Australia's National
Science Agency

Artificial Intelligence foundation models

Industry enablement, productivity growth, policy levers and sovereign capability considerations for Australia

March 2024

About this report

The report was prepared by CSIRO to inform industry, government and community organisations about the strategic implications of AI foundation models for Australia.

Citation

Hajkowicz SA (2024) Artificial intelligence foundation models: Industry enablement, productivity growth, policy levers and sovereign capability considerations for Australia. CSIRO, Canberra.

Acknowledgements

The author would like to thank the many experts and stakeholders we interviewed from Australian and overseas industry, government and academic institutes. This includes our expert committee who participated in two meetings to help us frame the questions and review early drafts. Gratitude is also expressed to staff within CSIRO who provided review and feedback on early draft versions of this report.

Copyright

Disclaimer

Accessibility

CSIRO is committed to providing web accessible content wherever possible. If you are having difficulties with accessing this document please contact csiro.au/contact

Report images – sources and credits

Cover image: A generative-AI created image showing the foundations of a large, impressive skyscraper created by the author using natural language prompts in ChatGPT with DALL-E 3. This picture is a metaphor for AI foundation models which support wide-ranging social and economic activity.

Three images in the report were made with ChatGPT (GPT4) via text prompts (as noted). Other images are from Shutterstock (with artist credits for Maxx-Studio, Ovchinnkov Vladimir, Bany's beautiful art).

Contents

What did AI have to say? 2
Executive summary 4
Why did we write this report? Why now? 6
What are AI foundation models? 7
How are AI foundation models made? 8
Working with existing models 9
A paradigm shift? 10
A surge of activity 11
Computing infrastructure needs 12
Productivity and economic impacts 14
Global policy responses 15
Australian AI product innovation 18
Sovereign capability considerations 19
Policy levers 22
A healthcare example 26
A weather and climate forecasting example 27
Other examples 28
Conclusion 29
References 30

What did AI have to say?

A conversation with ChatGPT (GPT-4)

Please keep responses to these questions below 30 words.

Certainly, I'll keep my responses concise and to the point. What would you like to know?

What are AI foundation models?

AI foundation models are large-scale machine learning models trained on diverse data. They serve as a base for specialized applications. Examples include GPT-3, BERT, and DALL-E.

What makes them different to other AI models?

Foundation models are more versatile, trained on extensive data, and can be fine-tuned for various tasks. They offer a general-purpose base, unlike specialized, task-specific models.

How will they impact the economy?

Foundation models can drive innovation, automate tasks, and improve decision-making. However, they may also disrupt job markets and exacerbate inequality. Economic impact is multifaceted.

Will they boost productivity and improve living standards?

Yes, foundation models have the potential to boost productivity by automating tasks and improving efficiency. They could also enhance living standards through better services and healthcare. However, equitable distribution is key.

Are they associated with market failure? What should governments do?

Foundation models could lead to market failures like monopolies or information asymmetry. Governments may need to regulate, ensure fair competition, and address ethical concerns for equitable benefits.

Given they're a new and disruptive technology what should Australian industry be doing about AI foundation models?

Australian industry should invest in AI research, collaborate with academic institutions, and adopt foundation models in key sectors like healthcare and energy. Regulatory preparedness and ethical considerations are also crucial.

What about government?

The Australian government should focus on regulation, fund AI research, and promote ethical AI practices. Collaboration with industry and academia is essential for maximizing benefits and minimizing risks.

What about universities and research institutes?

Universities and research institutes should focus on advanced AI research, ethical considerations, and workforce training. Partnerships with industry and government can facilitate real-world applications and policy development.

Should Australia make its own ChatGPT?

Creating an Australian ChatGPT could support local language and cultural nuances and contribute to national AI capabilities. However, it requires significant investment in data, computing power, and expertise.

What overall advice do you have for Australia when it comes to AI foundation models? The response to this question can be up to 60 words.

Australia should adopt a multi-stakeholder approach involving government, industry, and academia to harness AI foundation models. Prioritize investment in research, infrastructure, and skill development. Establish robust ethical and regulatory frameworks to ensure equitable benefits and mitigate risks. Collaboration, both national and international, is key for maximizing the technology's positive impact.

Executive summary

AI foundation models have 4 key ingredients

To make, adapt or fine tune AI foundation models you need these things – they form the basis of most national strategies.

1

Skilled workers – interdisciplinary and advanced capabilities in machine learning, data science, software engineering, mathematics, statistics and social, cultural and economic aspects of AI use.

2

Hardware – AI accelerator chips (graphics processing units) and high-performance computers to run the algorithms in a secure environment.

3

Datasets – for training and fine tuning the models.

4

Policies, awareness, regulation and AI literacy – to achieve safe, secure, beneficial and ethical application.

What are they? – AI foundation models are typically trained on vast troves of data using unsupervised learning at scale and support wide ranging downstream tasks, many of which may not have been envisaged by the developers. They're a powerful and novel development in the AI field. Examples include BERT (Google), GPT-4 (OpenAI), Llama2 (Meta AI). These models are made accessible via tools such as ChatGPT, Google Gemini and Microsoft Copilot. They can be multi-modal working with text, audio, images and video. Some AI foundation models have robotic capabilities such as DeepMind's RoboCat (Bousmalis *et al.*, 2023).

Surge of activity – Since the publication of the paper on transformer architecture (Vaswani *et al.*, 2017) there's been an explosion of AI foundation models. Researchers from Germany identified 125 such models with 73% from the US and 15% from China (AKI, 2023). The ChatGPT tool which is an interface for the large language model GPT-4 was released in November 2022. In 2 months, it had >100 million active monthly users. Google and Microsoft soon released Bing Chat and Bard into the public domain with hundreds of millions of users.

