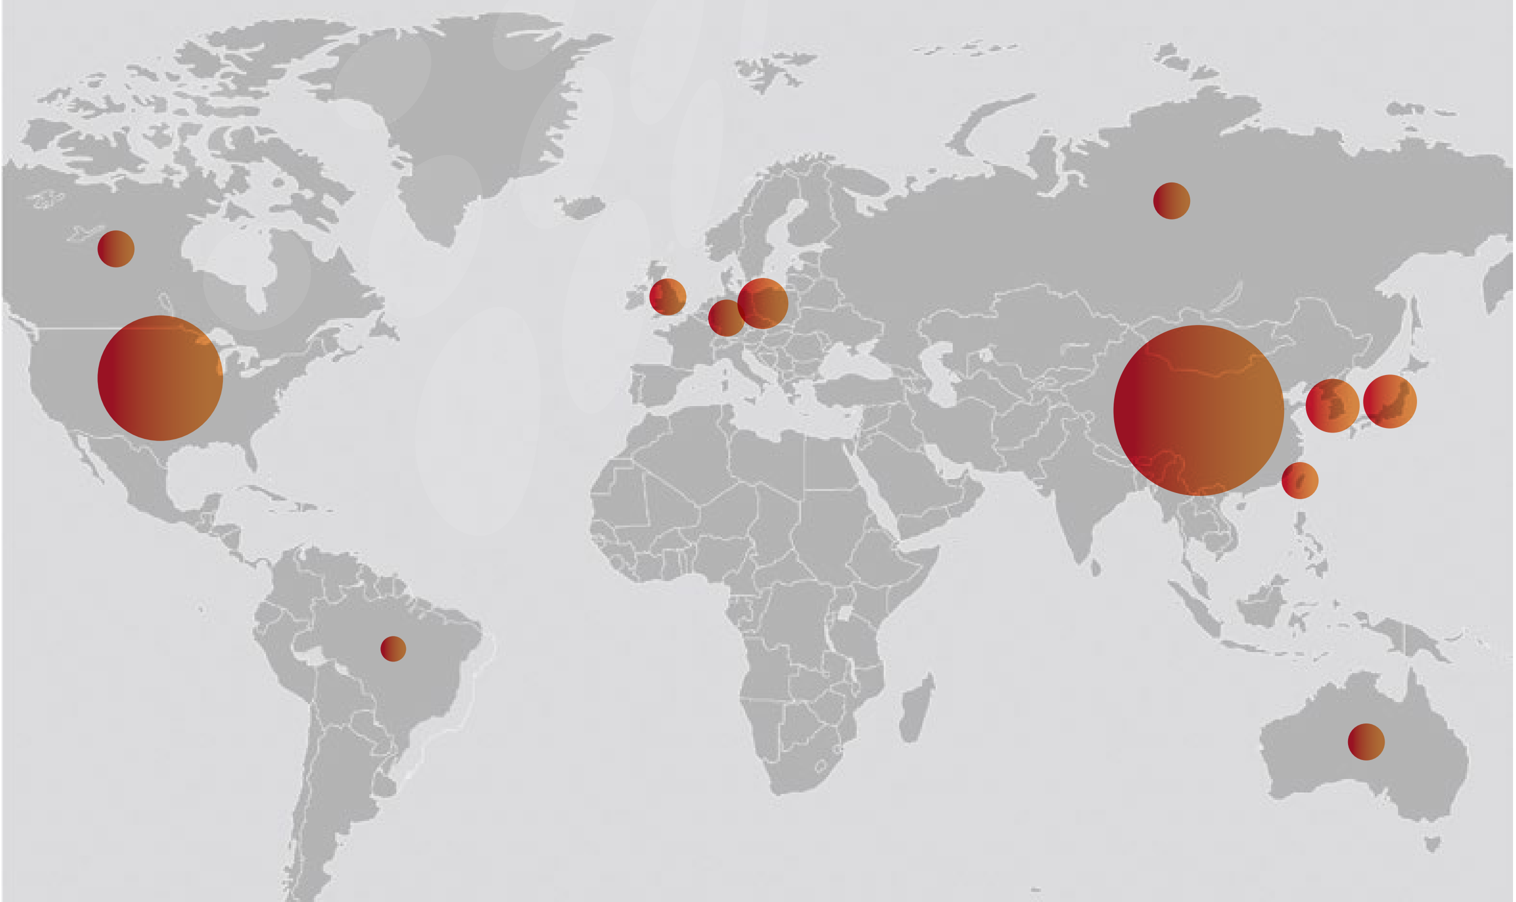

CHINA	26,758	UNITED KINGDOM	429
UNITED STATES	8,064	TAIWAN	384
SOUTH KOREA	1,823	CANADA	348
JAPAN	1,693	AUSTRALIA	295
EUROPE	1,611	RUSSIA	255
GERMANY	472	BRAZIL	107

Applicant origin

Analysis of patent applicant origin provides insights into countries with investment or interest in a specific area of innovation. Figure 3 shows the top 10 countries where patent families relating to machine learning originate. The outer bars represent the total number of families from a given country of origin and the inner bars represent the number of families from that location that have at least one member filed with an authority other than the home application authority (for example a Chinese applicant filing in the United States or filing a PCT application with WIPO).

Country information was derived from applicant address and applicant name harmonisation, where applicant address information was available. If a patent had applicants with addresses from more than one country, representing international collaboration, the patent was counted separately for each country of applicant origin. Where applicant country code information was not available (i.e. null data that would result in the patent application not being assigned to any country), the country of the application authority of the first application filed in the family was used as a proxy.

Full details of this approach are given in Appendix A: Definitions.

In the case of applicants who are resident in China, coverage of country information in PATSTAT is relatively poor. These applications confirmed the validity of the approach we used to derived applicant country information, because most of these patent families were also solely filed in China.

This method of applicant country assignation may account for the relatively larger number of Chinese applications found in the present dataset than in the recent WIPO[9] and UKIPO[10] patent analytics reports on AI.

FIGURE 3 PATENT FAMILIES RELATING TO MACHINE LEARNING BY APPLICANT ORIGIN

Source: PATSTAT 2018 Autumn Edition

CHINA	25,319	907
UNITED STATES	6,590	3,127
SOUTH KOREA	1,630	573
JAPAN	1,503	997
GERMANY	427	372
UNITED KINGDOM	338	248
TAIWAN	327	218
CANADA	239	221
SWITZERLAND	224	224
INDIA	144	130

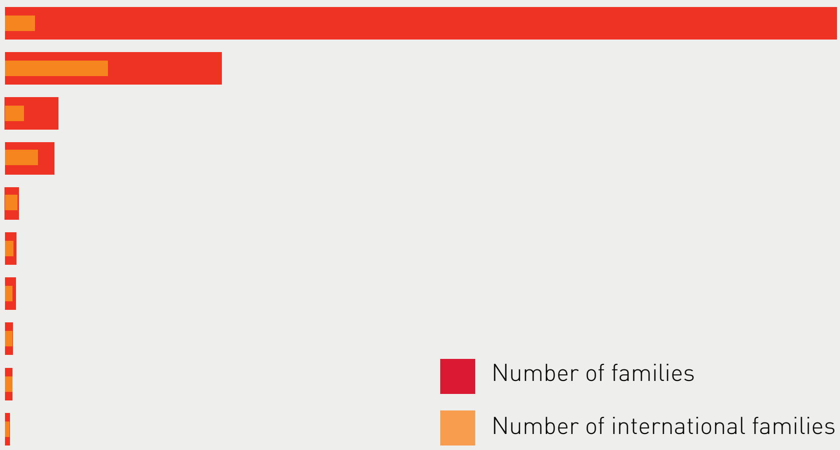

[9] WIPO (2019) WIPO Technology Trends 2019, Artificial Intelligence. WIPO, Geneva, https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf

[10] IPO (2019) Artificial Intelligence: A worldwide overview of AI patents and patenting by the UK AI sector. IPO, Newport https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/808891/AI-a-worldwide-overview-of-AI-patents.pdf

China is the largest source of patent filings, mirroring the result for patent destinations, with 69 per cent (25,319) of patent family filings having a Chinese applicant. Only 4 per cent of these families from Chinese applicants have international family members.

The number of Chinese-originating patents in machine learning may reflect Chinese government research and development initiatives. These have been particularly strong in the field of AI, with a government venture capital investment of US $30 billion and state and public-private investments totalling tens of billions more. [11]

As well as this, the State Council of China released the 'New Generation Artificial Intelligence Development Plan' in July 2017, which aims to build a domestic AI industry worth US $150 billion and to become the leading AI power by 2030.[12]

The United States comes in second with 18 per cent (6,590) of patent families originating from at least one American applicant. Forty-seven per cent of patent families from United States applicants have international family members.

The United States has a range of government measures supporting AI commercialisation. In 2015, the United States government invested approximately US $1.1 billion in unclassified research and development in AI-related technologies. In 2016, the National Science and Technology Council of the United States chartered a subcommittee on Machine Learning and Artificial Intelligence to help coordinate federal activities in AI.[13] The latest American AI Initiative,[14] issued by the United States President this year, is expected to drive further growth in AI and machine learning related inventions.

South Korea and Japan are in third and fourth place respectively, each with 4 per cent of global patent families. These originate from at least one applicant of either country (1,630 and 1,503 patent filings, respectively). Thirty-five per cent of patents originating in South Korea and 66 per cent of patents originating in Japan have international family members filed.

South Korea is one of the biggest government investors in AI per capita. The South Korean Ministry of Science, ICT and Future Planning defined an Artificial Intelligence Information Industry Development Strategy in 2016. The strategy considers the role of AI alongside other emerging technologies including the Internet of Things, cloud computing, big data analytics, and mobile technologies.[15] In 2016, South Korea also released a plan to invest US $2 billion by 2022 to strengthen its AI research and development capability and reach the 'global top four'.[16] The program included establishing six new AI research institutes to ramp up AI innovation.

Japanese government initiatives in AI research include the fifth Science and Technology Basic Plan 2016–20. This sets out a goal of Japan leading in this sector and transforming society with new information technologies and systems.[17] Japan also aims to train up to 250,000 AI experts a year to drive AI innovation.[18]

Germany ranks fifth globally for patent applicant origin with one per cent (427 patent applications) of global patents originating here, of which 87 per cent have international family members.

Germany has one of the largest AI research centres in the world, the German Research Center for Artificial Intelligence (DFKI), with an annual budget of around EUR 45 million.[19]

[11] Davenport, T. (27 February 2019) China is catching up to the US on artificial intelligence research, The Conversation https://theconversation.com/china-is-catching-up-to-the-us-on-artificial-intelligence-research-112119

[12] Future of Life Institute, AI Policy – China (accessed 9 October 2019) https://futureoflife.org/ai-policy-china/

[13] National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee (2016) The national artificial intelligence research and development strategic plan, https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf

[14] Future of Life Institute, AI Policy – United States (accessed 9 October 2019) https://futureoflife.org/ai-policy-united-states/?cn-reloaded=1

[15] Future of Life Institute, AI Policy – South Korea (accessed 9 October 2019) https://futureoflife.org/ai-policy-south-korea/

[16] Peng, T (17 May 2018) South Korea Aims High on AI, Pumps $2 Billion Into R&D, Medium, Synced Review, https://medium.com/syncedreview/south-korea-aims-high-on-ai-pumps-2-billion-into-r-d-de8e5c0c8ac5

[17] Harris, P (June 2017) Research brief: developments in artificial intelligence (AI) in Japan and implications for Australia, Australian Government, Department of Education and Training, https://internationaleducation.gov.au/International-network/japan/PolicyUpdates-Japan/Documents/AI%20in%20Japan%20research%20brief_07-2017.pdf

[18] Yamashita, M (30 March 2019) Japan aims to produce 250,000 AI experts a year, Nikkei Asian Review, https://asia.nikkei.com/Economy/Japan-aims-to-produce-250-000-AI-experts-a-year

[19] German Research Center for Artificial Intelligence (accessed 9 October 2019) https://www.dfki.de/en/web/

In 2018, the German government adopted a national AI strategy and invested EUR 3 billion toward AI research and development, with the aim of developing world leading AI capability.[20]

Germany may also benefit from the European Union's investment of EUR 1.5 billion in AI research.[21]

Australia ranks seventeenth globally, with 59 patent families having at least one Australian applicant, and 68 per cent of these having an international family member. More details on Australian applications are given in the final section of this reports, on Australian applicants.

A recent survey by Deloitte, State of AI in the Enterprise, shows that, although most Australian organisations are using AI service technologies, Australia has a perceived skills and capability gap in the AI sector, with Australians aiming to keep up with technological development, rather than take the lead.[22] The Australian Government has made a funding commitment of AU $29.9 million over four years announced in the 2018–19 Federal Budget to promote AI research and commercialisation. This is committed to initiatives including the Cooperative Research Centres Program, scholarships, education programs and developing a national AI framework.[23]

Patenting incentives in China

China is a world leader in filing patent applications across all technologies. This is partly driven by government policies to increase patent filings and enable research commercialisation. Initiatives include 'Measures for the Administration of Special Funds for Subsidising Foreign Patent Applications', issued by the Chinese Ministry of Finance in 2012. This helps domestic applicants enter the international market by funding PCT and Paris Convention patent applications.[24] This has pushed China into the second position globally in PCT filings in recent years, but does not appear to have stimulated large numbers of patents to be filed in other jurisdictions for patent protection.[25]

The Chinese patent office receives by far the largest number of patent filings of any office in the world, but - as is also reflected in the data presented here - more than 95 per cent of all Chinese-filed patents are seeking protection only within China. This means the their impact outside Chinese borders is relatively limited compared with that of jurisdictions such as the United States, which remained substantially the largest international filer to 2017.[26] This pattern of Chinese patent filing by Chinese applicants could also be driven by further Chinese government patent initiatives, ranging from subsidies for local filings offered by provincial governments, including reimbursement of fees and postponing of fees in case of financial difficulties, to tax breaks for companies with IP holdings, cash bonuses and funds for use in commercialising patented inventions,[27] and even reduced prison sentences for criminals who have granted patents.[28]

[20] Made in Germany, The national AI Strategy of the Federal Government (accessed 9 October 2019) https://www.ki-strategie-deutschland.de/home.html

[21] Dutton, T (June 29 2018) An Overview of National AI Strategies, Medium, https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd

[22] Jackson, P and Marshall, A (May 2019) Future in the balance? How countries are pursuing an AI advantage, Deloitte, https://www2.deloitte.com/au/en/pages/technology/articles/how-countries-are-pursuing-ai-advantage.html

[23] Budget Review 2018–19, Parliament of Australia, https://parlinfo.aph.gov.au/parlInfo/search/display/display.w3p;query=Id%3A%22library%2Fprspub%2F5982057%22

[24] Notice on Printing and Distributing the Measures for the Administration of Funding for Foreign Patent Special Funds http://www.gov.cn/zwgk/2012-05/31/content_2149501.htm

[25] China Drives International Patent Applications to Record Heights; Demand Rising for Trademark and Industrial Design Protection https://www.wipo.int/pressroom/en/articles/2018/article_0002.html

[26] WIPO Stats on Patent Application Filings Shows China Continuing to Lead the World https://www.ipwatchdog.com/2017/12/12/wipo-stats-patent-application-filings-china-lead/id=90855/

[27] How patent subsidies boost R&D https://www.managingip.com/Article/3538144/How-patent-subsidies-boost-R-D.html?ArticleId=3538144

[28] How get out jail early in China: buy an inventor's idea and patent it https://www.scmp.com/news/china/article/1681850/how-get-out-jail-early-china-buy-inventors-idea-and-patent-it

Technology overview

Analysing the technical content of patents can provide insights into specific technologies. These insights highlight areas of commercial interest and can be used to inform research and commercialisation strategies.

In machine learning, we have considered these developments in terms of four intertwined functional layers. While each of these layers plays an important role, they are not mutually exclusive, which is demonstrated in the examples of this report.

	FUNCTIONAL LAYER	
Applications	Core capabilities	Broad functions that use foundational algorithms to provide general outcomes
	Real world applications	Uses of core capability functions to provide solutions for real world problems
Foundational Technologies	Algorithms	Basic algorithms and techniques
	Architecture	Hardware and systems that allow new technologies to work

We start by broadly categorising the field into either applications or foundational technologies to give an idea of investment and the degree of focus on applications of these technologies. We then focus more specifically on examining real world applications and core capabilities (used to facilitate real world applications) and reviewing the state of play in foundational technologies (architecture and foundational algorithms).

The patent families identified in this analysis relate to machine learning applications, machine learning foundational technologies, or both. Figure 4 shows the focus of the 36,740 patent families, with 60 per cent referring to both a foundational technology and its application (see Appendix C: Technology Analysis).

Altogether, 66 per cent (24,091) of these patent families relate to foundational technologies. These include architecture and algorithms, which we will explore later in this report.

A total of 94 per cent (34,376) of these patent families discuss machine learning applications. These applications include the core capabilities of image and video analysis or speech and text analysis as well as the application of these technologies to industry. These core capabilities and their real world applications are discussed in detail in the following sections.

FIGURE 4 TECHNOLOGY OVERVIEW OF PATENT FAMILIES RELATING TO MACHINE LEARNING

Source: PATSTAT 2018 Autumn Edition

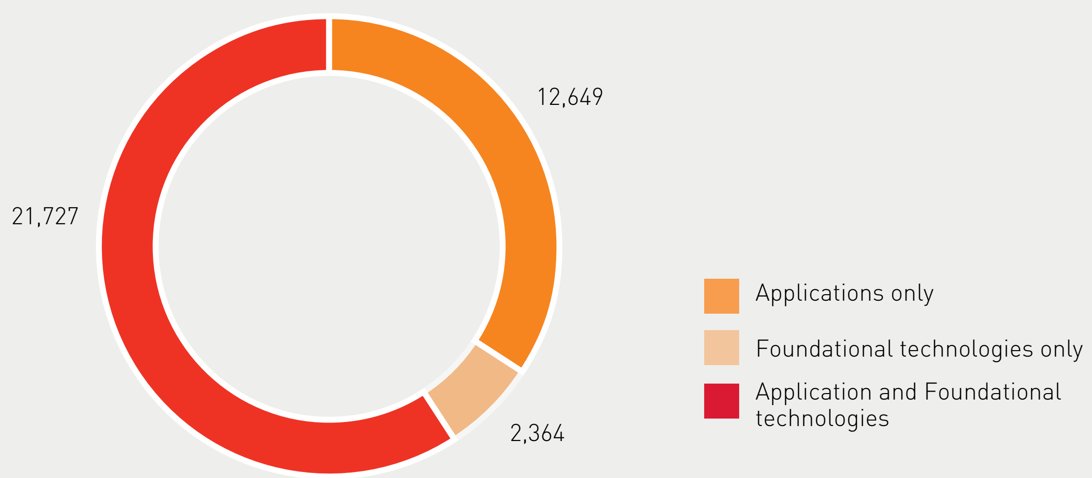

04

CORE CAPABILITIES

Core capabilities

There are 34,376 patent families that describe machine learning applications where machine learning is applied to solve problems. These patent applications can describe both the core capability (27,956 patent families) being applied e.g. image and video analysis, and/or its real world application (25,768 patent families) e.g. agriculture or healthcare.

Here we initially analyse core capabilities, followed by real world applications.

We have divided machine learning patent applications by their core capability according to patent technology classification, in the following hierarchal order. This hierarchy allows a comparison between broad sectors.

HIERARCHY	CORE CAPABILITY	EXAMPLES OF APPLICATIONS
1	Speech and text analysis	• Speech recognition • Natural language processing • Audio processing and pronunciation
2	Image and video analysis	• Computer vision • Object recognition and tracking • Image matching and captioning
3	Control and automation processes	• Systems control • Automation • Process optimisation
4	Signal analysis	• Optical, electric, and magnetic signals • Signal detection • Use in diagnosis
5	Data analytics and cognitive processing	• Data mining • Information retrieval • Forecasting

Figure 5 shows the number of patent families in each category by their grant status; granted (where at least one family member is granted in any jurisdiction) or patent application (a patent family with members seeking patent protection, or in which all members are lapsed/withdrawn without grant). The

FIGURE 5 PATENT FAMILIES RELATING TO MACHINE LEARNING CORE CAPABILITIES BY THEIR GRANT STATUS

Source: PATSTAT 2018 Autumn Edition

IMAGE/VIDEO ANALYSIS	2,970	10,352
ANALYTICS AND PROCESSING	1,290	4,301
SPEECH/TEXT ANALYSIS	1,125	2,756
CONTROL AND AUTOMATION	1,033	1,576
SIGNAL ANALYSIS	913	1,640

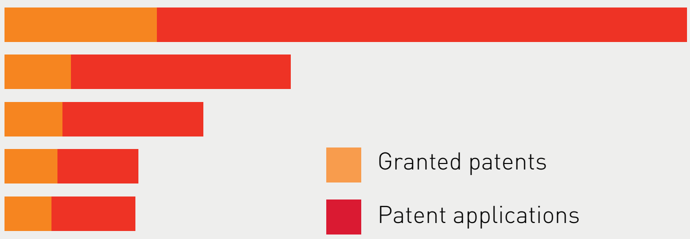

image and video analysis sub-category captures patent applications relating to computer vision, image processing, object recognition, object tracking, image matching and captioning. This sub-category accounts for 48 per cent of patent families describing machine learning core capabilities, totalling 13,322 patent families of which 22 per cent have a granted family member.

The analytics and processing sub-category covers patent applications relating to data mining, processing, conversion and filtering, and information retrieval and forecasting. This is the second largest category with a total of 5,591 patent families, of which 23 per cent have been granted in at least one jurisdiction.

Speech and text analysis includes natural language processing (NLP) and understanding, semantics, and pronunciation, and speech and audio processing. There are 3,881 patent families in this sub-category, of which 29 per cent have a granted family member.

The control and automation sub-category include systems control, automation and optimisation. There are 2,609 patent families in this category, of which 40 per cent have a granted family member.

Signal analysis includes optical (light), electric, electro-chemical and magnetic signals, their detection and their use in diagnosis. There are 2,553 patent families, of which 36 per cent have a granted family member.

Next we will take a closer look at each of these categories and their top applicants.

Top applicants

Chinese applicants have the most filings in each core capability, with 77 per cent of patent applications for image and video analysis, 72 per cent for control and automation, 71 per cent for signal analysis, 66 per cent for analytics and processing and 49 per cent for speech and text analysis originating from applicants in China. The United States is second for speech and text analysis with 35 per cent of patent families originating there. This is a focus of the United States, which has 23 per cent of total patent families in analytics and processing and less than 20 per cent of patent families in any other application field.

Figure 6 shows the top 10 applicants who filed patent applications relating to machine learning applications by core capability.

FIGURE 6 PATENT FAMILIES RELATING TO MACHINE LEARNING CORE CAPABILITIES BY APPLICANT

Source: PATSTAT 2018 Autumn Edition

	ANALYTICS & PROCESSING	CONTROL & AUTOMATION	IMAGE/VIDEO ANALYSIS	SIGNAL ANALYSIS	SPEECH/ TEXT ANALYSIS
STATE GRID	390	58	181	144	17
CHINESE ACADEMY OF SCIENCES	90	38	467	57	88
IBM	143	13	70	15	317
XIDIAN UNIVERSITY	19	4	414	17	8
ZHEJIANG UNIVERSITY	68	77	143	48	22
MICROSOFT	113	1	72	7	127
UESTC	27	16	220	30	25
TIANJIN UNIVERSITY	31	12	190	43	29
ALPHABET	62	8	86	7	138
TSINGHUA UNIVERSITY	63	24	119	28	45

The State Grid Corporation of China (State Grid), which is also the leading applicant in machine learning techniques, ranks top in patent filings relating to data analytics and processing, signal analysis and field only applications. State Grid also has a rich patent portfolio in the control and automation, and image and video analysis capabilities.

The Chinese Academy of Sciences rank top in patent filings for image and video analysis, and ranks in the top six for all other application sub-categories as shown in Figure 6. Notably, the Chinese Academy of Sciences also ranks second in machine learning foundational technologies, discussed later in this report.

IBM, Alphabet and Microsoft take the lead – first, second and third, respectively – in patent applications relating to speech and text analysis. Each company also has a strong patent portfolio in analytics and processing and image and video analysis.

Xidian University ranks second in the image and video analysis sub-category. This is clearly a focus of the university, comprising 83 per cent of their patent portfolio relating to machine learning applications. They are also the third top applicant in the category of machine learning foundational technologies.

Image and video analysis

Image and video analysis is the most rapidly growing core capability in this dataset, with a seven-fold increase from 2012–17. This is largely due to recent developments in machine learning that vastly improve the ability to process visual information. Video and image analysis is employed across a broad range of sectors including telecommunications and networks, medicine and healthcare, predictive maintenance, electrical power and energy, and transport and logistics.

Examples of representative patent applications relating to image and video analysis are held by NEC Corporation, Adobe Systems and Canon.

NEC, a Japanese multinational technology company, operates with a vision of continuously seeking technologies to create social value. NEC focuses research and development on developing recognition AI, analytics using AI, AI platforms, communications, cybersecurity, and data distribution platforms.[29] NEC has filed 67 patent families relating to image and video analysis. Application WO2017079522 describes a method for detecting objects in an image by using 'subcategory-aware' convolutional neural networks (CNNs). This is a two-step process consisting of: firstly, making a list of areas in the image that may contain an object of interest using a region proposal network (RPN) trained to sub-categories such as cars with frontal view or people standing; and secondly classifying and refining the list of possible objects with an object detection network (ODN). The ODN simultaneously assigns both category and sub-category classification and bounding boxes to objects. Applications of this technology focus on detection of cars, pedestrians and cyclists in outdoor scenes. This patent family has a member granted in the United States.

Adobe Systems is an American multinational computer software company focusing on multimedia and creativity software and digital marketing software.[30] Adobe Systems has 63 patent families filed relating to image and video analysis. The application US2017200065 describes techniques for automatically generating descriptions of images. The method uses weak supervision data (noisy data that is not closely curated and may include errors, such as tags from social networks) to provide information to supplement concepts extracted automatically from the image. A combined machine learning approach is taken, where a pre-trained CNN is used to assign descriptive terms and to produce a vector of visual features, which is fed into a recurrent neural network (RNN) to generate labels using language modelling and sentence generation. This patent family has a member granted in the United States and protection in being sought in Australia and China.

Canon, a Japanese multinational company specialising in imaging and optical products, host a healthcare optics research laboratory and an imaging systems research laboratory.[31] Canon has filed 62 patent families relating to image and video analysis since 2012. Canon's patent portfolio includes US2016155024, which describes classifying elements in a video as foreground or background, also known as scene modelling. The technique operates in two phases: first classifying all elements to be either foreground or background; and second zooming in on a foreground object and re-classifying elements in the frame. This second classification is more sensitive to foreground. The method is both computationally inexpensive and able to handle similarities between elements (a so-called camouflage scenario), and may be useful for activity detection, unusual object or behaviour detection, and scene analysis. This patent family has a member granted in the United States, with protection being sought in Australia.

[29] NEC Technology Vision (accessed 9 October 2019) https://www.nec.com/en/global/rd/technologyvision/index.html

[30] Adobe (accessed 9 October 2019) https://www.adobe.com/

[31] Canon Research and Development (accessed 9 October 2019) https://www.usa.canon.com/internet/portal/us/home/about/research-development/

CASE STUDY: GENERATIVE ADVERSARIAL NETWORKS (GANS)

Machine learning has generated such remarkable developments over the past five years that it may be difficult to grasp the breadth of their potential impact. The development of GANs is an example of the huge advances in this field.

GANs are a class of deep neural networks introduced by Ian Goodfellow and other researchers at the University of Montreal in 2014.[32] GANs have been identified as 'the most interesting idea of the decade' by Yann LeCun, who oversees AI research at Facebook.[33]

GANs have a wide range of applications, mostly related to images and video capabilities. They allow improvements in image resolution, reconstruction of scenes, three dimensional models, synthetic music and text, and even applications in physics and biology.[34,35,36] One of the most well-known uses of GANs in the public eye is generation of deepfakes. Deepfakes are synthesised images or videos synthesised using machine learning techniques called deep neural networks. This morphs one person's face to mimic someone else's features while preserving original facial expressions, without the need for any video editing skills. This technique allows the creation of photorealistic videos at a very low cost. These techniques have also been used to replicate audio in a similar manner. GANs are based on the idea of adversarial training,[37] using two neural networks; the generator and the discriminator. The generator captures the data distribution and creates samples that are almost identical to the training data set. The discriminator tries to differentiate between original samples and those created by the generator. During training, the generator tries to improve its synthetic samples and the discriminator tries to catch the fakes, and so the process is called adversarial training. Success occurs when samples from the generator are convincing to the discriminator.[38]

Patent applications relating to GANs were first filed in 2016, with only four patent families that year. Patent filings surged to 77 in 2017, with wider application of GANs in the areas of image and video processing, and speech and text analysis.

Two innovation patent applications filed in Australia relate to GANs. AU2017101166 relates to synthesising images and describes a method for transferring style between images. This method uses a conditional GAN (one that incorporates classification to exert some control on the output) to extract features of the input image and a 'U-Net' (a form of convolutional network, to decode the features). The conditional GAN allows for more explicit extraction of features from the input image. AU2018100325 discloses a method for colouring black-and-white films and photos in real time, also by using conditional GANs.

Patent application WO2018053340, filed by Twitter,[39] describes a method of enhancing low resolution images to higher resolution equivalents (known as super-resolution reconstruction). This method uses a GAN trained on corresponding low- and high-resolution images, followed by sub-pixel convolution to create high-resolution output. The ready accessibility of GANs means that both positive and negative social outcomes may result from their use in creating realistic synthetic content. To mitigate the resultant risks associated with GANs, researchers are also working to find markers of the use of GANs in images.[40] The power of this type of technology is vast, and as with any emerging technology, regulation lags behind early adoption. It remains to be seen what impacts GANs will have in commerce, science and society.

[32] Goodfellow, I.J. et.al., "Generative Adversarial Nets", NIPS, 2014 https://arxiv.org/pdf/1406.2661.pdf

[33] Yann LeCun Quora Session Overview (August 2016) https://www.kdnuggets.com/2016/08/yann-lecun-quora-session.html

[34] Hui, J (23 June 2018) GAN — Some cool applications of GANs, Medium, https://medium.com/@jonathan_hui/gan-some-cool-applications-of-gans-4c9ecca35900

[35] Mustafa et al., CosmoGAN: creating high-fidelity weak lensing convergence maps using Generative Adversarial Networks, Computational Astrophysics and Cosmology 6 (1) 2019 https://comp-astrophys-cosmol.springeropen.com/articles/10.1186/s40668-019-0029-9

[36] Zhavoronkov et al., Deep learning enables rapid identification of potent DDR1 kinase inhibitors, Nature Biotechnology, 37, 1038-1040, 2019 https://www.nature.com/articles/s41587-019-0224-x

[37] Adversarial Machine Learning Tutorial (accessed 9 October 2019) https://aaai18adversarial.github.io/

[38] Kang, N (May 2017) Deepfakes: The Good, The Bad and the Ugly, Medium, https://medium.com/twentybn/deepfake-the-good-the-bad-and-the-ugly-8b261ecf0f52

[39] About Twitter Inc (accessed 9 October 2019) https://about.twitter.com/en_gb.html

[40] Murphy, H (16 August 2019) Cyber security companies race to combat 'deepfake' technology, Financial Times, https://www.ft.com/content/63cd4010-bfce-11e9-b350-db00d509634e

Analytics and processing

Developments in data analytics and business intelligence methods using machine learning are evolving predictive and prescriptive analytics to gain insights from big data sources. Patent filings in this area have increased almost five-fold from 2012–16. Patent applications in this core capability predominately relate to electrical power and energy, business and finance, telecommunications and networks, retail and medicine and healthcare sectors.

Examples of patent applications in analytics and processing are held by Microsoft, Facebook and Yahoo.

Microsoft has 113 patent families in this core capability. WO2016164694 describes a system that presents the best information to a content feed. The system uses predictive models (such as neural networks or regression models) to predict the trending likelihood of digital entities (documents, people, locations etc) based on historical behaviour and context. In combination with information about the user and their network, this allows presentation of the most relevant and trending digital entities to that user. This patent family is granted in the United States and is seeking protection in China and Europe.

Facebook is the world's largest social networking company with a focus on connecting individuals with content and stories they care about most.[41] Facebook's researchers develop machine learning algorithms that rank feeds, ads and search results, and create new text understanding algorithms that keep spam and misleading content at bay.[42] Facebook has filed 54 patent families relating to analytics and processing. WO2015038297 discloses a method for understanding a user's sentiment toward a page or topic based on their interactions with it. Part of the method is the use of traditional sentiment analysis techniques to identify individuals with a strong polarity toward particular content to generate a list of 'trusted users'. The information from this set of users is then used to train a predictive machine learning module that can be used for predicting the sentiment of new users interacting with the content. This patent family is seeking protection in the United States, Canada, Japan, Korea and Australia.

Yahoo is an American web services provider. Yahoo researchers develop large scale machine learning algorithms relating to content ranking, powering ad selection, spam detection and abuse prevention.[43] Yahoo's patent portfolio includes 30 patent families on analytics and processing. US2017308537 describes a method for providing a customised news stream to a user, and measuring the amount of time that news items are displayed on a user's device. These times, along with the features present in each news item, are used to train a classifier. This provides a feed ranked by expected viewing time by the user. This patent family has a member granted in the United States.

Speech and text analysis

Development of machine learning techniques has enabled speech and character recognition to have human-like fidelity and the ability to dynamically adjust to accents, handwriting styles and variable layouts. Speech and text analysis core capabilities include NLP and understanding, which has rapidly progressed in the past five years. AI tools can now beat humans in comprehension tests.[44]

Our analysis shows a five-fold increase in speech and text analysis related patent filings between 2012–16, from 270 patent families to 1210. These patent applications are applied across broad sectors including telecommunications and networks, medicine and healthcare, business and finance, and retail.

Examples of patent applications in this sub-category are held by IBM, Alphabet, Apple and NTT (Nippon Telegraph and Telephone Corporation).

[41] Statista, Most famous social network sites 2019, by active users (accessed 9 October 2019) https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/

[42] Applying machine learning science to Facebook products (accessed 9 October 2019) https://research.fb.com/category/machine-learning/

[43] Yahoo Research Areas (accessed 9 October 2019) https://research.yahoo.com/research-areas

[44] Chong, Z (16 January 2018) AI beats humans in Stanford reading comprehension test, CNET, https://www.cnet.com/news/new-results-show-ai-is-as-good-as-reading-comprehension-as-we-are/

An IBM patent application US2017323011 describes methods for answering questions using a 'persona-based' NLP system. The user selects the persona for the machine to adopt, and the machine considers the personality attributes of that persona in addition to any factual information in formulating its response to queries. A member of this patent family has been granted in the United States.

One of Alphabet's patent applications US2017186420, originally filed by Google, describes a method for converting audio input into a sequence of phonemes (a set of unique sounds that cover a particular language). Acoustic features are extracted at each time step, and these features are processed using a long-short-term memory (LSTM) approach. LSTM neural networks are a very successful embodiment of an RNN, adapted to understand long-term dependencies in data, and are particularly useful for time-series data like speech processing. The output of the LSTM component of the system is used to generate the probability that the sound at any given time step is accurately represented by a particular phoneme, building a map of received sound into language. This method can be applied to speech recognition along with pronunciation and language modelling. This invention has been patented in the United States.

Apple are dedicated to developing machine learning technologies to incorporate intelligence into their devices and services, touching the lives of millions of users every day.[45] Apple has a total of 51 separate patent families filed in this sub-category. WO2016003509 relates to a virtual assistant for facilitating interactions with a television. The virtual assistant uses NLP and contextual information (such as items displayed) to understand a user's intention and to accurately play their desired media. This patent family includes members filed in the United States, Europe, Australia, China and Taiwan, with patents granted in the United States, Australia and Taiwan.

NTT is the fourth largest telecommunications company in the world, and is taking a leading role in digital transformation in the telecommunications sector. One example of a patent application held by NTT in speech and text analysis is WO2018062265 which describes the use of a predictive acoustic model to perform speech recognition with improved efficiency and less training data.

Control and automation

Patent filings in control and automation core capabilities tripled between 2012–16. This analysis found applications were concentrated around electrical, power and energy, industry and manufacturing, transport and logistics and predictive maintenance.

Examples of patent applications of this sub-category are held by State Grid and Ford Motor Company.

State Grid patent application CN105262113 sets out a method for controlling a photovoltaic power generation system using a probabilistic fuzzy neural network with an error backpropagation learning algorithm. The system is designed to rapidly adjust the in and out of phase power that the generator puts into the grid, in order to robustly handle changes in grid conditions such as faults, and to support grid voltage during failures. This patent family has a member granted in China.

Ford Motor Company's mission is to design smart vehicles for a smart world.[46] The company focuses on using AI and machine learning across different sectors from connected car solutions to autonomous vehicles. Their AI platform called CarStory analyses the inventory data of used cars to gain insights into what makes each previously owned car unique.[47] US2018099646 uses a neural network to classify levels of precipitation using a range of vehicle sensors such as vibration of the windshield and imaging (Figure 7). This classification is used to control vehicle functions such as windshield-wiper speed and traction-control. This patent family includes family members actively seeking protection in the United States, China and Germany.