Improved capability – Upgrades and add-ons emerged quickly. OpenAI released GPT-4 March 2023 and Advanced Data Analysis in July 2023. On 6 November 2023 they released the 'Assistants API' that allows people to make their own AI agents. Other technology companies are likely to keep releasing new generative and foundational AI products in the months ahead. There will be more tools. They're going to get better.

Productivity uplift – There's a growing body of evidence that foundation and generative AI (GAI) is associated with productivity uplift. One study found customer support staff resolved complex problems 14% faster when aided with GAI (Brynjolfsson and Raymond, 2023). Two separate studies of software developers using code- writing GAI observed productivity gains of 46% (Tonkin, 2023) and 56% (Kalliamvakou, 2022). Another study found productivity uplift from GAI of 40% for management consultants (Dell'Acqua et al., 2023).

A futuristic sky city (built on solid foundations) as generated by ChatGPT and DALL-E 3 with natural language prompts from the author. Just as building foundations can support wide-ranging social and economic activity in an entire (future) city so too can AI foundation models.

Solow's Paradox – However, productivity gains are not assured. When used to augment tasks beyond its known capabilities, GAI was associated with a 19% increase in errors by management consultants (Dell'Acqua et al., 2023). There's challenges and complexity associated with AI adoption and adaptation. Whilst researchers see productivity gains at task/firm level it's not observable at the economy-level; this is known as Solow's Paradox. It's also likely that AI success stories get publicised more than AI failures.

Sovereign capability and responses – Considerations relate to the ability of nation states to ensure security, reliability, industry competitiveness, worker benefits and cultural compatibility of AI foundation models. The vast bulk of AI foundation models are made and operated by overseas technology corporations. The US, the UK, Germany and Japan have announced initiatives, proposed legislation and whitepapers on AI foundation models. These initiatives centre on the 4 key ingredients of AI foundation models.

Policy levers – Governments seeking to harness the opportunities can consider actions such as developing AI foundation models to improve public sector functions, acquiring and democratising access to high performance computers, sharing datasets, skills uplift and international collaborations. Governments are applying these levers due to emerging observations of shortcomings in the AI foundation model marketplace and the data-security requirements of government (CMA, 2023; Brookings Institute, 2023).

Why did we write this report? Why now?

One of the (many) things we do at CSIRO is technology foresight. This involves scanning global knowledge-networks for emerging technologies that could impact Australia' economy, environment and society. We want to ensure our companies, workers, communities and governments are equipped with the best technologies and knowledge in a globally competitive marketplace. We want to keep Australia ahead of the technology frontier.

Recently our technology foresight radar, operated by our researchers specialising in the field of strategic foresight, lit-up with multiple bright flashing lights indicating something significant was unfolding. The lights were about developments in the field of artificial intelligence (AI). But they came with much uncertainty about potential impacts and future development trajectories. They continue to flash.

At the centre of the action are AI foundation models. These are machine learning models trained on vast swathes of data that can perform wide-ranging tasks at human levels of performance and often better. AI foundation models became much better when a scientific research paper on the transformer architecture was published several years ago. This substantially improved the performance of machine learning and ushered in the era of AI foundation models.

AI foundation models support tools like OpenAI's ChatGPT, Google's Gemini, Microsoft's Copilot, IBM's Watsonx and others. There's been over a hundred AI foundation models released into the world over the past several years. Many are multi-modal which means they can handle inputs and outputs in text, image, audio, video and robotic formats. They can perform wide ranging tasks with general, and ambiguous, instructions. And they're associated with major productivity uplift as they let people do things faster, cheaper, safer and overall better. The task-performance of these models has been stunning.

But they are rarely made in Australia. The vast majority come from the United States and China. Our industry is becoming increasingly dependent on these models. On one hand this is an opportunity. All these great tools are appearing at our fingertips. But there is also risk. Will these models boost the competitive interest of Australian companies and workers? Will they behave according to our norms, cultural and ethical expectations? Can we be adapting, fine-tuning or building our own AI foundation models to support our own industry and workers? If so, how should we go about doing this? What's Australia's best placement in the global value chain for AI foundation models?

There are many more questions. And we are at the beginning of a journey. We wrote this report to help frame the issues and alert Australian industry, community and government to the opportunities and risks associated with AI foundation models. We want to share what we're seeing and explore the sorts of strategies and policies organisations across the globe are implementing. We're at the early stages of this journey but AI foundation models might be Australia's best chance of reversing the productivity slump – something that's happening in most advanced-economies worldwide. It's a general-purpose technology that's useful in every company and every career.

AI has been with us for decades. What's new, and why we wrote this report now, is the emergence of AI foundation models stemming from the transformer architecture. These models are now at our fingertips. And they're widely seen to be an inflection point in the journey of AI. They've changed what's possible and they've changed the way the field of AI will develop into the future. We hope this report will help our colleagues in industry, government and community organisations understand the forthcoming world of AI foundation models and chart a pathway through the coming decade of AI-fuelled technology disruption.

Elanor Huntington
Director of Digital, National Facilities & Collections, CSIRO

What are AI foundation models?

The rise of foundation models is seen by many as a paradigm shift in the field of AI. But they're not completely new. They use machine learning which has been around for decades. What is new is the size of datasets they're trained on, the wide-ranging tasks they perform and their overall task performance and multi-modality which is the ability to handle text, images, video, audio and robotic movements as inputs and/or outputs.

Powerful and adaptable for diverse downstream tasks

Whilst most AI models perform a specific and well-defined task AI foundation models can handle generalised instructions and perform wide ranging downstream tasks oftentimes not envisaged by the model developers. A team of researchers at the Stanford University Center for Research on Foundation Models define an AI foundation model as 'any model that is trained on broad data (generally using self-supervision at scale) that can be adapted (e.g., fine-tuned) to a wide range of downstream tasks' (Bommasani *et al.*, 2023). Self-supervision means the model 'learns' from unlabelled data and makes predictions based on pattern recognition.