[45] Apple machine Learning Journal (accessed 9 October 2019) https://machinelearning.apple.com/

[46] Ford-About us (accessed 9 October 2019) https://corporate.ford.com/company.html

[47] Marr, B (17 May 2019)The Amazing Ways The Ford Motor Company Uses Artificial Intelligence And Machine Learning, Forbes, https://www.forbes.com/sites/bernardmarr/2019/05/17/the-amazing-ways-the-ford-motor-company-uses-artificial-intelligence-and-machine-learning/#313f79dee49a

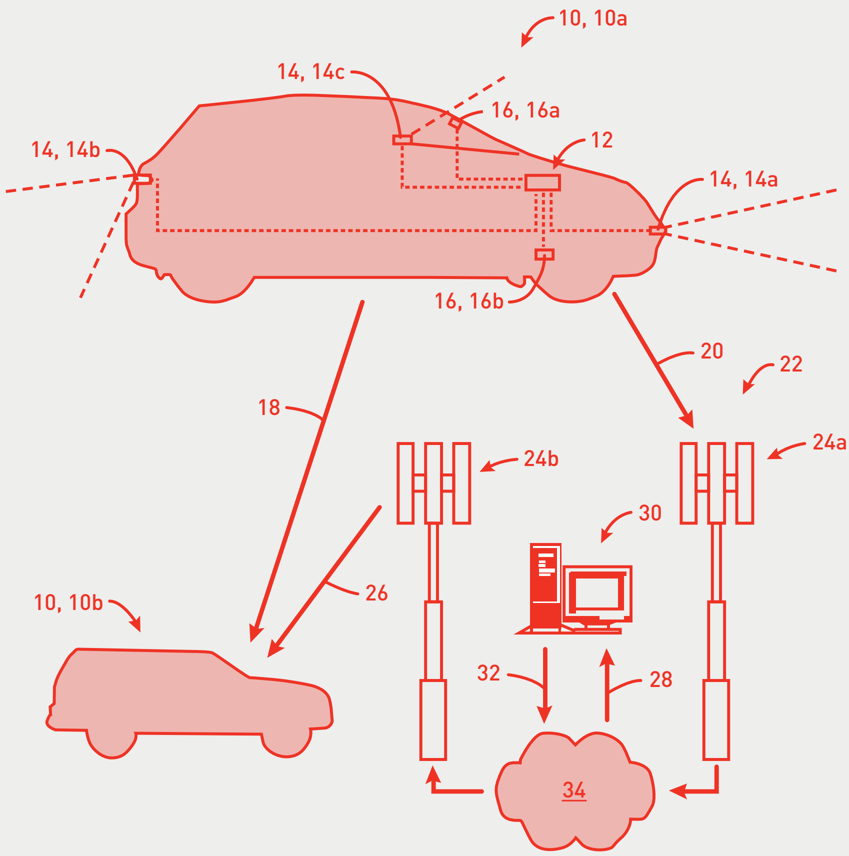

FIGURE 7 IMAGE FROM FORD MOTOR COMPANY'S PATENT APPLICATION US2018099646

This patent application shows the use of machine learning neural network in the core capability of control and automation. This links to the real world use of a neural network to control vehicle functions such as windshield-wiper speed and traction-control based on vehicle sensors such as windshield vibration and imaging.

Signal analysis

Patent families relating to signal analysis tripled from 2012–16, with 685 patent families filed in 2016. Signal analysis is widely used for detection and diagnosis in various sectors including medicine and healthcare, predictive maintenance, electrical power and energy, and telecommunication and networks.

Representative patent applications from this core capability are held by Intel and the California Institute of Technology.

Intel, the American computer company that provides computing, networking, data storage and communication solutions worldwide, uses Intel AI to handle their AI patenting. Intel combines multi-purpose, customisable and application-specific solutions to provide the best hardware for each job, along with software optimisation for tasks.[48] US2018206774 relates to a wearable device for sufferers of Parkinson's disease. This device is able to identify a specific motor symptom of the disease called Freezing of Gait (FoG) using a neural network, and uses heart rate detection and inertial sensors to identify FoG events. Patent protection is being sought in the United States.

The California Institute of Technology (Caltech) is a science and engineering institute with research interests ranging from quantum science and engineering to bioinformatics, human behaviour and economics, and energy and sustainability.[49] WO2018035257 describes an enhanced stethoscope with a range of sensors (such as stethoscopic and non-stethoscopic audio, ultrasound, temperature etc). Machine learning is used to correlate the outputs of these various sensors to provide clinically relevant information. This patent family includes a member granted in the United States and members actively seeking patent protection in Canada, China and South Korea.

Other applications

There are 6,420 patent applications, of which 35 per cent have a granted family member, that describe real world applications of machine learning, but which do not specify a core capability. While real world applications of machine learning will be discussed in detail in the next section, analysis of this group of patent families gives an idea of movements in development from foundational and core capabilities to real world applications.

[48] Intel AI (accessed 9 October 2019) https://www.intel.ai/
[49] Caltech at a Glance (accessed 9 October 2019) https://www.caltech.edu/about/at-a-glance

CASE STUDY: PRIVACY PROTECTION

Control of data privacy is increasingly important, as the global economy is increasingly driven by data. Tech corporations dominate market capitalisation lists,[50] and the major product of many of these companies is either data itself (including Alphabet and Facebook) or strongly dependent on the use of data, as well as the sale of data processing capability (including Microsoft, Amazon, Apple, Alibaba and Tencent). The devices we use, the cities we inhabit, the hospitals we visit and the shops we patronise are constantly collecting our data, both public and intimate. This data is used in machine learning applications to train advertising, prediction, language processing and diagnostic machines, amongst many others; but this data can also be highly sensitive.

Privacy considerations associated with data storage are highlighted by the use of machine learning techniques, which require a vast amount of data for training.[51] In addition, models trained on private data can be attacked in order to recreate the private data.

Machine learning itself can also be used to, deliberately or accidentally, create data links and inferences that would not be possible by humans. There is therefore an increasing interest in the development and use of privacy preserving machine learning. Privacy preservation goals in machine learning are approached in several ways, with leading fields including differential privacy, secure enclaves, federated learning and homomorphic encryption.

Differential privacy is a requirement that an algorithm is not overly dependent on the information from any individual in producing its output. A differentially private algorithm makes it unclear whether any individual's data was used as the input leading to a particular statistical output.[52] Relevant patent applications include CN107368752 from the Beijing Technology and Business University,[53] directed toward protecting data used for training deep learning models from attacks that use GANs to recreate training datasets. This is achieved by anonymising the training data, running simulated attacks, and using the degree of recreation of the input as feedback for the anonymisation procedure.

Machine learning models can be executed in secure enclaves, which provide trusted and isolated environments[54] WO2017222902 from Microsoft discloses secure execution environments where third parties can separately upload data and machine learning code, and the code can be run on the data in a secure environment without either party having access to both components. The output is then provided to the third parties.

To reduce the vulnerability of stored data, machine learning algorithms can also be run as federated learning methods in a distributed fashion, reducing the need for transferring data into a centralised data store giving data suppliers more control over how their data is used.[55] WO2018057302 by Alphabet relates to efficient communication techniques for performing distributed machine learning on unreliable networks, such as mobile networks with intermittent coverage.

Homomorphic encryption allows certain computations to be performed on the encrypted data itself, without requiring access to the raw data.[56] CN107241182 from the University of Electronic Science and Technology of China (UESTC) relates to a method of clustering vectors where the vectors are homomorphically encrypted before being passed to a machine for clustering calculations and then the processed clusters are passed back to the client for decryption.

[50] Desjardins, J (21 June 2019) The Front of the Pack, How the make-up of the world's largest companies has changed since 1999, Visual Capitalist, https://www.visualcapitalist.com/a-visual-history-of-the-largest-companies-by-market-cap-1999-today/

[51] Mancuso, J (11 January 2019) Privacy-Preserving Machine Learning 2018: A Year in Review, Medium, https://medium.com/dropoutlabs/privacy-preserving-machine-learning-2018-a-year-in-review-b6345a95ae0f

[52] Ketchell, M (19 March 2018) Explainer: what is differential privacy and how can it protect your data?, The Conversation, http://theconversation.com/explainer-what-is-differential-privacy-and-how-can-it-protect-your-data-90686

[53] Beijing Technology and Business University (accessed 9 October 2019) http://english.btbu.edu.cn/

[54] The Oasis team (1 September 2018) Towards an Open-Source Secure Enclave, Medium, Oasis Labs, https://medium.com/oasislabs/towards-an-open-source-secure-enclave-659ac27b871a

[55] Tramel, E 2019 Federated Learning: Rewards & Challenges of Distributed Private ML, https://www.infoq.com/presentations/federated-learning-distributed-ml/

[56] Homomorphic Encryption (accessed 9 October 2019) https://brilliant.org/wiki/homomorphic-encryption/

05

Real world applications

The use of machine learning techniques in real world applications demonstrates how comprehensively machine learning developments affect our lives. Analysis of these applications gives us a glimpse of what the future holds as this dramatic growth continues to develop.

Real world applications (25,768 patent families) of machine learning build on the development of core capabilities and foundational technologies. In this section, we will show how machine learning is affecting specific industries. As part of the layered technical development of machine learning, this section demonstrates the potential and complexity of this field.

Here we have divided machine learning patent applications by their real world application according to patent technology classification, in the following hierarchal order. This hierarchy allows a comparison between broad sectors, despite obvious cross-over between different applications.

HIERARCHY	REAL WORLD APPLICATION	EXAMPLES OF APPLICATIONS
1	Predictive maintenance	• Monitor equipment, structures or vehicles • Predict when maintenance is required • Faults detection
2	Telecommunication	• Managing mobile and computer networks • Internet and data security • Automated telemarketing and polling
3	Agriculture	• Field robots • Crop and soil monitoring • Predictive analytics
4	Energy	• Generating energy • Automated management of grids • Renewable energy
5	Environment and mining	• Mineral exploration • Automated sorting and optimised metallurgy • Earth observation
6	Security and military	• Biometric scanning • Fraud detection • Reading sentiment in written text
7	Retail	• Product recommendations • Product identification and personalisation • Demand prediction
8	Health care	• Diagnostics • Augmented decision making • Administration management
9	Industry and manufacturing	• Industrial automation • Intelligent supply chain optimisation • Product design
10	Transport and logistics	• Autonomous vehicles and driver assistance • Monitoring engine performance • Enhanced safety and efficacy
11	Business and finance	• Algorithmic trading • Robo-advisors • Deep data analytics

Figure 8 shows the number of patent families in each category of real world application by their grant status: granted (where at least one family member is granted in any jurisdiction) or patent application (a patent family with members seeking patent protection, or in which all members are lapsed/withdrawn without grant).

These areas showcase current industries where machine learning is being developed and is having an impact. These industries are – or will soon be – transforming through the use of AI and machine learning in all aspects of their systems and production. This gives a preview into how the global future of work will soon change.

FIGURE 8 PATENT FAMILIES RELATING TO MACHINE LEARNING REAL WORLD APPLICATIONS BY THEIR GRANT STATUS

Source: PATSTAT 2018 Autumn Edition

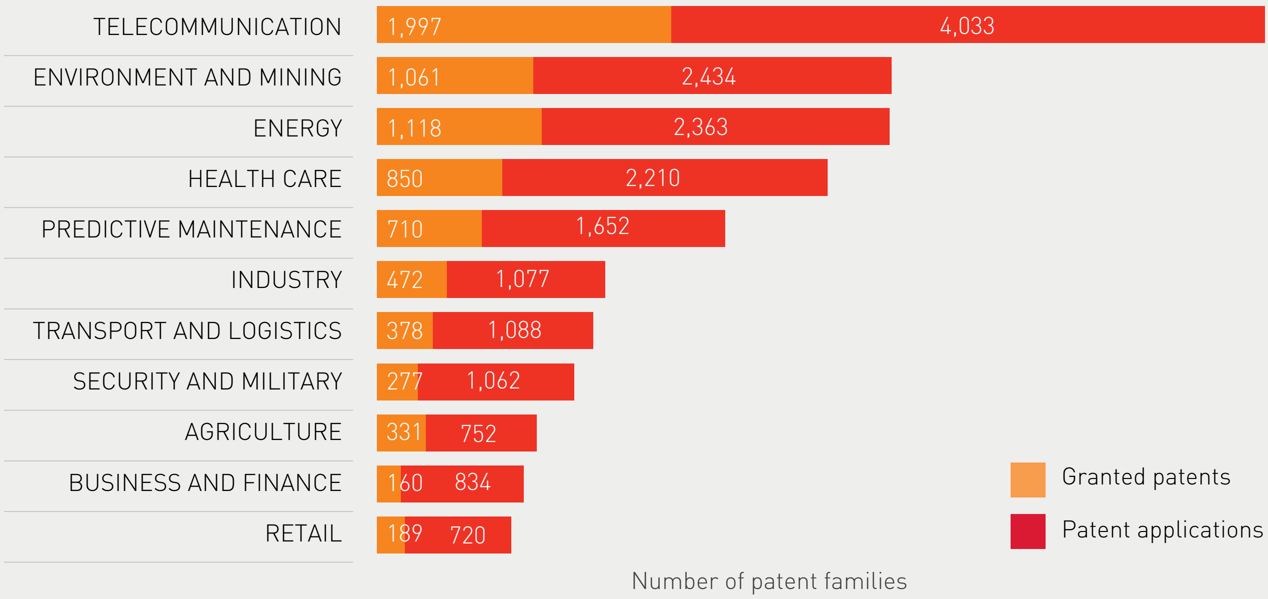

Telecommunication

The telecommunication sector is the leader in real world applications of machine learning. With over 6030 patent families, this sector is maturing fast and has nearly 2000 granted patents. This category includes patent applications for managing mobile and computer networks, internet and data security, as well as inventions that allow service providers to measure or predict user satisfaction with web-based services and content.

Examples of innovation in this sector are the use of metrics from home internet network operation, such as streaming rates and signal strength (WO2017220428 from Alcatel Lucent, acquired by Nokia in 2016[57]); and media-specific measurement using metrics such as length and frequency of video buffering events and audio quality (WO2017144506 from Ericsson).

WO2016040590, filed by Qualcomm, describes a method of removing noise caused by interference between two channels when a device is connecting to two or more mobile communication networks. As the interference is non-linear, it is mathematically very expensive to calculate all possible cases of the interference, but the expected value of the interference can be learned and accounted for by using a multi-layer perceptron.

This patent family includes a member granted in the United States.

Massachusetts Institute of Technology (MIT) is a private research university in

[57] Nokia finalizes its acquisition of Alcatel-Lucent, ready to seize global connectivity opportunities, Nokia media release (2 November 2016) https://www.nokia.com/about-us/news/releases/2016/11/02/nokia-finalizes-its-acquisition-of-alcatel-lucent-ready-to-seize-global-connectivity-opportunities/

Cambridge, Massachusetts and hosts an active group of researchers working on different aspects of machine learning including theoretical foundations, optimisation algorithms, and a variety of applications such as vision, speech, healthcare, materials science, NLP and biology.[58] WO2016054384 describes a method to assess the risk of downloading and/ or using mobile applications, in which operational features of mobile applications and malware are extracted and used to train classifiers. The assigned classifications and the correlations between them are used to generate a risk rating. User-defined security guidelines can also be incorporated into the risk classification scheme.

Environment and mining

Environment and mining is a broad category. In the mining sector, machine learning is being applied to mineral exploration, self-driving vehicles and automation of mineral processing. This sector includes 3495 applications, with 30 per cent granted.

Environment and mining real world applications include inventions that assist in extracting natural resources, and in protecting and monitoring environments. In the mining industry, machine learning is being used to assist with prospecting for minerals, oil and gas; and automation of drilling and mineral processing. Applications for environmental science focus on the use of image analysis and remote sensors to measure pollutants in air and soil, and to monitor and/or predict natural disasters.

WO2018072815 (Schlumberger Holdings) is an example of a real world application of machine learning in the mining sector. This invention relates to a method of identifying the location of a subsurface seismic layer for optimal drill placement. A subsurface apparatus generates a three dimensional image of the subsurface area and provides it to a machine learning process that then predicts the presence of a subsurface seismic layer in the three dimensional image.

In environmental science, US2014099020 (Keimyung University Industry Academic Cooperation Foundation) shows a method to predict the occurrence of a forest fire and eliminate false alarms using spatiotemporal Bag of Features (BoF) and random forest algorithms.

Energy

Real world applications of machine learning in the energy industry include a range of technologies both for generating energy and for the automated management of grids. These technologies maximise output from generators, and minimise the impact of blackouts. The category is the third largest area of machine learning, with 3481 patent families, and over 1000 patent families granted.

State Grid have filed 17 per cent of patent filings (590 patent families) in the energy applications category. State Grid controls the world's largest electricity supply network, providing power transmission, power system maintenance and power network maintenance, as well as other services across China. State Grid also has substantial global holdings in the energy sector, including the National Grid Consortium of the Philippines and investments in Australia, Brazil, Italy, Portugal and Greece.

Of these 590 patent families, five (or less than one per cent) have family members filed in countries other than China. These five international applications by State Grid cover a range of applications, and include WO2016141739 for establishing a supply and demand early warning model in the electricity market; WO2015081660 for forecasting residential short-term load using neural networks, and WO2016033883 for providing power transmission line gallop risk early warning prediction using adaptive boosting (AdaBoost) with historical and meteorological data. State Grid's WO2018045642, for predicting electrical busbar loading using random forest classification and data mining to provide more accurate decision support for power grid planning and real time scheduling, was also filed in Australia (AU2016325186), but lapsed due to lack of acceptance in examination in early 2019.

The renewable energy sector is also well represented. Algorithms have been developed that use multiple datasets to predict weather patterns and forecast the availability

[58] MIT Machine Learning Group (accessed 9 October 2019) http://machinelearning.mit.edu/

of renewable resources, for solar (WO2016210102, Qatar Foundation) and wind (AU2013207551, Tata Consultancy Services Ltd; WO2014063436, State Grid) power.

Energy applications of machine learning also include smaller scale domestic uses. For example, WO2018105784, filed by the South Korean company Ontest Corporation, describes a home energy management system using sensors that record wall temperatures and use machine learning techniques to control heating and cooling of the indoor spaces.

Health care

The health care category captures a variety of inventions that help to deliver better diagnostics and treatment for patients and streamline administration of hospitals. There are 3060 patent families for real world machine learning applications in healthcare. Technologies include using machine learning to assist with medical imaging (e.g. reading magnetic resonance imaging (MRI) scans) and biological sample testing.

Machine learning is also augmenting medical decision making by predicting how a patient's condition will change based on their medical history and what side-effects a particular treatment could cause (WO2017019706, held by Alphabet, the parent company of patent applicant Google), and potential risks associated with genetic variants identified through genomic screening (WO2016209999, Counsyl).

Predictive maintenance

The use of machine learning in predictive maintenance has applications across the whole economy. This sector covers technologies that are critically important to manufacturing and other industrial sectors, in which machine learning can be used to monitor equipment, structures or vehicles, and to predict when maintenance or repair work is required. These methods can include the capture and analysis of acoustic signals to predict when a machine is likely to fail. This category contains 2362 patent families, about one third of which have been granted.

Clearly, patents relating to predictive maintenance applications can also be applied to other applications. In this analysis, we set out to identify applications that primarily related to predictive maintenance systems as a priority.

An example patent application related to predictive maintenance is WO2015070513 (State Grid). This patent application describes the use of machine learning employed to detect patterns in insulation defects related to gas insulated switchgear (GIS) used in ultra-high voltage grids. Pattern recognition allows control of partial discharge of power, which prevents electrical breakdown or failure.

More broadly, WO2017213810 (Exxonmobil Research and Engineering Company) describes the use of machine learning techniques to detect and analyse visual and audio events. This information can be used to control a range of automated industrial processes.

Industry

The industry and manufacturing real world applications include patent applications for industrial automation, which.

The industry category includes technologies of key importance for manufacturing and other industrial sectors. Patent applications cover addition of sensors to industrial equipment for optimising performance and monitoring, optimising the supply chain through logistical planning, and use of machine learning to help design better products. This category includes 1549 patent families.

For example, machine learning programs are being used to improve workflows in the oil and gas industry, allowing users to more easily interrogate information repositories and recommending courses of action to improve workflows (for example WO2016153470 from Haliburton Energy Services Group).

An example in this category is a patent application that describes the use of sensors and data analysis to measure the health of structures. National ICT Australia (NICTA) (now part of Commonwealth Scientific and Industrial Research Organisation (CSIRO)) (WO2014176625) measures vibrations caused by traffic for structural health monitoring for bridges. Ambient factors such as wind are measured, and machine learning is used to determine the health of the structure. This cross cutting technology is also relevant to other applications such as predictive maintenance and transport and logistics.

CASE STUDY: MACHINE LEARNING-ASSISTED MEDICAL IMAGING

A large amount of diagnostic effort across most fields of medicine relies on the interpretation of various forms of imaging. Image analysis is key to areas of medicine from cancer detection and characterisation (both internal and on the skin) to fracture evaluation, diagnosis of infection, circulatory studies, blood clot and aneurism detection and many more.

Machine learning is already applied to medical imaging problems, including ophthalmology and skin lesion detection,[59] and research suggests this market will grow to $2 billion USD by 2023.[60] Anticipation of this emerging market has led large computer manufacturers such as graphics processing specialist NVIDIA to release purpose built platforms for medical image applications.[61]

The medical profession stores data that is highly detailed, systematically archived and classified by professionals with years of professional experience and training. This gives rise to a large amount of pre-tagged ground-truth data. The combination of highly specialised pattern recognition tasks, vast data streams and a corpus of tagged training data means that medical imaging is an area well suited to enhancement by machine learning. As well as this, improvements in medical imaging are very important to both patients and to society. Enhancement of the capabilities of a limited supply of medical professionals promises to be one of the most valuable and rapidly adopted applications of machine learning.

Patenting in this area is active, with more than 1250 families filed since 2012 and 403 in 2017 alone. The global leader in medical imaging patenting is the German company Siemens AG with 86 patent families. Their patent US2014185888 relates to non-invasive lesion detection from medical images such as computed tomography (CT) and MRI. Their technique consists of two major components. Image-related information such as absorption and geometry are collected and interpreted. The analysed images are examined together with patient medical history data regarding conditions that may cause abnormalities in the images in consideration (for example osteoporosis in bone images). This combination of patient-specific information with medical image information is expected to give more reliable information than image information alone.

Australia has a technology strength in medical technologies in general, and medical imaging in particular. An example of an Australian patent application comes from the CSIRO (WO2016019439). This patent application relates to imaging the interior of a body cavity, with a specific use in effectively detecting colorectal abnormalities. A machine learning system is trained with real image data and produces simulated images that can be used to generate models. These models can be used for a variety of purposes including for training purposes (reducing the burden on real patients) and for enhancing the capability of clinicians by showing in real time any areas that have not been examined.

[59] How AI and Deep Learning Are Revolutionizing Medical Imaging, GEHealthCare (08 April 2019) https://www.gehealthcare.com/long-article/how-ai-and-deep-learning-are-revolutionizing-medical-imaging

[60] Harris, S (16 October 2018) What's new for machine learning in medical imaging, Signify Research, https://www.signifyresearch.net/medical-imaging/signify-ai-medical-imaging-white-paper/

[61] Salian, I (4 March 2019) How AI is changing medical imaging, NVIDIA blogs, https://blogs.nvidia.com/blog/2019/03/04/how-ai-is-changing-medical-imaging/

Transport and logistics

The transport and logistics area is active in the development of driverless vehicles, from cars to trains and aeroplanes, and sensors for monitoring engine performance and predicting the need for maintenance. It also includes artificial intelligence used in space technologies and spacecraft. This category has 1466 patent families.

A patent application filed by Netposa Technologies Ltd is an example of an image analysis systems that uses deep learning techniques to read license plate numbers (WO2018028230).

Improving safety is an area of commercial interest. For example, Uber Technologies have filed patent application WO2018073662 for a travel safety system that collects data, such as dangerous driving or interpersonal conflict incidents, to predict future safety incidents.

Security and military

The security and military category covers a broad range of inventions, including facial recognition, biometric scanning (facial, fingerprint and voice recognition), fraud detection, reading sentiment in written text, and defence applications. This category includes 1339 patent families.

Examples of the use of these technologies include a patent application by Shanghai Jiao Tong University for a system to lip reading based on sound received by a smart device. It identifies the characteristics of a person's lip motion to detect individual words and uses deep learning methods alongside sound wave signal processing technology to model and predict sounds CN107784215.

An example of facial scanning includes a patent application by Facebook for a face tracking system that predicts face alignment by using facial anchor points to create a mesh shape and to track facial movement (Figure 9). This system reduces to a model that can be used on devices to predict facial movement WO2018118099.

FIGURE 9 IMAGE FROM FACEBOOK'S PATENT APPLICATION WO2018118099

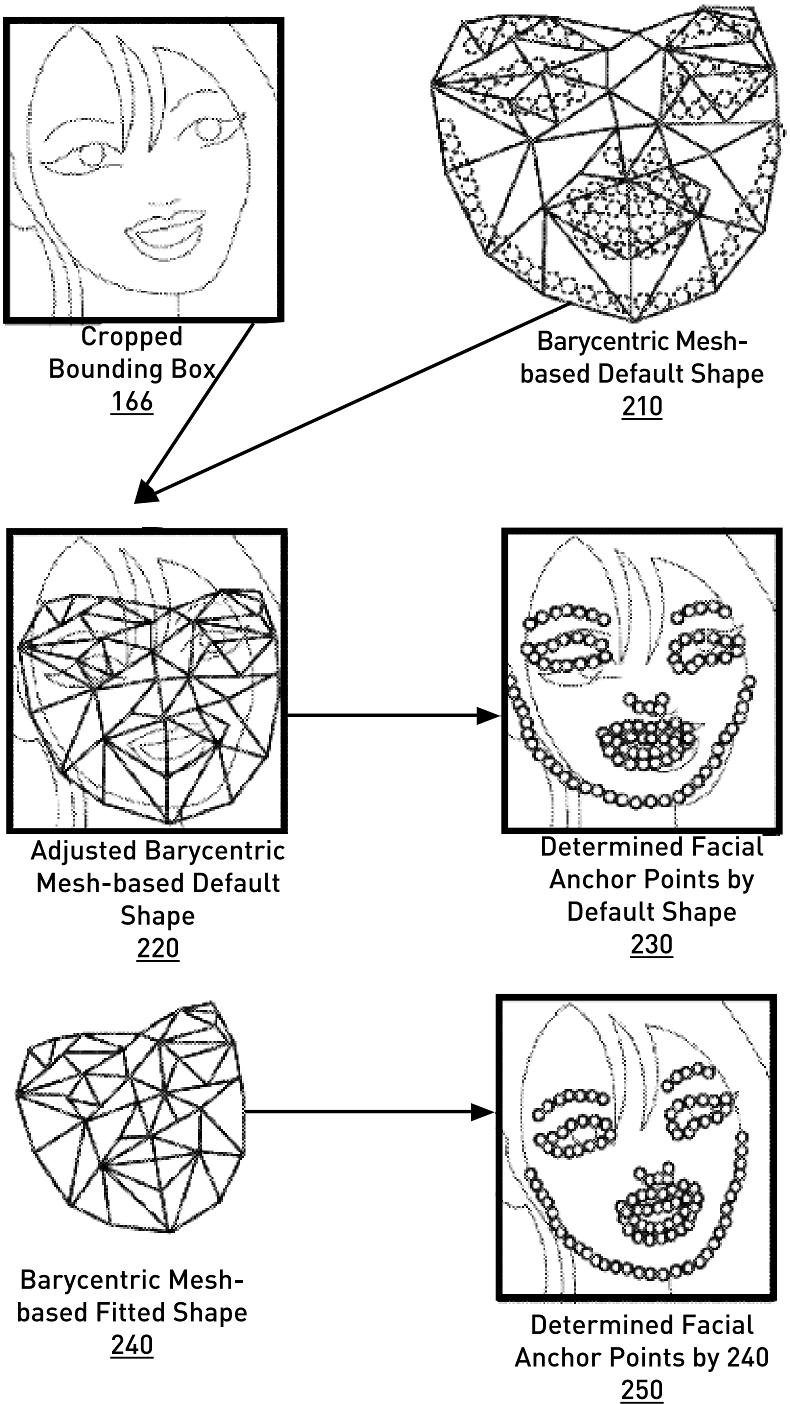

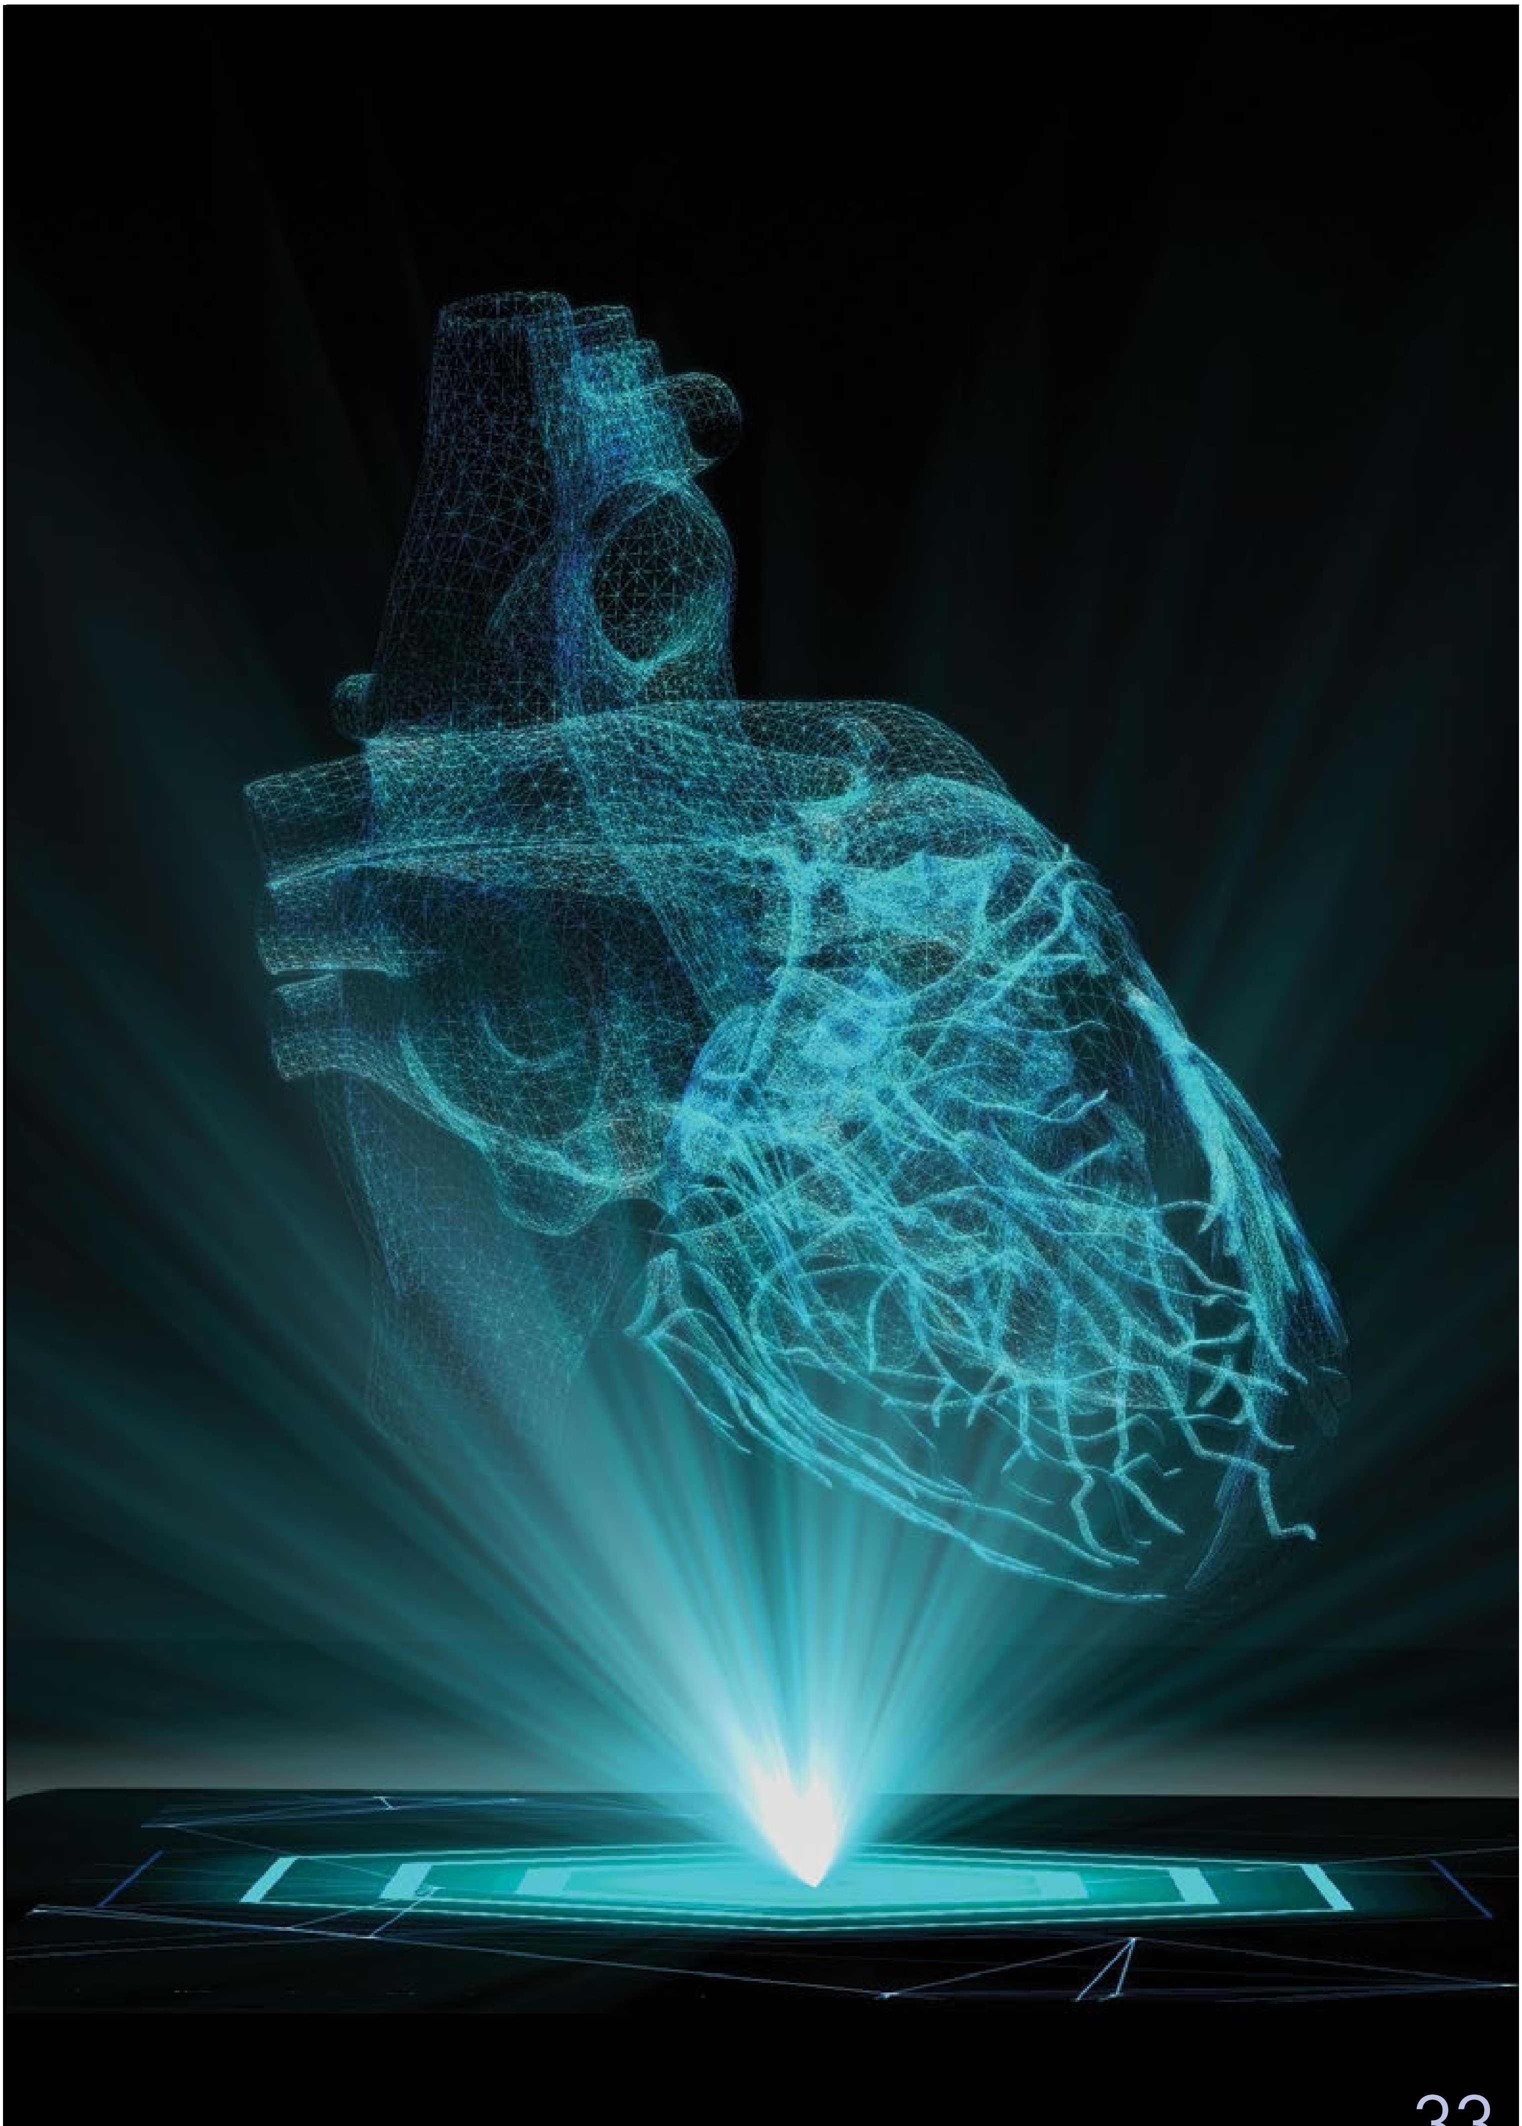

Agriculture

In agriculture, machine learning is being used in robots that reduce the cost of field labour, using image analysis to monitor crops and soil quality, and for quarantine and food inspection.