Generative AI is a related (but distinct) concept to foundation AI, and many generative AI models are also foundation AI models. A generative AI model uses machine learning trained on (typically human generated content) to generate realistic text, audio, imagery or video.

Examples of AI foundation models

Over the past few years there's been an explosion of AI foundation models. Most of the world's large technology corporations have developed foundation models along with many startups and government agencies. Some examples include GPT-4 (OpenAI), BERT (Google), Llama2 (Meta) and the Watsonx.ai geospatial model (NASA/IBM). These models are often made available to the public via chatbots and other user interfaces such as ChatGPT or Microsoft Copilot. Some are commercially available others are provided free of charge in open access.

Characteristics of AI foundation models

- Typically, pre-trained on vast troves of data using hundreds of millions, billions or trillions of input parameters. However, it's worth noting that some foundation AI models can be pre-trained on smaller well curated datasets.
- Typically use deep learning approaches which are a type of machine learning inspired by the structure and function of neural networks in the human brain.
- Typically use unsupervised machine learning with unlabelled data. However, not always. ImageNet and GLUE are examples of labelled datasets for AI foundation models.
- May have human-intuitive user interfaces via text, image, audio and video modes and can also be multi- modal where text is used to generate images (or videos, audio) and vice versa.
- Enable transfer learning so that the skills and knowledge developed in the pre-training phase can be used for smaller datasets and task-specific problem solving reducing the requirement for (time consuming and costly) labelled data in the downstream application.
- Are typically multi-purpose and can be fine-tuned and adapted to perform wide-ranging downstream tasks which may not have been envisaged by the model's creators.

How are AI foundation models made?

There isn't yet a well-accepted generalised process for building a foundation AI model. The topic is explored by CSIRO researchers Lu *et al.* (2023) in a recent paper. Here we make some broad generalisations about the stages typically involved in making foundation AI models.

1 Collect and prepare the dataset

This is a time consuming and costly task. The data may come from a wide variety of sources. In many AI foundation models it's unlabelled data but labelled data can also be used. Using fit for purpose data is critical. The data is broken up into smaller pieces called AI tokens which can be organised into sequences which are fed into the model at the training stage. The dataset used to train an AI foundation model is called the 'corpus'.

2 Choose the model architecture

This involves choosing the number of layers for the neural network and the types of layers and setting the number of nodes for each layer. At this stage a set of hyperparameters are also chosen relating to learning rate, batch size and optimisation algorithms.
The main different architectures that can be used include feedforward neural networks, convolutional neural networks, recurrent neural networks, long short-term memory and gated recurrent units, transformer architecture, graph neural networks, reinforcement learning models, autoencoders and generative adversarial networks and hybrid models.

3 Train the model

This can take hours, days weeks or months depending on model complexity, data size, hardware resources, the optimisation algorithm and other factors. Model training involves a forward pass where the sequence of tokens is entered into the model, a backward pass which calculates the loss (actual outcome versus predicted) and the model's parameters are updated using a technique called backpropagation. The forward and backward passes are then repeated for multiple epochs (entire passes through the model).

4 Evaluate and improve the model.

The model is then tested on an entirely different dataset that has not yet been used. The model is tested for precision and recall. Precision is the ratio of true positives to total positives. Recall is the ratio of true positives to actual positives. Recall and precision are used to calculate what's called an F1 score. There's typically a trade-off between precision and recall when designing and testing a foundation AI model. If the F1 score isn't high enough, then the model is recalibrated.

5 Fine tune and deploy the model

Although not always necessary, the model is sometimes fine-tuned on a smaller specific dataset to solve a specific problem. This involves making minor adjustments to parameters. The model is then deployed into an environment where it can be used by people to inform decision making, make predictions or provide useful information.

Working with existing models

The previously listed steps are what's involved in making an AI foundation model from scratch. However, it's not always necessary to do this (and incur the substantial costs). It might be possible to start at the last step and do the fine tuning and deployment of an existing AI foundation model. For example, this could be done with an opensource AI model which can sometimes be freely downloaded along with the model weights and computer code. The model weights are important information – they encode the relationships between features and target variables. The weights capture the patterns in the training data. One example of an open-source model that can be freely downloaded is the Llama2 large language model by Meta AI. They've made both the model and weights available as open-source products.

Furthermore, the providers of commercial foundation AI models are frequently fine tuning their models for specific applications which are available to users as add-ons. These are accessible to users via natural language interfaces or application programming interfaces (APIs). For example, in July 2023 OpenAI released Code Interpreter which is now called Advanced Data Analysis. OpenAI made this extension to ChatGPT by fine-tuning the GPT4 model on a large dataset of data analysis tasks (*e.g.* computer code, natural language descriptions of data analysis and the results of those analyses). It is likely that companies that have developed (pre-trained) large foundation AI models will continue to fine tune them for specific applications where there's a strong market demand.

Meta AI's open-access large language model is called Llama 2. This image was generated using ChatGPT and DALL-E 3.

A paradigm shift?

One of the key developments in the field of AI foundation models was the paper 'Attention is all you need' by Vaswani *et al.* (2017). According to Google Scholar this paper has 91 thousand citations (as of 5 October 2023) making it amongst the most highly cited research papers ever. The paper introduced a neural network architecture for machine learning called the transformer which achieved state of the art performance for many natural language processing tasks such as question & answering and summarising documents. The transformer architecture enabled much faster (and cost effective) training of machine learning models. The transformer architecture uses an attention mechanism to focus the model on the most salient inputs. This new capability saw an explosion of machine learning models developed using the transformer architecture (Amatriain et al, 2023). It was a game changer.