Agriculture machine learning real world applications reduce the cost of farm labour through AI-controlled field robots, use of imaging to monitor soil and crops, use of predictive analytics tools to maximise production, and also development of improved quarantine and food inspection processes. We found 1083 patent families relating to the use of machine learning in agriculture, indicating that this area has strong potential growth.

An example of the use of machine learning in water resource use is from CN107578116 (National Dong Hwa University, Taiwan). This invention aims to improve the distribution of agricultural water resources during drought using an immune particle swarm optimisation algorithm. This maximises crop yields even in drought conditions.

Vision enhanced drones for precision farming have been developed by Intel (US9898688). This system uses neural networks to analyse and classify agricultural conditions based on data collected by drones and generates a three dimensional map showing agricultural conditions. This method is sensitive enough to distinguish between crop disease, crop damage caused by animals and water damage.

Business and finance

The business and finance category covers patent applications relating to AI used in business investments, such as algorithms for making trades, AI assistants for managing portfolios and assessing risk, fraud detection, and predictive analytics. This category has 994 patent families. The business and finance category brings together a range of technologies involved in finance, entertainment, and education.

Business applications include, for example, the analysis of customer behaviour. Machine learning is used to analyse customers' credit card transactions to determine their interests, and then recommends credit card purchase special offers that are aligned with those interests (WO2014204693). Personalised recommendations are also important in emerging inventions in entertainment, as are virtual and augmented reality for games.

In finance, machine learning is being used in business investments, such as algorithms for making trades, automated assistants for managing portfolios and assessing risk, and predictive analytics.

Patent application WO2018074902 by Samsung describes a system for mobile payments that, when a request to send money from a sender to a receiver is made, a computer program generates a list of eligible providers based on the transaction, and then notifies a provider to transfer the money.

Fraud detection is vital to a viable business and finance sector. Patent application US2017200164 (Korea Internet and Security Agency) is an example of machine learning in detection of fraudulent transactions.

CASE STUDY: PROPERTY VALUATION

A niche application of machine learning techniques is in the valuation of property. Property is big business in Australia. According to some sources the total value of the Australian residential property market alone is more than $7.5 trillion AUD.[62]

Property valuation is difficult due to high variability between properties. Properties must also be valued at different times and by different parties, not only at the time of sale but also for assessments such as land taxes. There are many online AI tools available for valuing residential and commercial properties.[63,64] In the United States, real estate databases such as Zillow, Trulia and Redfin provide AI-assisted home valuations as part of their database offering.[66] Recently the University of New South Wales has developed a suite of automated land valuation software tools called Value Australia.[66]

While patenting activity appears relatively muted in property valuation overall, some patent applications from across the world were identified in this analysis. Examples included US2018068329 from IBM relates to automated property valuation using neural networks. This uses images at various levels of zoom to see features of the property and its surrounding area to predict the price of the property based on training data from the same geographical area.

Patent CN107230113 from the Jiangxi University of Science and Technology is also related to the valuation of a property based on information for recorded transactions for similar properties.

Ongoing activity in property valuation seems poised to enable rapid, accurate and often free valuations of property, both residential and commercial, resulting in useful outcomes for property investors as well as those who need information about the property market, such as government agencies.

Retail

The real world applications of machine learning in retail, with 909 patent families identified, includes inventions where machine learning is used to create personalised online product recommendations, personalised advertising and customer engagement. It covers all web-based shopping and e-commerce platforms. It also captures applications for predicting product demand.

The retail category captures technologies that are being used to increase customer engagement and improve sales. This includes machine learning to generate product recommendations, chatbots for customer assistance, and also for the prediction of product demand and supply management.

The biggest retail players are well represented here. Ebay holds patent applications such as WO2018089762, for personal assistance with online shopping. This system leverages existing inventories and curated databases to provide intelligent answers to a human user, using image analysis including text-in-image searching.

Yahoo (US2016125314) describes the use of machine learning techniques to analyse information about a user to provide targeted advertisement selection and display, which assists customers as well as improving sales.

Overall, this section on real world applications is exciting and demonstrates the diversity and significance of developments in machine learning across all areas of our lives and economy. In the next section, we will take a closer look at the development of foundational machine learning technologies that are allowing these changes.

[62] Yardney, M (21 December 2018) State by state: An end-of-year update on Australia's property markets, Smart Company, https://www.smartcompany.com.au/industries/property/end-year-update-australia-property-markets/

[63] Chagani, E (28 March 2018) The potential of machine learning real estate valuation models, Cornell University, https://blog.realestate.cornell.edu/2018/03/28/machine-learning/

[64] Wiggers, K (24 January 2019) Geophy raises $33 million to apply AI to commercial property appraisal, Venture Beat, https://venturebeat.com/2019/01/24/geophy-raises-33-million-to-apply-ai-to-commercial-property-appraisal/

[65] Tarver, E (17 April 2019) The 5 Best Alternatives to Zillow and Trulia, Investopedia, https://www.investopedia.com/articles/markets/100215/5-best-alternatives-zillow-trulia.asp

[66] Johnston, M (4 October 2019) UNSW partners with govt, industry for AI land valuation, ITNews https://www.itnews.com.au/news/unsw-partners-with-govt-industry-for-ai-land-valuation-531896

06

Foundation technologies

Now we have considered the applications of machine learning – both core capabilities and real world applications, we will turn to the foundational technologies that support the continued development of machine learning.

We identified 24,091 patent families relating to machine learning foundational technologies, and divided these into two categories: algorithms and architecture (see Appendix C: Technology Analysis for more information about the sub-technology classification).

Figure 10 shows the number of patent families in each category by their grant status: granted (where at least one family member is granted in any jurisdiction) or patent application (a patent family with members seeking patent protection, or in which all members are lapsed/withdrawn without grant).

Overall 26 per cent of the patent families relating to the machine learning techniques have at least one family member granted.

The algorithms category includes patent families that describe supervised, semi-supervised or unsupervised learning, reinforcement learning, other niche learning methods such as ensemble learning and extreme learning, and general machine learning techniques. This category accounts for 85 per cent of patent families relating to machine learning techniques, totalling 20,491 patent families, of which 30 per cent have a granted family member.

The architecture category includes patent families that describe physical realisation or hardware implementation of neural networks, neural network architecture and its optimisation. Only 15 per cent of the patents families in machine learning techniques relate to this category, with 3,600 patent families, of which 20 per cent have a granted family member.

In the next section of this report, we will take a closer look at each of these categories and their top applicants.

FIGURE 10 PATENT FAMILIES RELATING TO MACHINE LEARNING FOUNDATIONAL TECHNOLOGIES BY THEIR GRANT STATUS

Source: PATSTAT 2018 Autumn Edition

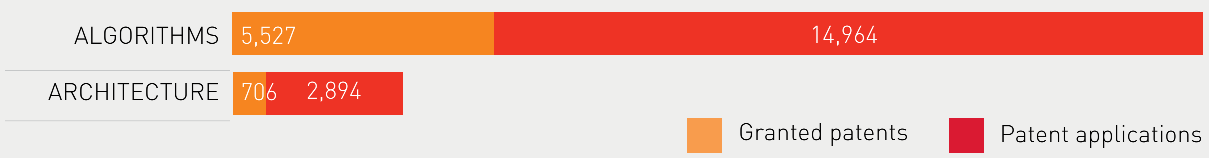

Top applicants

The number of patent families filed by an applicant in a specific technology can be indicative of their interest and market presence or desire to build and maintain a market share. Figure 11 displays the top applicants by their patent families on machine learning foundational technologies.

FIGURE 11 PATENT FAMILIES RELATING TO MACHINE LEARNING FOUNDATIONAL TECHNOLOGIES BY APPLICANT

Source: PATSTAT 2018 Autumn Edition

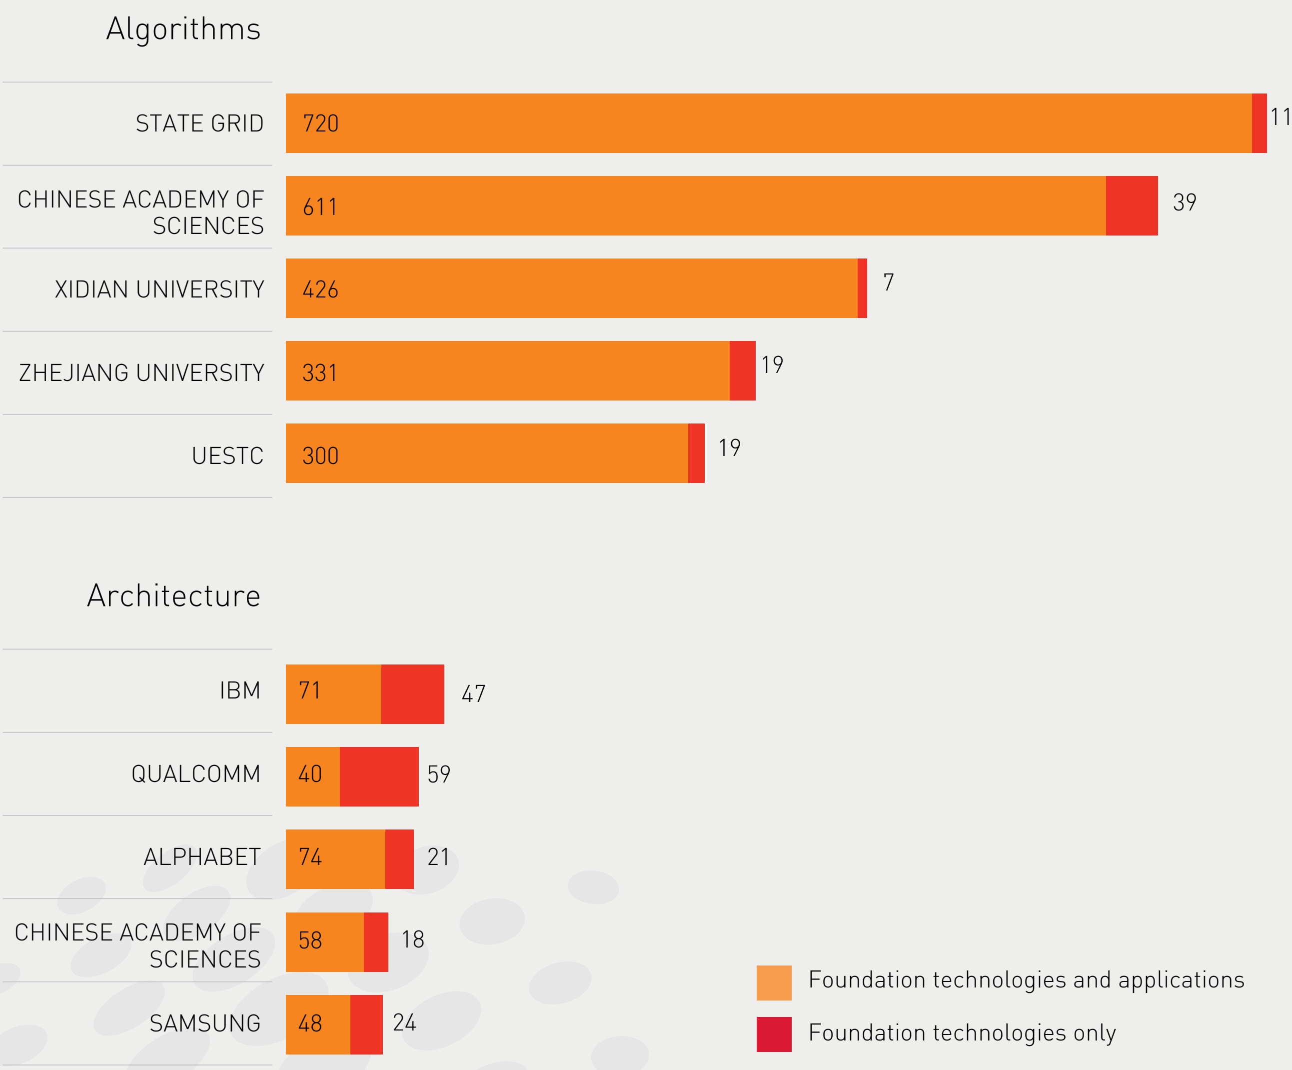

The top five filers of patent applications relating to algorithms are Chinese. Most patent families filed by these applicants describe both algorithms and their applications, which may be an indication of interest in applying machine learning techniques to real world applications.

Patent filings relating to architecture are led by American applicants. Architecture patent applications show a greater proportion of filings in techniques development only.

State Grid is the top applicant relating to algorithms, with a total of 763 patent applications in either foundational technologies or foundational technologies and applications. State Grid is the world's second largest company on the Fortune Global 500 list,[67] and is the largest public utility company in the world, supplying power to over 1.1 billion people in 88 per cent of Chinese national territory.[68]

State Grid focuses on AI developments in power grid control, power distribution and utilisation networks, and green energy. Their patent portfolio is focused on application of machine learning techniques, especially in energy and predictive maintenance. The company also contributes to the development and use of AI-related smart algorithms and robotics. State Grid's grid operation and management systems involve the collection and analysis of various data, using AI technologies including image processing, voice recognition and big data analysis.[69] State Grid is one of the two state-owned enterprises that operate the power network in China, the other being China Southern Power Grid (CSPG), which has a service area spanning one million square kilometres and a population of 252 million people.[70] CSPG is also an active player in the machine learning field with a total of 263 patent families filed since 2012.

The Chinese Academy of Sciences is active in both algorithms and architecture, and hosts 104 research institutes, including the Institute of Intelligent Machines.[71] This institute has become one of the leading national research centres in sensor technology and AI. The research centre focuses on biomimetic sensing and intelligence related to robotics and human-computer interactions, aiming to advance digital agriculture, sensor performance, control systems and administration.[72]

Xidian University is a public research university located in Xi'an, China,[73] ranking third in algorithms related inventions. Their portfolio also includes inventions relating to neural network architecture and implementation. In 2017, Xidian University established the School of Artificial Intelligence, which focuses on innovative research and development in AI, including fields such as intelligent information processing, intelligence and image understanding, visual computing and collaborative cognition.[74]

Zhejiang University is a public university located in Zhejiang, China. Their patent portfolio predominantly includes algorithms related inventions and they rank fourth worldwide in this category. The university's Institute of Artificial Intelligence focuses on research areas including AI theory, computer aided design and computer graphics, client information management systems, and other advanced manufacturing technologies.[75]

UESTC ranks in fifth place for filing algorithms related patent applications. This is a public university located in Sichuan province, China. In 2014, UESTC established an AI research institute called the Statistical Machine Intelligence and Learning (SMILE) Lab. The lab studies algorithms for solving various data analytics problems relating to AI and machine learning.[76] Their patent portfolio also includes inventions relating

[67] Global 500 (accessed 9 October 2019) http://fortune.com/global500/

[68] State Grid Corporation of China, Corporate Profile (accessed 9 October 2019) http://www.sgcc.com.cn/html/sgcc_main_en/col2017112307/column_2017112307_1.shtml

[69] China Institute for Science and Technology Policy at Tsinghua University (July 2018) China AI Development Report 2018, http://www.sppm.tsinghua.edu.cn/eWebEditor/UploadFile/China_AI_development_report_2018.pdf

[70] China Southern Power Gridgrid, Company Profile (accessed 9 October 2019) http://eng.csg.cn/About_us/About_CSG/201601/t20160123_132060.html

[71] Chinese Academy of Sciences, CAS Institutes (accessed 9 October 2019) http://english.cas.cn/institutes/

[72] Chinese Academy of Sciences, Introduction for the Institute of Intelligent Machines (accessed 9 October 2019) http://english.hf.cas.cn/r/ResearchDivisions/IIM/

[73] Xidian University, Overview (accessed 9 October 2019) https://en.xidian.edu.cn/About_Xidian/Overview.htm

[74] Xidian University, School of Artificial Intelligence (accessed 9 October 2019) https://en.xidian.edu.cn/info/1018/1945.htm

[75] Zhejiang University, Institute of Artificial Intelligence (accessed 9 October 2019) http://fit.zju.edu.cn/english/redir.php?catalog_id=2443

[76] University of Electronic Science and Technology of China (accessed 9 October 2019) https://en.uestc.edu.cn/

to neural network architecture and implementation.

International Business Machines Corporation (IBM), the top applicant in neural network architecture and implementation, is an American computer and information technology consulting company. The company's research centre, IBM Research, explores AI and machine learning techniques, with a focus on three areas: advancing AI, scaling AI, and trusting AI.[77] The company's patent profile includes fundamental developments in machine learning techniques as well as applications of machine learning techniques in telecommunication and networks, retail, medicine and healthcare.

Qualcomm is an American telecommunication technology company, with a focus on creating, developing and commercialising power-efficient on-device AI to make mobile phones a pervasive AI platform.[78] In this analysis, Qualcomm is the second ranked applicant for inventions relating to neural network architecture and implementation. More of their inventions focus on the developments in machine learning foundational technologies than on applications of machine learning.

Alphabet is an American multinational conglomerate that owns multiple subsidiary companies including Google, Waymo and DeepMind Technologies.[79] Google already uses machine learning to power their products including search engines, Google maps, YouTube, Gmail and Google drive, with new techniques leading the way from mobile-first to AI-first in the world.[80] Waymo, Google's self-driving car, is committed to the use of applied AI and machine learning techniques to make roads safer and improve mobility for everyone.[81] DeepMind focuses on developing algorithms that can learn to solve any complex problem without needing to be taught how.[82] One success story is DeepMind's AlphaGo, the first computer program to defeat a Go world champion.[83] Alphabet has a strong patent portfolio across the algorithms and architecture categories, and ranks third for inventions relating to neural network architecture and implementation.

Samsung, a South Korean multinational conglomerate, is developing technologies including AI, data intelligence, intelligent machines and next generation communication, with the aim of creating cutting-edge AI technologies for application in all of its products and services.[84] Samsung ranks in fifth place for patent applications relating to neural network architecture and implementation. They also have filed patent applications relating to algorithms and real world application of machine learning techniques.

Algorithms

Foundational technologies related to algorithms are sub-divided hierarchically into supervised and semi-supervised learning, unsupervised learning, reinforcement learning, niche learning techniques and general machine learning techniques (see Figure 12).

FIGURE 12 PATENT FAMILIES RELATING TO ALGORITHMS

Source: PATSTAT 2018 Autumn Edition

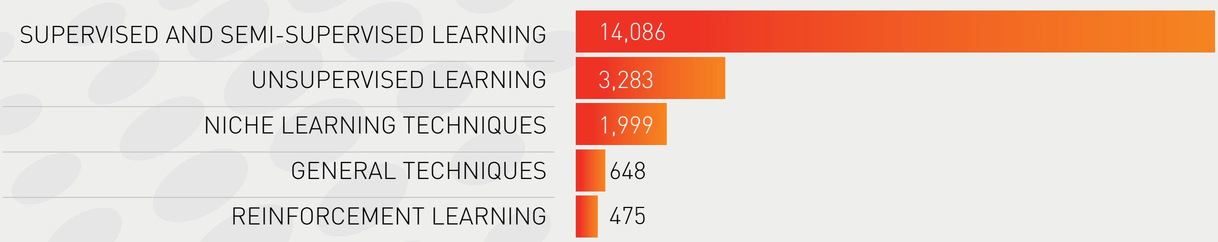

[77] IBM Research (accessed 9 October 2019) https://www.research.ibm.com/artificial-intelligence/
[78] Qualcomm Research & Invention (accessed 9 October 2019) https://www.qualcomm.com/invention/artificial-intelligence
[79] Alphabet, G is for Google (accessed 9 October 2019) https://abc.xyz/
[80] AI Multiple (21 february 2019) Google is AI first: 12 AI projects powering Google products, https://blog.aimultiple.com/ai-is-already-at-the-heart-of-google/
[81] AI Multiple (11 December 2018) Alphabet's AI Investments in 38 Companies: In-depth Analysis, https://blog.aimultiple.com/alphabet-ai/
[82] DeepMind, About Us (accessed 9 October 2019) https://deepmind.com/about/
[83] AlphaGo (accessed 9 October 2019) https://deepmind.com/research/alphago/
[84] Samsung Research Areas (accessed 9 October 2019) https://research.samsung.com/research-areas

Supervised and semi-supervised learning is the largest sub-category, with a four-fold increase in patent family filings between 2012–16. This category includes classification and regression techniques, supervised dictionary learning, supervised neural networks (including supervised deep learning) and boosting algorithms.

The second largest sub-category is unsupervised learning, including cluster analysis techniques, blind signal separation techniques, unsupervised dictionary learning, and unsupervised neural networks (including unsupervised deep learning, autoencoders, self-organising maps and GANs).

Reinforcement learning accounts for two per cent of the patent families in algorithms and has also increased four-fold in five years. This sub-category includes dynamic programming, learning automata and reinforcement neural networks.

The niche learning techniques sub-category includes patent families that discuss ensemble learning, extreme learning, transfer learning, genetic algorithms (with machine learning), rule learning, swarm optimisation, online learning, feature learning, graphical models, mixture models and relational learning.

The general techniques sub-category contains patent families that mention machine learning or neural networks and do not fall into other categories of machine learning foundational technologies or machine learning applications (see Appendix C).

The following patent applications are examples of algorithms applications in these sub-categories.

Patent application CN102831404, filed by the Chinese Academy of Sciences, describes a gesture detection method using machine learning classifiers (supervised learning) on both colour and grayscale versions of an image. The method employs a sliding window to extract features from a region of an image, an adaboost classifier for classifying features, and a random forest classifier for classifying point-pair features (a type of image feature that is used for rigid object detection in point cloud data). This patent family has a member granted in China.

An patent application describing an unsupervised learning technique is WO2015088780, by the University of Southern California, which hosts the Center for Artificial Intelligence in Society for conducting research in AI to help solve difficult social problems facing the world.[85] This application proposes a noise enhanced clustering technique, consisting of applying an iterative clustering method (k-means clustering algorithm or competitive learning algorithm) along with artificial noise, which is used to speed up convergence and increase the robustness of the final clusters. This clustering technique may find application in a range of areas including computer vision systems, target tracking, data/document clustering, credit risk analysis, speech recognition and bioinformatics. Patent protection for this patent family is being sought in the United States.

Patent application US20180129974 discloses a deep reinforcement learning technique relating to control systems. The technique determines an action (for example the control of speed and angle of a nozzle in a material deposition process) based on environment observations (e.g. LiDAR or video images). The outcome of the action is then assessed against the desired outcome and this is fed back into the control system to inform the choice of subsequent actions, along with continued environmental observation. The deep reinforcement learning technique uses a convolutional neural network or a deep autoencoder. This patent application was filed by United Technologies Corporation, an American aerospace and building technology company with headquarters in Connecticut.[86] Their research centre focuses on developing machine learning techniques for intelligent systems that allows for self-diagnosis and repair.[87] Patent protection for this invention is being sought in the United States and Europe.

[85] USC Center for Artificial Intelligence in Society (accessed 9 October 2019) https://www.cais.usc.edu/
[86] United Technologies (accessed 9 October 2019) https://www.utc.com/
[87] United Technologies Research Center (accessed 9 October 2019) http://www.utrc.utc.com/platform-technologies.html

CASE STUDY: TRANSFER LEARNING – THE NEXT EDGE OF MACHINE LEARNING

If AI is the simulation of human intelligence processes by machines, then machine learning could be considered a simulation of human learning by machines.

Humans don't learn everything from scratch. We leverage our previous experiences and apply earlier information to new tasks and environments. Inspired by this process, transfer learning leverages and transfers knowledge acquired from solving one task in one domain to solve other tasks in other domains.[88] Transfer learning marks a paradigm shift in machine learning, from models designed to learn in isolation to models designed to learn from shared past knowledge.

For example, a model trained using traditional techniques to detect pedestrians in night-time images would fail to detect pedestrians on a day-time image and would require a new model to be trained separately with day-time images. Transfer learning can instead leverage the knowledge gained from training on night-time images and apply it to training for detecting pedestrians in day-time images, thereby using fewer training images and reducing the time to train.

Transfer learning can be applied to machine learning techniques using algorithms designed for each type of technique (i.e. to supervised techniques using inductive transfer learning, unsupervised techniques using unsupervised transfer learning, and reinforcement techniques using transductive transfer learning) and is divided across the categories in Figure 12, including niche learning techniques where supervision is not specified. Methods of transfer learning can be interchangeably used in these categories between models in different domains. These methods include instance transfer, feature representation transfer, parameter transfer and relational knowledge transfer.

Our analysis shows dramatic growth in filing trends for transfer learning from eight patent families in 2012 to 53 patent families in 2017, with the number of filings doubling between 2016–17. Transfer learning is widely used with unstructured data and this is reflected in its application. Of the 134 transfer learning patent families, 54 per cent are used for image and video analysis while 19 per cent relate to speech and text analysis.

Examples of patent applications on transfer learning include WO2017173489, originating from Australia and filed by Deakin University, which hosts an AI and Data Analytics Research Cluster.[89] This patent application describes a method for calculating control parameters to make an optimised manufactured product. Designed for use in early product development, this method helps reduce the number of experiments required to determine optimal control parameters, by leveraging knowledge gained from past experiments to adjust control parameters in subsequent experiments.

General Electric Company[90] has filed patent application US2017300605, describing the use of transfer learning for predictive maintenance. This application describes a method for creating a predictive damage model for aviation assets, by predicting behaviours of a given physical asset and then using transductive transfer learning from this model (with labelled data only available in the first model) to create a second predictive damage model to predict the behaviour of a second physical asset in a separate domain, such as new or expanded asset types or new flight conditions.

Another example is granted patent US10007863 by applicant Gracenote, a subsidiary of Nielsen Holdings. This patent describes a method for accurately detecting logos in media content on media presentation devices using neural networks and transfer learning.

As predicted by the Chief Scientist at Baidu and professor at Stanford University, Andrew Ng, transfer learning is becoming the next driver of machine learning commercial success after supervised learning.[91]

[88] Pan, S. J. and Yang, Q., 'A Survey on Transfer learning', IEEE Transactions on Knowledge and Data Engineering, 22 (10) Oct 2010 https://ieeexplore.ieee.org/document/5288526

[89] Deakin University Artificial Intelligence and Data Analytics Research Cluster (accessed 9 October 2019) https://www.deakin.edu.au/information-technology/research/artificial-intelligence-and-data-analytics

[90] Company Information (accessed 9 October 2019) https://www.ge.com/news/company-information

[91] Kim B-H (6 May 2018) NIPS 2016 tutorial: 'Nuts and bolts of building AI applications using Deep Learning' by Andrew Ng, https://www.youtube.com/watch?v=wjqaz6m42wU

Architecture

Foundational technologies related to architecture are sub-divided into neural network architecture, physical realisations, and architecture optimisations (Figure 13).

FIGURE 13 PATENT FAMILIES RELATING TO ARCHITECTURE

Source: PATSTAT 2018 Autumn Edition

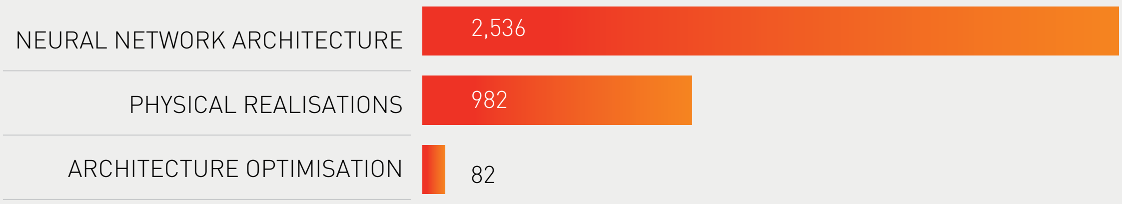

Neural network architecture is the largest sub-category, with 70 per cent of patent families. This sub-category includes patent families describing the architectural layout of neural networks and how the neurons are interconnected.

Physical realisations are the second largest sub-category, with 27 per cent of patent families in the architecture category. This sub-category includes implementation techniques used to physically construct the neurons or neural networks.

Architecture optimisation includes patent applications that describe optimisation of neural network architecture by adding or deleting nodes or connections. This sub-category represents only two per cent of patent families in the architecture category.

The following patent applications are examples of architecture applications in these sub-categories.

Patent application WO2016146278, by IBM, describes a system that maps a neural network onto a neurosynaptic substrate (a computing architecture specifically optimised for use with neural network algorithms) based on information about a graph (or 'adjacency matrix') representation of that neural network. Users can interact with the mapping system in order to trade-off between accuracy and resource utilisation of the substrate. This patent family has a member granted in the United States.

Qualcomm has filed WO2015038311, a patent application describing optimisation of a neural network system in which computational resources are limited by structural plasticity. Structural plasticity is the ability for elements of a neural network to change their shapes and thereby connections in the network. This invention alters the location of synapses connecting particular neuron pairs based on timing information from previous executions. One family member of this patent family has been granted in the United States.

Patent application WO2016123409, originally filed by Google, relates to a neural network system including a normalisation step (a 'normalisation layer') between successive layers of the neural network. When the neural network is being trained, the normalisation layer receives the output of a given layer of the neural network for each of the training examples in a given batch, and then normalises these outputs based on the statistics of the batch as a whole, before providing the normalised output to the next layer of the neural network. The application claims this technique increases the training speed of a neural network, and can also reduce the need for regularisation (mechanisms to restrict the degree of overfitting of a given model). This patent family has members granted in Europe, Australia and Spain.

07

Australian Applicants

To better understand Australian patenting activity, we analysed filings by Australian applicants.

This report captures 59 patent families filed by Australians, with the largest growth between 2015 (10 patent families) and 2016 (19 patent families). Australia is the primary patent destination for Australian applications, with 54 per cent of families filed here, followed by the United States with 34 per cent. Figure 14 shows the divide across applications and foundational technologies.

Applications of machine learning are clearly the focus for Australians. Figure 14 shows 97 per cent of patent families discuss an application and only 44 per cent discuss a technique, including a 40 per cent overlap. Figure 15 shows a further analysis of the filings by Australians, by application type.

Australian applicants favour analytics and processing (41 per cent) and image and video analysis (31 per cent) in the development of machine learning core capabilities.

Patent applications in real world applications of machine learning by Australian applicants are filed across all technology areas, with a focus on health care (Figure 16), which is an Australian technology strength.

Figure 17 shows the top five Australian applicants. The top patent filer is CSIRO, with five patent families. In second place is HRO Holdings with three patent families. In shared third place are four entities: Atlassian, CRC Care, NewSouth Innovations and the University of Technology Sydney, with each having two patent families.

Apart from the top five Australian applications, all other Australian applicants in our data set have each filed a single patent application.

The top applicants vary greatly in size and in their technology focus, from multi-billion dollar companies to universities, publicly funded research organisations and start ups.

FIGURE 14 TECHNOLOGY OVERVIEW OF MACHINE LEARNING PATENT FAMILIES BY AUSTRALIAN APPLICANTS

Source: PATSTAT 2018 Autumn Edition

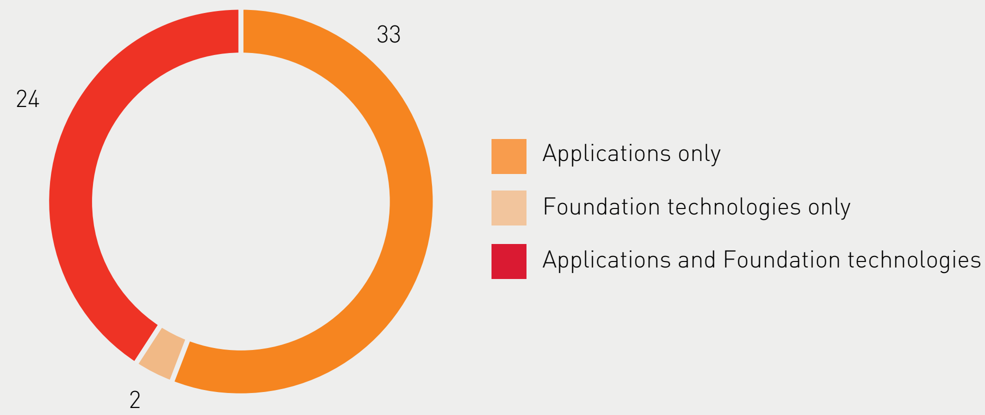

FIGURE 15 PATENT FAMILIES BY AUSTRALIAN APPLICANTS BY CORE CAPABILITIES

Source: PATSTAT 2018 Autumn Edition

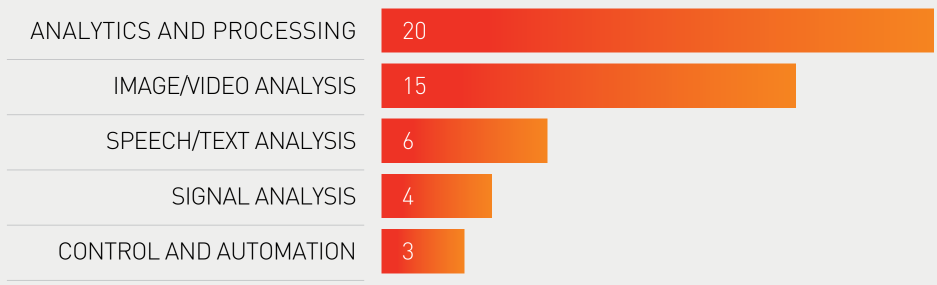

FIGURE 16 PATENT FAMILIES BY AUSTRALIAN APPLICANTS BY REAL WORLD APPLICATIONS

Source: PATSTAT 2018 Autumn Edition

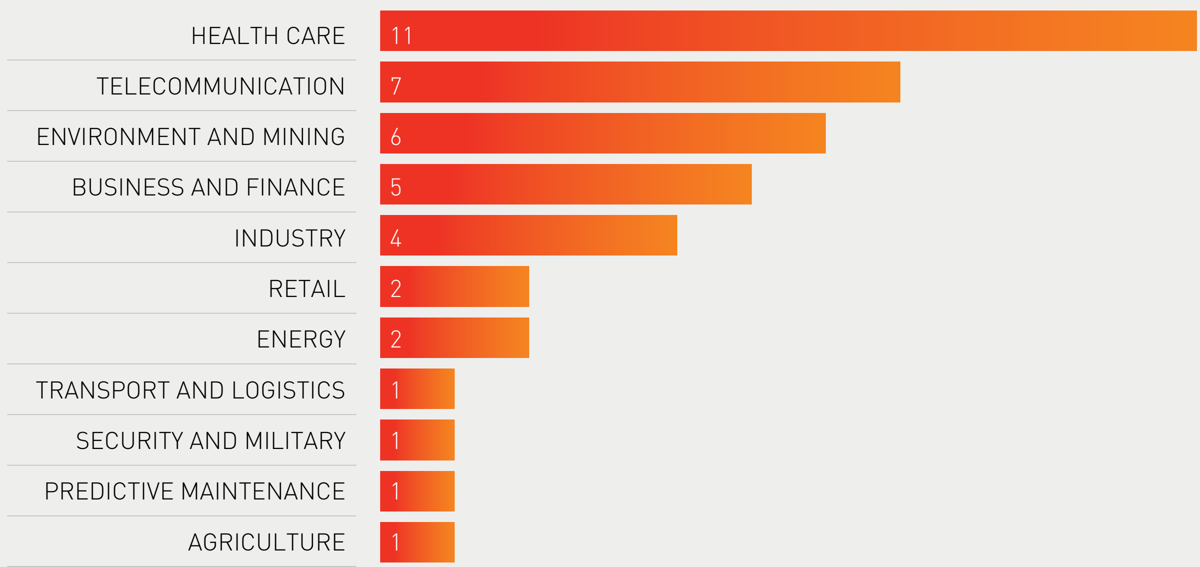

CSIRO

CSIRO is Australia's national science research agency and largest patent holder, with more than 1800 patent applications across all technologies.[92] CSIRO is Australia's largest public research organisation, with around 5 500 employees and funding of approximately AU $600 million in 2017.[93] In November 2018, CSIRO invested AU $19 million in frontier research in AI and machine learning to target AI-driven solutions for areas including health and wellbeing, sustainable energy and resources, and Australian and regional security.[94] CSIRO filed machine learning patent applications related to social media text analysis (WO2017106904, as NICTA) data visualisation, privacy preserving machine learning WO2016061628,as NICTA) image analysis (WO2014201521), medical imaging and diagnostics (WO2016019439) and structural health monitoring (WO2014176625, as NICTA) from 2013–15.