Some experts consider foundation models to be a paradigm shift for the field of AI (Bommasani *et al.*, 2023). Recent developments have demonstrated that AI models can be pre-trained on one dataset and then applied to solve wide-ranging downstream problems with different data. It's been shown that AI can generate text, images, videos and audio based on natural language commands. The quality of the output can match, or exceed, the quality of human generated content. This includes tasks that involve ideation, creativity, emotional intelligence, common sense and complex problem solving. The extent of recent advances has led some AI scientists (Bubeck *et al.*, 2023) to question whether we're seeing the 'sparks' of artificial general intelligence (AGI). AGI is an unproven concept. An AGI is generally understood to be an AI system that can adapt, learn and problem solve with creativity and reasoning equal to, or better than, humans. The achievement of AGI is a stated objective of the company OpenAI.

However, technology development seldom follows a linear trajectory and there's many blockages on the way. The astonishing progress of the past few years was partly achieved by using much bigger training datasets. OpenAI has not confirmed the number of parameters in GPT4 (released March 2023), the estimate on Wikipedia is 1.76 trillion parameters which is much more than the 175 billion parameters used in GPT3 (Alarcon, 2020). That's why GPT4 was so much better than GPT3. However, OpenAI CEO Sam Altman has stated that the gains from using larger models, with more parameters, have now been mostly achieved for OpenAI (Knight, 2023). There's diminishing marginal returns to model size. Future advances may come from increased computational power, knowledge discovery and other innovations. There's considerable uncertainty about the future pathway for AI foundational modelling capability. Nevertheless, even if nothing new were invented (which is not going to happen) there's decade's worth of industry adaptation and adoption from existing AI foundation modelling capabilities.

A surge of activity

Following the release of the transformer architecture by Vaswani *et al.* (2017) the number of AI foundation models has risen sharply. A recent review paper by Amatriain *et al.* (2023) identifies 77 AI foundation models based on the transformer architecture released since 2018.

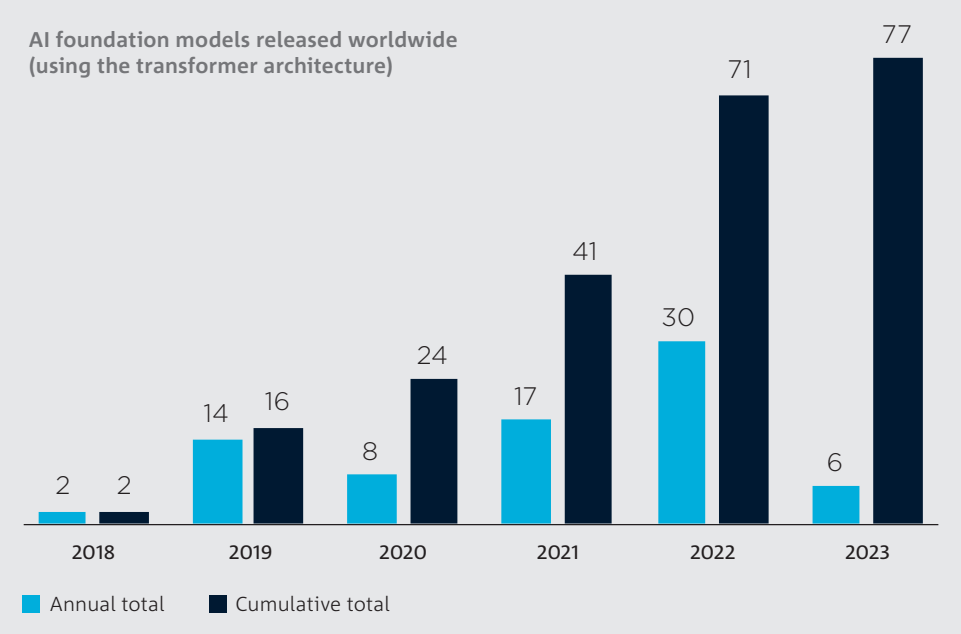

In addition to an increasing number of foundation models the models are being trained on large datasets. The paper by Amatriain *et al.* (2023) provides estimates on the number of parameters used in training for 39 (of the 77) models for the years 2019 to 2023.

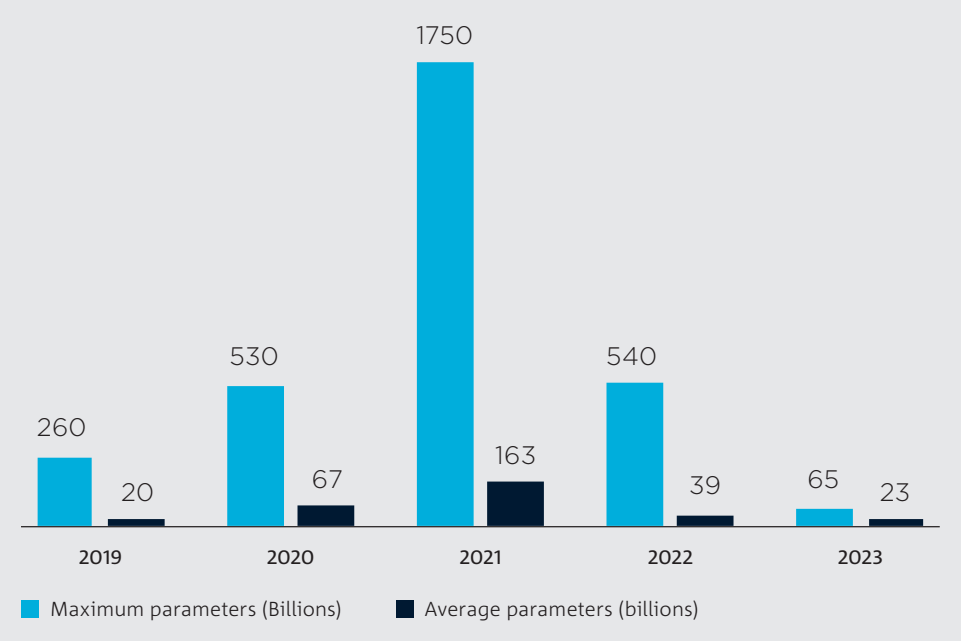

Computing infrastructure needs

Making and fine tuning a foundation AI model requires high performance computing infrastructure. Demand for these resources is high. Supply is limited. And prices are high. Getting access to compute power is probably one of the main challenges Australian start-ups and researchers face for developing AI foundation models.