HRO HOLDINGS

Founded in 2015, inventor Jock Lawrie - through Australian company HRO Holdings[95] - has filed three machine learning patent families in 2016, in the areas of patient procedure scheduling and outpatient clinic attendance (WO2018058189, WO2017120646 and AU2017202982).

ATLASSIAN

Founded in 2002, Atlassian is an Australian success story. This software company offers products for software developers, project managers and other business users.[96] Atlassian is the first 'unicorn' start-up company (a start-up valued at more than US $1 billion) to emerge from Australia.[97] The company's products include Loom, an AI-powered log analysis solution that informs the user of any problems within a digital system; Talla, an AI-powered information technology service management assistant; and Jira, a service desk software with automation systems.[98] They filed machine learning patent applications related to managing conversations in messaging applications (US9804752) and the return of user-specific content in searching (US2018189282) in 2016.

[92] CSIRO, Our mission (accessed 9 October 2019) https://www.csiro.au/en/About/We-are-CSIRO

[93] CSIRO 2017-2018 Annual report (accessed 9 October 2019) https://www.csiro.au/~/media/About/AnnualReport/Files/2017-18/18-00410_CORP_AnnualReport2017-18_WEB_180928.pdf?la=en&hash=38A094CABD16B4D934CD44C309B8F95CB700F78E

[94] CSIRO News Release (19 November 2018) https://www.csiro.au/en/News/News-releases/2018/CSIRO-invests-35M-infuture-of-space-and-AI-for-Australia

[95] ASIC Connect search results for HRO Holdings Pty Ltd (accessed 9 October 2019) https://connectonline.asic.gov.au/RegistrySearch/faces/landing/SearchRegisters.jspx?_adf.ctrl-state=xixewdp8p_4

[96] Atlassian, Overview (accessed 9 October 2019) https://www.atlassian.com/software/jira/guides/getting-started/overview

[97] Noyes, K (11 December 2015) Atlassian president on IPO - 'We're overjoyed', https://www.pcworld.idg.com.au/article/590624/atlassian-president-ipo-we-re-overjoyed/

[98] Atlassian, Software (accessed 9 October 2019) https://www.atlassian.com/software

CRC CARE

Founded in 2005, the Cooperative Research Centre for Contamination Assessment and Remediation of the Environment (CRC CARE), performs research, develops technologies and provides policy guidance for assessing, cleaning, and preventing soil, water and air contamination.[99] Their partner organisations include many of the world's largest resource companies such as BHP Billiton, Rio Tinto and Chevron, as well as universities across Australia, government bodies and CSIRO.[100] In 2015, CRC CARE filed machine learning patent applications related to determining the concentration of ions in liquids (WO2016145481) and for determining the presence of volatile organic compounds in samples (WO2016187671).

NEWSOUTH INNOVATIONS

NewSouth Innovations Pty Ltd is the commercialisation arm of the University of New South Wales, one of Australia's largest universities and one of the Group of Eight research intensive universities, with laboratories devoted to computer vision, AI and robotics.[101,102] NewSouth Innovations filed machine learning patent applications related to computer security (WO2014094034) and real time moving image recognition (WO2014194345) in 2012 and 2013, respectively.

UNIVERSITY OF TECHNOLOGY SYDNEY

The University of Technology Sydney established a research centre, the Centre for Artificial Intelligence (CAI), in 2017. CAI is a world leading research centre in AI, with a vision to develop theoretical foundations and advanced algorithms for AI and to drive significant progress in related areas including machine learning, business intelligence, computer vision, data science and social robotics.[103] The University of Technology Sydney filed machine learning patent applications related to traffic lane detection (US9286524) and object tracking (US2016342837), both co-filed with Toyota in 2015.

Overall, this data shows Australia is a minor player in machine learning patenting globally. The picture of Australian patent filings in machine learning is one of diversity, both in the type of applicants, and in the subject matter of their innovation. This highlights a broad range of experience here, and indicates strong potential for development of Australian research in this field.

[99] CRC CARE, About (accessed 9 October 2019) https://www.crccare.com/about

[100] CRC CARE, Partners (accessed 9 October 2019) https://www.crccare.com/about/the-company/partners

[101] UNSW, AI & Robotics Research (accessed 9 October 2019) http://robolab.cse.unsw.edu.au/

[102] UNSW Computer Science and Engineering, Research Activities (accessed 9 October 2019) https://www.engineering.unsw.edu.au/computer-science-engineering/research/research-activities

[103] UTS, Centre for Artificial Intelligence (accessed 9 October 2019) https://www.uts.edu.au/research-and-teaching/our-research/centre-artificial-intelligence

08

Conclusion

This report set out to explore the patent data for innovation relating to machine learning. The analysis demonstrates ongoing commercial development in this sector, with a four-fold global increase in patent filings over five years from 2012–16. Ninety-two per cent of patent families relating to machine learning remain in an active state, indicating ongoing commercial value of these patent applications.

We found 36,740 machine learning patent families filed since 2012, and analysed inventions by machine learning techniques and their applications. Real world applications of machine learning show where machine learning is having an impact in current industries. The telecommunications sector leads development of real world applications of machine learning, with 70 per cent of all patent families in this sector alone. Examples of innovation in telecommunications include internet network operation, measurement of video buffering events, prediction of interference between two or more mobile communication networks and risk assessment for downloading and using mobile applications.

Image and video analysis is the most important machine learning core capability, occurring in 36 per cent of patent families, and the foundational technologies of supervised or semi-supervised learning techniques are the most popular approach, accounting for 38 per cent of all patent filings in machine learning since 2012.

China is both the largest patent filing destination with 73 per cent (26,758) of machine learning inventions filed here, and the largest source of innovation in machine learning, with 69 per cent (25,319) of patent families having a Chinese applicant. Patent applications that originate in China stay in China - only 4 per cent of applications from Chinese applicants are filed internationally. The largest applicant is the State Grid, which has filed 1018 patent families related to machine learning since 2012.

Following China, the United States, South Korea and Japan are the second, third and fourth top patent destinations and countries of origin, respectively, for patent applications relating to machine learning. The data shows some technological specialisation by country; for example the United States is particularly prominent in methods of speech and text analysis.

Australia is an important international market in this sector, ranking tenth in the world as a patent destination for

PATENT FILINGS

4x GLOBAL INCREASE

FROM 2012-16

machine learning innovation. Australia ranks seventeenth worldwide as a source of innovation in machine learning, with 95 per cent of Australian patent applications relating to the application of machine learning methods. CSIRO is the leading Australian patent filer in machine learning.

Recent publications by WIPO and the UKIPO give a broader overview of technology trends across the whole field of AI. The WIPO report emphasises the dominance of machine learning techniques, and the major role of China, and also shows that patent filings are becoming increasingly important in this sector relative to scientific publications.

In the field of machine learning, it is also important to be aware of difficulties in patenting computer algorithms. Patentability of computer-implemented inventions varies across jurisdictions by definition and interpretation of subject matter eligibility. The WIPO report helpfully sets out an analysis of subject matter eligibility in the United States, Europe, Japan, China and Korea, and the definition of patentable subject matter in Australia involves similar considerations. While this legal hurdle might be expected to contribute to patent applications lapsing before grant, the data in this study shows that only eight per cent of machine learning patent families are not in an active state, which indicates a strong global interest in pursuing patent protection – and commercialisation - in this sector.

The WIPO report also includes further analysis on market trends related to AI, including company mergers and acquisitions, patent oppositions, patent litigation, public funding, and investment in start ups.

As well as this, the report includes contributions from experts in AI, data, IP, policy, and innovation, adding context to the information revealed from the patent data, and addressing issues such as existing and potential uses and impact of AI technology, legal and regulatory questions, data protection and ethical concerns.

This patent analytics report on machine learning shows a sector with strong growth and potential to boost Australian innovation and leverage global investment and growth for the benefit of the broader Australian economy. This report is a catalyst to consider the future impact of AI and inform policy and decisions to ensure that Australians benefit from machine learning.

09 Appendix A: Definitions

Patent, applications and publications

A patent is a right that is granted for any device, substance, method or process that is new, inventive and useful. Australian patent rights are legally enforceable and give the owner, or patentee, exclusive rights to commercially exploit the invention in Australia for a period of up to 20 years. In this report, an application refers to a single patent filing. A patent application is usually published within 18 months of its earliest filing date (also known as the priority date). We consider that the priority date is most relevant to our analysis as it is the closest date to that when the invention occurred.

There are two major routes for filing a patent application: the international route and direct filing. The international route involves filing a PCT application, which establishes a filing date in all 152 contracting states. Subsequent prosecution at national patent offices, referred to as national phase entry, is made at the discretion of the applicant. A patent can only be enforced once it has been granted and a PCT application must enter the national phase in each country to proceed toward grant. Alternatively, applications can be filed directly in the countries of interest, without using the PCT system.

Patent families

Applications with the same priority, but filed in different jurisdictions, are known as patent families. Patent families enable us to analyse inventive activity regardless of the number of countries in which protection is sought. Patent families are used in analytics to represent a single invention. We determine patent families based on DOCDB database definition, with a unique family ID for patents that share a common priority document. The number of patent families is typically used as a metric. There are some exceptions when reporting individual applications, as each application represents a legal right in an individual country. When analysing applicants, related commercial entities are grouped by a singly, harmonised name. When individual publication numbers are quoted, we have chosen a representative publication from the patent family, typically US or WO English language documents.

Classification

Patents are hierarchically classified by technology into the hierarchical IPC or CPC systems. The CPC began in 2013 and provides significantly more depth to the hierarchy of the IPC. For more information on the coverage of the CPC, see the CPC Annual Report 2017–18.

Patent status determination

A patent application is considered as 'being sought or in force' when it has not lapsed (due to expiry or non-payment of renewal fees), been revoked or withdrawn. A family is designated as 'being sought or in force' if it contains at least one application that is 'being sought or in force'. The PATSTAT database is used to define the legal events 'Application Discontinuation' and 'IP Right Cessation'. These are considered to make an application lapsed, expired or withdrawn. Applications without these events are considered as 'being sought or in force'.

Appendix A: Definitions

Country of origin

Many records in the PATSTAT database (around 50 per cent) do not have any associated country code information for applicants/inventors. In previous patent analytics studies, null data has generally been excluded from the country of origin analysis.

In this patent analytics report, we have used other data available in PATSTAT to provide additional insights for patent families that do not have applicant origin data available. This method of country code assignment leverages all data from the entire PATSTAT database, and not only country codes information from a subset of the data extracted for a specific technology area.

The country of origin of a patent family is assigned using a three-step process.

1. Where country code is available for applicant data, this is used as the country of origin.
2. If the country code is 'null' and the applicant name is associated with a standardised or cleaned version of their name (e.g. the PATSTAT standardised PSN name), there may be several versions of an associated applicant name with several entries for country code. A count of the total number of applications applied for with each country code is used to determine the most representative country and that is applied. This is an independently developed technique analogous to the technique detailed by the European Commission.
3. If the country code for the standardised name is 'null' or where no applicant details are recorded for the patent application, then the first application authority is used. This is the country the earliest priority document was filed into for the patent family and it is used as a country of origin. This authority is the first by date where the authority is a country and not an international body. For international applications, the receiving office is used where possible.

It is important to note that as patent families are used, most families that are filed internationally have reliable country code data, either through WIPO, EPO, USPTO or other large patent offices. Applicants that often file internationally will have standardised name-associated country information for the same reason, even for local applications. The first application authority is therefore generally only used as a proxy for origin in the case of patent families that are filed to one jurisdiction only, increasing the probability that the country is correctly assigned.

Appendix B: Search Strategy

Searching patent information to identify relevant records for analysis requires a stepped approach to identify broad categories of relevance, and then specific records within them that meet the technology brief.

The following details outline the search and analysis process conducted.

Data extraction and analysis

We used five phases of data extraction and analysis.

- Phase 1: Development of a search strategy (below).
- Phase 2: Data mining using the database PATSTAT 2018 Autumn edition. The unique DOCDB family members relating to machine learning were identified and used as the basis of the analysis.
- Phase 3: Data cleaning, focusing on removing errors and harmonising applicant names and ensuring the return of correct records. This was done in Derwent Data Analyser version 7.
- Phase 4: Data categorisation according to the technological focus of the patent families. The technological focus was determined using the CPC and IPC symbols.
- Phase 5: Data analysis using Tableau 2018.3 for calculation and visual presentation of patent metrics.

Search strategy

The search was limited to patent applications that had an earliest priority date between 2012 18 and a CPC symbol, IPC symbol, keywords (in the title or abstract of the patent) or a combination of both symbol and keyword as listed here. ("%" indicates unlimited truncation.)

Classification Symbols (where a truncation is given, all elements in the hierarchy beginning with that stem are included).

1. IPC Symbols:
   G06N3/02, G06N3/04, G06N3/06%, G06N3/08, G06T1/40
2. CPC Symbols:
   G06N3/02, G06N3/04%, G06N3/06%, G06N3/08%, G06N3/10, G06N3/105, G06N7/046, G06N 99/005, G06T2207/20081, G06T2207/20084, G06T3/4046, G06T9/002, G05B13/027, G05B13/0285, G05B13/029, G10L17/18, G01N2201/1296, G01N29/4481, G01N33/0034, G01S7/417, G06F11/1476, G06F11/2263, G06F2207/4824, G06N7/046, G11B20/10518, G10H2250/311, H02P21/0014, H02P23/0018, H03H2017/0208, H03H2222/04, H04L2012/5686, H04L2025/03464, H04L2025/03554, H04L25/0254, H04L25/03165, H04L45/08, H04N21/4662, H04N21/4663, H04N21/4665, H04N21/4666, H04Q2213/13343, H04Q2213/343, B60G2600/1878, B23K31/006, B29C2945/76979, B29C66/965, Y10S128/925, F02D41/1405, F05B2270/709, F16H2061/0084, F05D2270/709, G10K2210/3038, G10L25/30, G10L15/16

Keywords:

An exact keyword search involves tokenising the field in question and then finding the exact word in the list. In the majority of cases unlimited left and right truncation was used. This has been summarised for clarity, but it means that the words in the order provided were found anywhere in the field in question in exactly the order shown with any number of characters on either side,

Appendix B: Search Strategy

including prefixes or suffixes. An underscore represents exactly one wildcard character.

1. Searched exactly:
   convnet, convnets, LSTM, LPboost, Totalboost, brownboost, madaboost, Logitboost
2. Searched with unlimited left and right truncation:
   neural net, multi_layer perceptron, multi_layer perceptron, machine learning, training algorithm, deep learning, deep belief net, deep_belief net, generative adversarial net, long short term memory, long_short_term memory, multi label classif, multi_label classif, random forest, reinforcement learning, semi supervised learning, semi_supervised learning, supervised learning, unsupervised learning, similarity learning, dictionary learning, Hebbian learning, feature learning, auto encoder, auto_encoder, support vector regression, support vector machine, rule learning, rule based learning, rule_based learning, decision tree, classification tree, regression tree, swarm intelligence, swarm optimisation, swarm optimization, swarm behavio_r, kernel learning, k-nearest neighbo_r, k-mean, hierarchical cluster, naïve Bayes, naive Bayes, Bayesian net, Bayesian learning, mixture model, mixture_model, Blind signal separation, independent component analysis, self_organising map, self organising map, Markov% decision process, natural language processing, chatbot, Virtual assistant, probabilistic graphical model, Instance_based learning, Instance based learning, multi_task learning, multi task learning, statistical relational learning, transfer learning, ensemble learning, gradient boosting, gradient tree boosting, xgboost, adaboost, adaptive boost, adaptive_boost, rankboost, boosting algorithm, Bootstrap aggregat, stacked generali_ation, stacked_generali_ation, logistic regression, stochastic gradient descent, latent semantic analysis, latent dirichlet allocation, meta learning, Q-learning, temporal difference learning, learning automata, learning_automata, extreme_learning_machine, extreme learning machine, logic learning machine, single_linkage clustering, single linkage clustering, conceptual clustering, multilayer perceptron, deepbelief net, longshort_term memory, long_shortterm memory, longshortterm memory, multilabel classif, semisupervised learning, autoencoder, rulebased learning, swarm behavior, k-nearest neighbor, mixturemodel, selforganising map, self organising map, selforganizing map, self organizing map, Instancebased learning, multitask learning, adaptiveboost, learningautomata, extremelearning_machine, extreme_learningmachine, extremelearningmachine, extremelearning machine, extreme learningmachine, singlelinkage clustering
3. Searched in a combination of exact (without % wildcards) and with unlimited left and right truncation, as listed:
   (inference & %learn%), (inference & %train%)

Appendix C: Technology Analysis

Patent families relating to machine learning were first categorised into two groups, machine learning techniques and machine learning applications, using IPC/CPC symbols and keywords. This classification was done in Tableau Desktop.

Multiple classification symbols may be present for a single patent family and in categorising patent families, overlap where families could be placed in multiple sub-categorises were noted. Sub-categories were prioritised as per the order of the tables below to ensure a patent family appears only in one sub-category. Prioritisation was not done between machine learning foundational technologies and machine learning applications as one patent family could inherently discloses both the foundational technologies and their applications.

Machine learning foundational technologies categorisation

The machine learning techniques category includes patent families that fall into algorithms and architecture.

Algorithms

Algorithms further includes supervised and semi-supervised learning, unsupervised learning, reinforcement learning, niche learning techniques and general techniques sub-categories as per the priority order shown below.