AI accelerators – Graphics processing units (GPUs)

AI foundation model development requires computer systems containing GPUs that can handle a vast volume of matrix algebra calculations.

Examples of GPUs for doing this include the NVIDIA A100, NVIDIA H100, AMD Instinct MI300X, Google Tensor Processing Unit (TPU) v4 and Cerebras WSE-2. Without these types of products standard computers would take centuries to build an AI model like GPT-3.

Clusters and supercomputers

Supercomputers for training machine learning models contain clusters of AI accelerator chips. According to NVIDIA a cluster of 1,024 A100 chips were used to train OpenAI's GPT3 model in 24 days. The H100 GPU is a higher performing GPU that would complete the task quicker. It is estimated that Meta AI's Llama model was trained on 2,048 A100 GPUs in 21 days, Google's LaMDA model on 1024 TPUs in 58 days and the MT-NLG model by Microsoft/NVIDIA on 4,480 A100 GPUs in 90 days (CMA, 2023).

Demand for GPUs outstrips supply

For the first time Microsoft's annual shareholder report (30 June 2023) alerts investors to the operational risk of insufficient GPUs to support AI development. Only a limited number of companies manufacture high performing GPUs to support AI applications. Microsoft and the AI ecosystem can't get enough GPUs, and they're concerned about supply due to the limited number of products and manufacturers. If Microsoft is struggling to get GPUs then governments, universities and startups will struggle even more to acquire GPUs for AI development in this market.

Technology will improve, supply should increase

But so too may demand. The global queue for GPUs is long and will take some time to clear. The future is likely to see more manufacturers start supplying GPUs with improving speed and performance. There's also paradigm shifting technologies such as quantum computing and biological computing which could be associated with huge increases in computational power. However, there is much uncertainty about the timing and capability of these emerging technologies. Demand is likely to grow faster than supply for some time to come.

Prices for GPUs will come down if supply shifts right. But they go up if demand shifts.

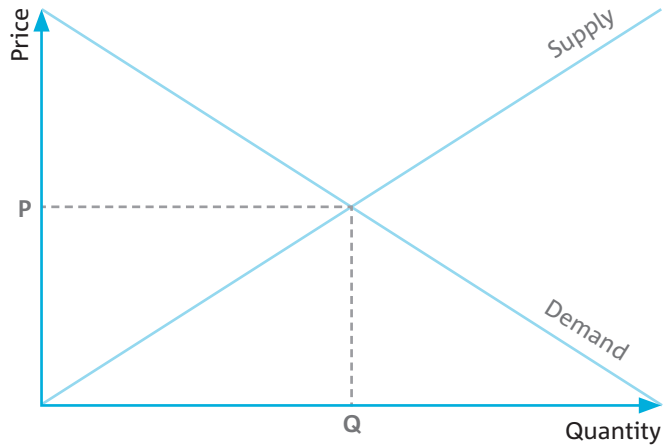

Build or buy?

One of the main strategic questions for AI model developers is whether to build their own compute infrastructure or buy access to computational resources from cloud-service providers. Both are limited by GPU availability and costs. If system use is continuous, it has been shown that building high-performing computers is more cost effective than buying cloud services (Villa and Troiano, 2020). Other considerations relate to scalability, useability, flexibility, security and reliability. Both cloud-computing services and in-house high performance computing infrastructure will play a role in supporting AI foundation model development.

Consideration of energy consumption and greenhouse gas emissions are becoming an increasingly important consideration in the build or buy decision.

Productivity and economic impacts

AI foundation models are a type of general-purpose technology applicable in all industry sectors and capable of performing wide ranging tasks. As such they have potential to achieve economy-wide productivity uplift and improved living standards.

Productivity growth	GDP growth
1.5%	3.1%
⌄	⌄
1.2%	2.2%

The productivity slump – can be changed with digital technology

The Australian Treasury (2023) recently downgraded forecast productivity growth from 1.5%/yr to 1.2%/yr. Partly due to ongoing productivity decline GDP growth over the coming 40 years is forecast at 2.2%/year compared to 3.1%/year for the previous 40 years. The Treasury report states that the 'expanded use of digital technology' could be one of the key (few) factors which increases productivity, income growth and living standards. Given its novelty and broad applicability AI, and in particular foundation models, can be considered one of the most prospective ways of shifting the dial on productivity growth.

Research into productivity impacts

- A study of 5,257 global firms with AI patenting activity during 2000-2016 finds that AI has a positive and significant impact on labour productivity. Benefits were greatest for SMEs and service-sector industries (Damioli *et al.*, 2021).
- A study of 1,353 AI innovating firms in Taiwan's electronics sector during 2002–2018 found AI was positively associated with productivity and employment (Yang, 2022).
- A study of 5,851 firms in Germany (409 of which were AI users) found 'positive and significant associations between the use of AI and firm productivity' (Czarnitzki, 2023).
- Brynjolfsson and Raymond (2023) studied 5,179 customer support agents in a software company and staff using Generative AI (GAI) could resolve 14% more customer problems. They observed improved customer sentiment, less requests for management intervention and improved employee retention.
- A study of management consultants using GAI by Dell'Acqua *et al.* (2023) found consulting tasks were performed 25% faster and 40% better using GPT-4. However, consultants were 19% less likely to perform a task correctly when using GPT-4 for tasks beyond its capability envelope. Lower skilled consultants had a greater productivity boost from using GPT-4.
- Microsoft and the Tech Council of Australia (2023) estimate GAI could automate 44% of tasks for Australian firms and value-add A$115 billion to the Australian economy by 2030. They cite recent studies finding that GAI computer coding tools reduce staff task time by 56% and GAI writing tools reduce staff task time by 37% with improved quality (Kalliamvakou, 2022; Noy and Zhang, 2023).
- Australian bank Westpac reported that early experiments with GAI tools resulted in a 46% increase in the productivity of software engineers in coding tasks with no loss of quality (Tonkin, 2023).
- The Australian Government Productivity Commission (PC, 2024) finds that AI 'could significantly improve the productivity of the services sector, which has a history of poor productivity growth' but adjustment and experimentation is needed. It will also take time for economy-wide productivity gains to be evident.

Solow's paradox and productivity challenges

However, analysts still struggle to measure productivity benefits of AI and digital technology at the economy level. This is known as Solow's Paradox (Capello, 2022); widespread evidence on the increased use of digital technology but with overall productivity growth still showing no change from its long-term declining trend (Australian Treasury, 2023). Recent reports from the association of Certified Practicing Accountants (CPA, 2023), Australian Computer Society (Braue, 2023) and the Reserve Bank of Australia (Nguyen and Hambur, 2023) examine mixed (positive and negative) productivity experiences associated with AI adoption by companies. These 'coal face' reports seem to conflict with some of the productivity outcomes observed by researchers above. Overall, it's clear that the productivity gains of AI aren't assured and depend upon the nature of tools used, the quality of training data and the human-AI interfaces (Hajkowicz and Whittle, 2023).

Global policy responses

The rapid ascendancy of AI foundation models has been met by policy responses and proposed new legislation by governments worldwide focused on ensuring the broad-based benefit and competitive advantage of these models for citizens (OECD, 2023a). Here we examine policies from the United Kingdom, the United States, Japan and Germany which hold relevance to Australia.

The United States

A white paper published by the Stanford University Human-Centred Artificial Intelligence research unit (Ho *et al.*, 2021) presents a blueprint for the US National AI Research Resource (NAIRR). The NAIRR is part of the National Defense Authorization Act (NDAA) enacted by US Congress in January 2021. The NAIRR is 'a system that provides researchers and students across scientific fields and disciplines with access to compute resources, co-located with publicly available, artificial intelligence-ready government and non-government data sets' (Ho *et al.*, 2021). It designed to support 'basic scientific AI research, the democratization of AI innovation, and the promotion of US leadership in AI'. It will support the development of AI models; including foundation models.

The bipartisan bill to fund the NAIRR is currently before congress and was introduced in July 2023 and is called the Creating Resources for Every American To Experiment with Artificial Intelligence Act of 2023 (CREATE AI Act). This is the result of an ecosystem of AI experts and researchers working over several years. The idea of a national research cloud was conceived in 2019 by researchers at the Stanford University Human-Centred Artificial Intelligence research unit. A Stanford University media release argues that the CREATE AI act is necessary for the United States to ensure the entire research community has democratised access to models, compute-power and datasets to develop public-good AI models and tackle blue-sky research problems (Wald, 2023).

On 24 January 2024 the United States National Science Foundation (NSF, 2024) announced the commencement of the NAIRR pilot program. It involves 10 federal government agencies and 25 private-sector firms. The press release indicates the 'pilot is a proof of concept to ignite the level of investment needed to realize the full NAIRR vision'. The CREATE AI Act is currently being reviewed by both Chambers of the US Congress and has bipartisan support. If enacted it will become one of the first and largest public sector initiatives to provide both computational resources and datasets to democratise AI development across the entire national R&D ecosystem.

US tech industry is also making commitments to support the broader AI R&D ecosystem. For example, Microsoft has made a commitment to 'increase investment in its academic research programs to ensure researchers outside Microsoft can access Microsoft's foundation models and the Azure OpenAI Service to undertake research and validate findings' (Microsoft, 2023). And many of the most significant AI discoveries, including the transformer architecture (Vaswani *et al.*, 2017), have been published freely in the academic literature by Google. Meta AI has made its large language model Llama2 freely available to download via open access. OpenAI, the creators of DALL-E and ChatGPT, has a stated mission to 'ensure that artificial general intelligence—AI systems that are generally smarter than humans—benefits all of humanity'. They've also published a charter which has the principles of (a) broadly distributed benefits; (b) long-term safety; (c) technical leadership; and (d) cooperative orientation.

Whilst profit drivers may play a role there's also considerable free and public release of valuable AI resources by the private sector. Policy instruments such as the CREATE AI act are about making datasets and compute-power available to the entire AI ecosystem. This will allow blue-sky research with no immediate commercial driver. It will also allow AI models to be built to address wide ranging social problems. Many academics in the US AI community (Wald, 2023) argue that profit-drivers will dominate the actions of technology corporations and that public sector investment is needed to ensure broad-based benefits and ethical application of AI foundation models.

In addition to these developments issues of safety and security relating to AI model development and application are receiving increased attention. On 30 October 2023 the United States President issued an executive order 'the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence' (White House, 2023). This statement identifies the 'promise and peril' of AI. It is weighted more towards risk mitigation and harm prevention, rather than capability development.

The United Kingdom

In April 2023 the UK Government (GOV.UK, 2023) announced an '*initial £100 [A$191] million for expert taskforce to help UK build and adopt next generation of safe AI*'. The press release notes this is in addition to £900 million (A$1.7 billion) invested into high performance computing infrastructure. The press release indicates that:

'*The investment will build the UK's 'sovereign' national capabilities so our public services can benefit from the transformational impact of this type of AI. The Taskforce will focus on opportunities to establish the UK as a world leader in foundation models and their applications across the economy, and acting as a global standard bearer for AI safety. The funding will be invested by the Foundation Model Taskforce in foundation model infrastructure and public service procurement, to create opportunities for domestic innovation. The first pilots targeting public services are expected to launch in the next six months.*'

Based on the publicly available information it appears that this initiative will involve building a set of foundation models which initially will be designed to improve public sector service delivery. The UK Government policy objectives in this initiative relate to productivity, economic growth, the competitiveness of UK industry and the ethics safety, security and reliability of AI foundation models used in the UK.