PRIORITY	SUB-CATEGORY	COMPONENTS	IPC	CPC	KEYWORDS
1	Reinforcement learning	[Dynamic Programming]			OR (dynamic programming, markov decision process%)
		[Learning Automata]			(learning_automat%)
		[Reinforcement Neural Networks]			OR ((reinforce% AND (neural_ network%)), deep reinforce%, deepmind, deep q network%, DQN, DDPG, deep deterministic policy gradient%)
		[Other Reinforcement learning]			OR (reinforce% learn%, q_learn%, sarsa) AND NOT ([Dynamic Programming] OR [Learning Automata] OR [Reinforcement Neural Networks])

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	COMPONENTS	IPC	CPC	KEYWORDS
2	Supervised semi-supervised learning	[Boosting]			OR (boosting, meta_algorithm, adaboost, lpboost, totalboost, brownboost, xgboost, logitboost, adaptive_boost, rankboost)
		[Classification and Regression]			OR (decision tree, random forest, classif% tree, svm, support vector machine%, k_NN, nearest_neighbo[, u]r, similarity learning, naïve bayes, bayes% network%, regression, linear discriminant analysis, gaussian process%, bayes% learn, binary_ classification%, multi_label_ classification, kernel, classification model%, classifier%)
		[Supervised Dictionary]			(supervi[s,z]% AND (dictionary learning))
		[Supervised Neural Networks]		G06N3/084	OR (supervised neural net, convolution% neural net, inferred function, convnet, CNN, error signal, label% training data, label% data, supervisory signal%, supervised teach%, multi-layer perceptron%, back-prop%, radial basis function%, RBF)
		[Supervised Learning-General]			((supervis%, semi_supervis%, weekly_supervis%) AND NOT ([Supervised Neural Networks] OR [Classification and Regression] OR [Boosting] OR [Supervised Dictionary]))
3	Unsupervised learning	[Blind Signal Separation]			OR (blind signal separation%, principle component analysis, independent component analysis, matrix factori[s,z]ation%)
		[Cluster Analysis]			OR (cluster%, hierarchical, k_ mean%, mixture model%, dbscan, optics algorithm%)
		[Dictionary Learning]			(dictionary learning) AND NOT [Supervised Dictionary]

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	COMPONENTS	IPC	CPC	KEYWORDS
		[Unsupervised Neural Networks]		G06N3/088	OR (non_supervised neural_net, unsupervised neural_net, auto_encoder, unlabel_ed data, deep belief net, hebbian learning, self_organi[s,z]ing map, generative adversarial network, competitive_ learn)
4	Niche learning techniques	[Transfer Learning]			OR (transfer learn, inductive learn, inductive transfer, transductive transfer)
		[Rule-Based Learning]			OR (rule_based_machine_learning, RBML)
		[Relational Learning]			(relation% learn%)
		[Particle Swarm Optimisation]			(particle swarm, swarm intelligence, swarm optimi%, swarm behavi%)
		[Online learning]			OR (online learn%, online machine learn%)
		[Graphical Model]			(graphical model%)
		[Genetic Algorithms]	G06N3/12%	G06N3/086, G06N3/12	(genetic algorithm%)
		[Gaussian Mixture Model]			(mixture model%)
		[Feature Learning]			(feature learn%)
		[Extreme learning]			OR (extreme learn%, ELM)
		[Ensemble learning]			(ensemble learn%)
		[BCM Rule learning]			OR (bcm rule, ibcm rule, bcm learning%)
		[Ant Colony Optimisation]			(ant_colony)

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	COMPONENTS	IPC	CPC	KEYWORDS
[Learning Paradigms] = [Supervised Learning] or [Unsupervised Learning] or [Reinforcement Learning] or [Niche Learning]					
8	General techniques	[Machine-Learning-general]	G06N99%, G06N3/02%, G06N3/10%, G06N3/08%	G06N3/10%, G06N3/08%, G01N33/0034%, H04L45/08%, G06N3/105%, G06F11/1476%, G06F11/2263%, G06F2207/4824%, H04L2012/5686%, H04L45/08%, H04Q2213/13343%, H04Q2213/343%, B60G2600/1878%, B23K31/006%	(OR (machine_learn%, deep neural net%, deep learn%, deep belief net%, DNN, neural_net%. Learning_machine%, training algorithm%) AND NOT ([Learning Paradigms] OR [NN Architecture])
		[Gradient Descent]			OR (stochastic gradient descent%, regulari[s,z]ation%)
	[General techniques] = ([Machine-Learning-general] OR [Gradient Descent]) AND NOT [Applications]				

% is a wildcard for zero or more characters
_ is wildcard for zero or one character

Architecture

The architecture sub-category further includes neural network architecture, architecture optimisation and physical realisations of neural networks as per the priority order shown below.

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
5	Architecture optimisation		G06N3/082	
6	Neural network architecture	G06N3/04%	G06N3/04%	
7	Physical realisations	G06N3/06%	G06N3/06%	(OR (FPGA, fixed point, hardware implement%, hardware reali[s,z]) AND (neural net%)) OR (graphic% process% unit%)

% is a wildcard for zero or more characters
_ is wildcard for zero or one character

Appendix C: Technology Analysis

Machine learning applications categorisation

The machine learning applications are categorised by core capabilities and/or real world applications.

Core capabilities

The machine learning applications category includes patent families that fall into speech and text analysis, image and video analysis, control, optimisation and automation, signal analysis and data analytics and processing as per the priority order shown below.

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
1	Speech and text analysis	G10L%, G06F17/2	G10L%, G06F17/2	OR (natural language process%, language model%, natural language understand%, NLP, NLU, word embedding, tokeniz%, semantic pars%, entity extraction%, entity name extraction%, entity resolution%, part_of_speech, POS, LSTM, long_short_term_memory, attention model%, pronunciat%, speech%, audio process%, chat_bot, virtual assist%)
2	Image and video analysis	G06T%, G06K9%	G06T%, G06K9%	OR (image analy%, image classif%, image process%, image compar%, image mat%, object recognition, object classif%, object match%, target recognition%, target classif%, target match%, spectral imag%, super resolution%, depth perception%, pose estimate%, captioning, adagrad%, adam gradient%, (OR (audio, media) AND OR(feature extract%, feature analy%, feature classif%)))
3	Control, optimisation and automation	B25J9%, B21B37%, B29C64/393%, B33Y50/02%, G05B%, G05D%, H03L%, H02P%, F24F11%	B25J9%, B21B37%, B29C64/393%, B33Y50/02%, G05B%, G05D%, F01N2900/0402%, F02D41/2438%, F02D41/2441%, F02D41/1405%, F03D7/046%, F05B2270/709%, H03L%, H02P%, F24F11%, B29C2945/76979%, B29C66/965%, F16H2061/0084%	OR (automatic control%, control system%, control method%)

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
4	Signal analysis	A61B5%, A61B8%, G01N21%, G01N22%, G01N27/90%, G01R%	A61B5%, A61B8%, G01N21%, G01N22%, G01N27/90%, G01R%, G01N29/4481%, G01N33/0034%, G10H2250/311%, G01S7/417%, H04L25/03165%, H04L25/0254%, H04L2025/03464%,H04L2025/03554%, G11B20/10518%, H03H2017/0208%, H03H2222/04%, G10K2210/3038%	OR (signal% AND OR (process%, analy[s,z]%))
5	Data analytics and processing	G06Q%, G06F3%, G06F5%, G06F7%, G06F16%, G06F17/14%, G06F17/15%, G06F17/16%, G06F17/30%, G06F17/40%, H04N21/466%	G06Q%, G06F3%, G06F5%, G06F7%, G06F16%, G06F17/14%, G06F17/15%, G06F17/16%, G06F17/30%, G06F17/40%, G06F2216%, Y10S706%, H04N21/466%	OR (data mining, data analy%, analy% data, cognitive, data filter%, filter% data)

% is a wildcard for zero or more characters
_ is wildcard for zero or one character

Real world applications

The machine learning applications category includes patent families that fall into predictive maintenance, telecommunication, agriculture, energy, environment and mining, security and military, retail, health care, industry and manufacturing, transport and logistics and, business and finance as per the priority order shown below.

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
1	predictive maintenance	G05B23/0283%	G05B23/0283%	OR ((predict% AND maintenance), (OR (fault%, wear%, tear%, acoustic%, failure%, defect) AND OR (diagnos%, detect%, predict%, identif%, calculate%, monitor%))) OR (diagnos% and current%, damage% and identif%)

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
2	telecommunication	H04B%, H04H%, H04J%, H04K%, H04L%, H04M%, H04N%, H04Q%, H04R%, H04S%, H04T%, H04W%	H04B%, H04H%, H04J%, H04K%, H04L%, H04M%, H04N%, H04Q%, H04R%, H04S%, H04T%, H04W%	OR (telecom%, celllular network%, mobile network%, computer network%, internet security, data security, LTE, LTE-A, lte_advance%, malware%, spam%, IOT%, social network%, communication network%, positioning, cloud service%, network system%, server system%, audio coding, video coding, network resource%, communications network%) OR (webpage, coaxial_cable, communicaiton, mimo, multiple_input_multiple_output, optic__fiber, optic__fibre, radar)
3	agriculture	A01%	A01%, G05D2201/0201%, Y02A40/1%, Y02A40/2%, Y02A40%, G06T2207/30188%	OR (plant cultivat%, farming, irrigate%, agriculture%, crop, grain, live_stock%, vegetation%, honey) AND NOT (honeycomb%) OR (fruit, maturity, milk, cotton, plant, grapefruit, beef, brooder, meat, garlic, wheat, tobacco, animal feed, fish, sorghum, grain, algae, tea, microalgae, peanut%, protein, food, ferment, gelatin, liquor, vintage, drink, hogwash, gutter, oil, pork, apples, shrimp, chicken, sausage%, citrus, rice, maize, chlorophyll, eggs, baijiu, cooking apparatus, straw, sugar, wine, pear tree, cooking appliance, red date, cyclotella, baking method, cooking utensil, pest, vaporariorum, quarantine, larvae, pesticide)
4	energy	F01%, F02%, F03%, F04%, F15%, F17%, H01%, H02%, H03%, G06Q50/06%, G01R21%	F01%, F02%, F03%, F04%, F15%, F17%, H01%, H02%, H03%, Y10S706/915%, Y10S706/907%, Y10S60/00%, G06Q50/06%, G01R21%	OR (voltage%, power consumption%, power distribut%, power generat%, power transmit%, electric%, transformer%, power station%, power plant%, harmonic emission%, partial discharge%, battery, multimeter%, power grid%, energy management%, alternating_current%, solar, wind power, energy saving%, microgrid|transmission line%,transmission wire%, wind farm%, photovoltaic%, line_to_ground%) OR (OR(power,energy) AND OR(load predict%, output predict%, load forecast%)) OR (gyro, diode, fpga, circuit board, semiconductor, grounding grid, electromagnetic, boiler, analog circuit, analogue circuit, coaxial cable, integrated circuit, electronic circuit, line loss, energy billing, 10kv line, resistive, energy network, power system, microwave filter)

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
5	environment and mining	E21%, G01W%, G01V%, G01C11%, G01C13%, G01C15%	E21%, G05D2201/021%, Y10S706/928%, Y10S706/929%, Y10S706/93%, Y10S37/00%, G01W%, G06T2207/30192%, G01V%, G01C11%, G01C13%, G01C15%	OR (excavat%, geology%, seismology%, weather%, meteorology, seismic%, earthquake, (OR (coal%, opencast%, mineral%, fossil%, oil, gas, LNG, gold) AND (mine%, mining, drill%, mill%)), (earth% and drill%)) OR (oil_gas_water, landslip, forest fire, soil, oxygen, housing, biomass furnace, fire alarm, garbage calorific, wind, typhoon, thunder, seabed, sea surface, fire flame, environment, water level, water tank, algal bloom, fire disaster, fire-hazard, trash content, gask leak, water quality, smog, rock, well line, gis, atmospheric, ground feature, soil, river, remote sensing, coal, forest, landmark, aerial photograph, areally photograph, oil_producing well, oil well, oil overflow, air_quality, geomagnetic, plankton, pressure vessel, terrain, aerial scene, polari[sz]ed sar, sar image, sediment, bird population, pollutant, heavy metal, natural gas, ore body, water analysis, flame, fire AND (detect, chemical) AND (contamination))
6	security and military			OR (forgery, copy protection, public sentiment, public monitor, intrusion, alarming, surveillance, counterfeit, ground target, camouflaged, trojan, detecting attack, fire control, authentication, spoofing, gun detection, war, crime, missile, submarine, violent behaviou?r, sql injection, threat analysis) OR (face, facial, finger_print, iris, voiceprint, speaker recognition, speaker age, handwriting, handwritten, hand recognition, palm_print, palm detection, hand_print)
7	retail	G06Q30%, G06Q50/28%	G06Q30%, G06Q50/28%	OR (product% AND OR (recommend%, identif%, personali%), demand predict%, predict% demand, customer engage%, customer satisf%, personali% market%, personali% adverti%, telemarket%, water meter%, amazon_go) OR (e_commerce, electronic_commerce, e_shopping, electronic_shopping, internet_commerce, internet_shopping, web_commerce, web_shopping, shopping mall, clothing, textile, fabric, clothes)

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
8	health care	A61%, G16B%, G16H%, G06Q50/24%, G06F19%, C12M1%	A61%, G16B%, G16H%, G06Q50/24%, G06F19%, G01N2800%, Y10S706/924%, Y10S128%, C12M1%	OR (medicin%, medical%, biology, biological, patient%, health_care, prosthesis, neural oscillator%, tumour%, tumor%, disease%, surgery, surgical, clinical, cancer, MRI, imag%, genetic toxicity, IVF, cardiac%, peptide, influenza, bacteria) OR (urine, endoscopic, protein process, blood cell, visual cortex, x-ray, emotion, cleavage cell, mammogram, vivo, electromyography, gastroscope, eye fundus, developmental disorder, nuclear magnetic resonance, magnetic resonance imag%, tongue, ecg, stem cell, marrow, biomarker, patholog%, cell outline, cell boundary, vein, vascular, bone, cilia, pelvic, neurotox%, thyroid, health management, lung, plasmid, retina, blood, cerebral, vessel, artery, fmri, mammary, adipose, abdominal, lesion, pulmonary, cervical, tooth, humoral, breast, tissue, fall, tumble AND (detection, monitor, brain machine interface, brain computer interface, encephalogram, eeg))
9	industry and manufacturing	B01%, B02%, B03%, B04%, B05%, B06%, B07%, B08%, B09%, B21%, B22%, B23%, B24%, B25%, B26%, B27%, B28%, B29%, B30%, B31%, B32%, B33%, B65%, B66%, B67%, B68%, C02%, C03%, C04%, C05%, C10%, C13%, C21%, C22%, C23%, D06%, D21%, F01%, F02%, F03%, F04%, F05%, G06Q50/04	B01%, B02%, B03%, B04%, B05%, B06%, B07%, B08%, B09%, B21%, B22%, B23%, B24%, B25%, B26%, B27%, B28%, B29%, B30%, B31%, B32%, B33%, B65%, B66%, B67%, B68%, C02%, C03%, C04%, C05%, C10%, C13%, C21%, C22%, C23%, D06%, D21%, F01%, F02%, F03%, F04%, F05%, F16%, F22%, F23%, F24%, F25%, F27%, G05D2201/0201%, G05D2201/0216%, G06Q50/04	(manufactur%, process % tech%, galvaniz%, aluminium cast%, copper concentr%)

Appendix C: Technology Analysis

PRIORITY	SUB-CATEGORY	IPC	CPC	KEYWORDS
10	transport and logistics	B60%, B61%, B62%, B63%, B64%, G05D1%, G08G%, G06Q10/08%, G06Q50/28%, G06Q50/30%	B60%, B61%, B62%, B63%, B64%, G05D1%, G05D2201/0212%, G05D2201/0213%, G06Q2240/00%, G08G%, Y10S706/905%, G06Q10/08%, G06Q50/28%, G06Q50/30%, G06T2207/30181%	OR (autonomous vehicle%, autonomous car%, driverless, driver assist%, traffic%, supply chain%, automatic driving, auto driving, autonomous driving) OR (driver, pedestrian, car, road sign, license plate, steel rail, taxi, postal, transport, driving, truck, plate number, vehicle color, chassis, ship, vehicle, vehicular, high_speed, rail, railway, road, aerocraft, airplane, aeroplane, aircraft, rotorcraft, navigation, automobile, flight control, uav, locomotive, helicopter, cockpit, bus, busses, tire, aerospace, lane AND (recogni[sz]%, detect%, recognit%, elevator, escalator) AND (visitor, people, passenger)) OR (satellite, spaceflight, spacecraft) OR ((OR (earth,planet) AND OR (observ%, monitor%)), (space_station%)
11	business and finance	G06Q10/02%, G06Q10/04%, G06Q10/06%, G06Q10/10%, G06Q20%, G06Q40%, G09B%, G06Q50/20, A63%	G06Q10/02%, G06Q10/04%, G06Q10/06%, G06Q10/10%, G06Q20%, G06Q40%, Y10S706/925%, Y10S706/926%, G09B%, G06Q50/20%, Y10S706/927%, A63%	OR (trading,robo advisor%, finance, financial, transaction%, banking, audit, (OR (risk%,credit%,portfolio, workforce) AND (analysis, assessment%, approv%, management%, planning)), (OR (currency,money,stock%) AND (exchang%, laundering)) OR (advertisement filtering, movie expectation, house property, merchandise, insurance, stock news, investor) OR (education%, teaching%, guided learning%) OR (pronunciation, student, classroom) OR (video game%, board game%, sport%, ball game%, skate%, skating%, card game%, indoor game%, outdoor game%, bowling, racing, riding) OR (sport, augmented reality, head_ mounted_ display, virtual reality, badminton, basketball, online game, computer game)"

Abbreviations

ACS	Australian Computer Society
AI	Artificial intelligence
BoF	Bag of features
CAI	Centre for Artificial Intelligence
CNN	Convolutional neural networks
CSIRO	Commonwealth Scientific and Industrial Research Organisation
CSPG	China Southern Power Grid
CT	Computed tomography
DFKI	Deutsches Forschungszentrum für Künstliche Intelligenz (German Research Center for Artificial Intelligence)
FoG	Freezing of Gait
GAN	Generative adversarial network
GIS	Gas insulated switchgear
IBM	International Business Machines
IP	Intellectual property
IPO	Intellectual Property Office (UK)
LSTM	Long-short-term memory
MIT	Massachusetts Institute of Technology
MRI	Magnetic resonance imaging
NLP	Natural language processing
NICTA	National ICT Australia
NTT	Nippon Telegraph and Telephone Corporation
ODN	Object detection network
PCT	Patent Cooperation Treaty
RNN	Recurrent neural network
RPN	Region proposal network
SMILE	Statistical Machine Intelligence and LEarning (UESTC)
State Grid	State Grid Corporation of China
UESTC	University of Electronic Science and Technology of China
UKIPO	United Kingdom Intellectual Property Office
WIPO	World Intellectual Property Office

About ACS

ACS is the professional association for Australia's Information and Communication Technology (ICT) sector. More than 44,000 ACS members work in business, education, government and the community. ACS exists to create the environment and provide the opportunities for members and partners to succeed. ACS strives for ICT professionals to be recognised as drivers of innovation in our society, relevant across all sectors, and to promote the formulation of effective policies on ICT and related matters.

Visit **acs.org.au** for more information.

ACS
International Tower One
Level 27, 100 Barangaroo Avenue
Sydney NSW 2000

P: 02 9299 3666
F: 02 9299 3997
E: info@acs.org.au
W: acs.org.au