Another set of activities are underway in the UK Government by the Competition and Market Authority (CMA). This agency is responsible for ensuring the efficient and fair operation of competitive markets and protects people from unfair trading practices and/or anti-competitive practices. The CMA released a report on AI Foundation models in September 2023 (CMA, 2023) which found different market outcomes.

1. A positive market outcome associated with efficiency, productivity and fairness will occur if there are multiple independent developers competing to produce and deploy AI foundation models into the marketplace using open-source and closed-source approaches.
2. A 'concerning' market outcome will occur if access to inputs (datasets, compute power, expertise) are limited to only a few firms providing models on a closed-source basis and imposing unfair prices and terms given the lack of competition which would result in decreased productivity and economic growth.

The CMA (2023) report suggest that where on the spectrum of these two outcomes eventuates will depend on access to data, compute power, first mover advantage and the existence of open-source models. They also note the emergence of proprietary data which can be accessed by only some firms to train AI foundation models; this could have the effect of locking other firms out of the market. It is in these spaces that the UK Government could intervene to ensure the market for AI foundation models is efficient and fair and maximises productivity and economic growth.

The CMA (2023) review identifies a series of next steps to ensure 'positive outcomes' in the marketplace from AI foundation models including consultations with experts, consumer groups, model developers and deployers, academics and governments. The CMA plans to publish the results of this analysis in 2024 and this may lead to policies or regulations to guide the AI foundation model market.

Japan

The Liberal Democratic Party (LDP) is a political party in Japan that has held a governing majority in both houses of the Japanese parliament almost continually since its formation in 1955. Policy positions published by the LDP are significant for the Government of Japan. In April 2023, the LDP published a whitepaper on 'Japan's National Strategy in the New Era of AI' (LDP, 2023). The whitepaper makes opening references to 'ChatGPT' and notes that the 'new era of AI' is associated with foundation models which could be a catalyst for new sources of economic growth.

The first and main recommendation of the LDP whitepaper is that the Japanese government should 'build and strengthen AI model development capabilities, including foundation models'. This is to be achieved by building on overseas made foundation models or by forming partnerships with leading developers worldwide. The LDP whitepaper notes that 'it will not be easy' to develop AI foundation models with domestic resources in the short term.

The remainder of the LDP paper makes recommendations about access to skills and expertise, datasets and computational resources for Japan's AI ecosystem. This includes resources that 'can be shared and utilized by all related public and private entities'. The LDP paper suggests new approaches to AI regulation to ensure sovereign capability, security and protect the well-being of Japanese citizens. An article in Reuters reports Hideki Murai, a ruling Liberal Democratic Party's (LDP) lawmaker and special adviser on artificial intelligence saying that '*the government's key priority is computing power. We feel a real sense of crisis about that' and that 'we want to create the foundations for an AI era*' (Kelly and Nussey, 2023).

In September 2023 an article in nature news (Hornyak, 2023) reported that the Japanese government and big technology firms such as NEC, Fujitsu and SoftBank are building their own Japanese versions of ChatGPT. The large language models (LLMs) under development in Japan aim to capture the intricacies of Japanese language and culture which are not fully handled by other LLMs developed offshore. There is also a hope that Japanese LLMs can have a positive impact on joint international research for Japan.

Germany

In Germany, the feasibility study of the LEAM initiative (Large European AI Models) by the German AI Association for the German Government published a report on AI foundation models (AKI, 2023). It notes the dominance of US and China in AI Foundation models, and that Germany risks missing the 'paradigm shift'. The LEAM report proposes a framework to support the development of AI foundation models in Germany. This involves developing and making available to academia and industry skilled expertise, datasets and compute power. The aim is to make trustworthy AI foundation models to support German businesses. The German LEAM initiative is mostly about making new AI foundation models and building sovereign capability.

The German whitepaper notes that since 2017, 73% of AI foundation models come from the United States and 15% are from China. Most of the remaining models come from Europe with very little emanating from other countries. They also note that most of the models come from industry with 86% of models developed by the business sector and 13% from the academic sector. Furthermore, the German whitepaper presents an analysis of AI foundation model capability by source country. They find that all the best-known models, with the most advanced capabilities, come from the United States and China. The analysis is based on 125 AI foundation models from the transformer family.

Australian AI product innovation

In many fields of science and technology Australia has high levels of skills, knowledge and capability but comparatively low levels of product innovation. A commonly cited example is that of solar panel technology (Watt, 2003). The University of New South Wales invented the solar cells which are used in over 80% of the world's solar panel today (Egan, 2022). However, the domestic solar panel-making industry is small with 99% of Australia's solar panels manufactured offshore (DCCEEW, 2023).

The challenges of product innovation and commercialisation associated with Australian R&D has been observed by researchers for decades (Martin, 1991; Harman and Harman, 2004). Recognising the issues, Australian governments have significant policy instruments designed to catalyse R&D commercialisation (DE, 2022).

The same challenge exists for AI R&D as per other fields. An analysis of research activity and patent applications related to AI in Australia using The Lens database found that whilst Australia's contributed 1.6% of all research publications on AI globally, we contributed 0.2% of all AI patent application inventions (Hajkowicz *et al.*, 2023). This trend has persisted over the past 20 years.

Patent applications don't capture all product innovation, some research suggests about one-third (Arundel and Kabla, 1998). But they are an important metric about the conversion of research outcomes into commercial products. The gap between Australian research on AI and product innovation is likely to exist for AI foundation models. From these data Australia is more of a downstream user, rather than maker, of AI products. However, there's every possibility that Australia's R&D sector could develop AI foundation models in target/niche areas and increase the number and value of exportable commercial AI products.

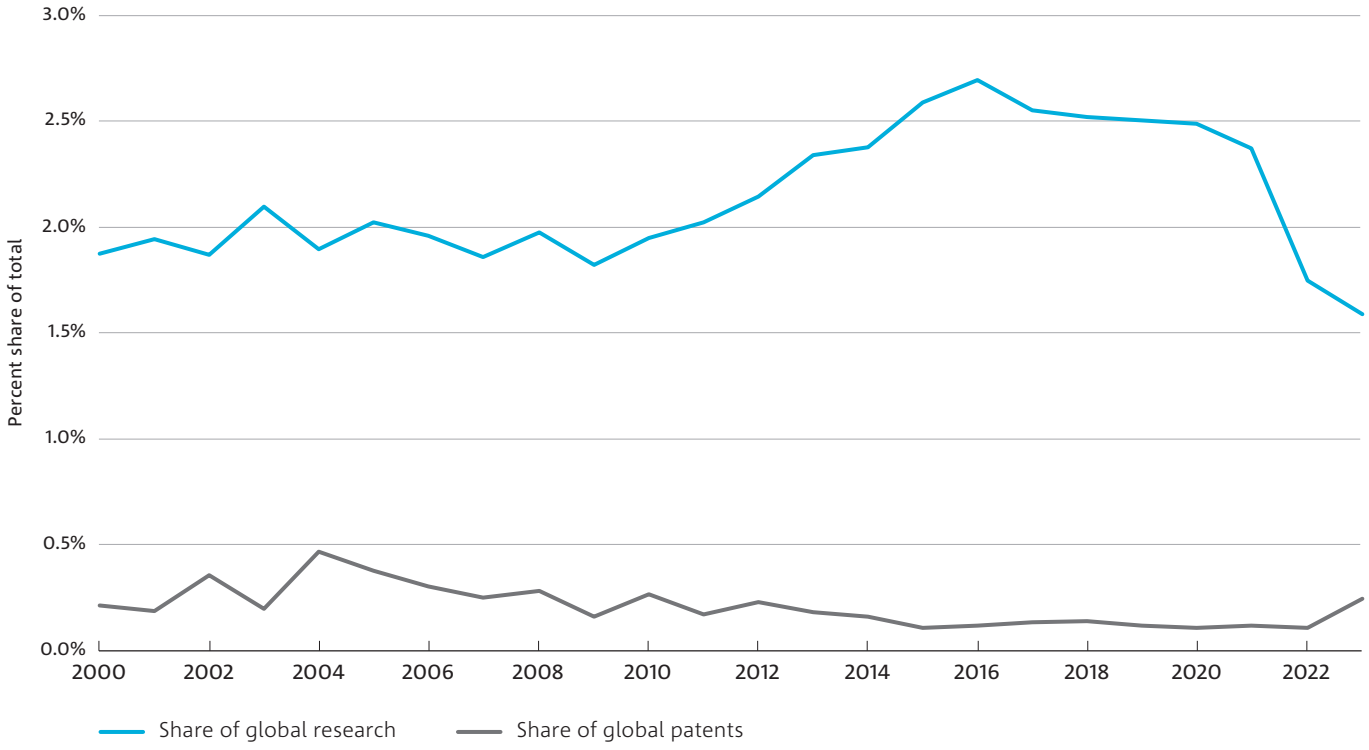

Australia's share of global AI research outputs and patent application (inventions)*

* The last data point is for 2023 and is based on data only from January to June 2023.

Sovereign capability considerations

Sovereign capability is a concept receiving increased attention in across most industries in Australia, especially the defence sector (DoD, 2021). In a generalised sense sovereign capability is the ability of nation states and governments to maintain societal, industry and economic functioning, and ensure the needs of industry and society are met, if trade and technology-sharing with the outside world is shutdown. In this paper we're interested in the sovereign capability dimensions of AI foundation models for Australia. We identify two dimensions of sovereign capability relating to AI foundation models:

1. The capability of nation states to build, operate and manage AI technology drawing upon data, skills, knowledge, models and computational resources within the nation or its direct jurisdictional reach.
2. The capability of government to deliver functions and services using AI technology despite changes in the ecosystem of private-sector AI suppliers. And the capability of government to guide and regulate AI development and application by the private sector.

Sovereign capability doesn't necessarily mean the whole AI model is developed and managed from within Australia; it's about our ability to manage the way the model is used and our ability to maintain socio-economic activity if the model is made too costly, inaccessible or abruptly changed in some way. Sometimes this might mean building and operating the model from within Australia; other times it may mean having the skills, resources and optionality to manage models built offshore.

Given that 88% of AI foundation models are from the US and China (AKI, 2023), with the remaining 12% emanating from Europe and a few other countries, many national governments are considering issues of sovereign capability. These issues become more important when the AI models attain widespread use and high levels of dependence for information, decision making and important tasks. We can already see this happening with AI foundation models released into the Australian marketplace this year.

There isn't anything inherently wrong with using AI models made and managed offshore, it's both an opportunity and risk. Here we identify some of the reasons governments and societies are considering sovereign capability when it comes to AI foundation models.

Competitive markets and pricing

Researchers have observed that markets for digital platforms, machine learning and AI tend to form natural monopolies where one supplier (or a few suppliers) dominate market share (Ducci, 2020; Narechania, 2021). There's a possibility that AI foundation model markets will evolve in a similar manner. Dominant market entities may have restricted access to training datasets, computational power and a broad user base. This can shut out the competition. The owners of the AI foundation models can set prices for API access at their discretion. If they have monopolistic presence in the marketplace, they're able to set high prices and extract rents because there's insufficient competition.

Whilst there are benefits of monopolies, they are mostly seen as harmful to consumers and economic growth (Ducci, 2020). Most governments in advanced economies have agencies mandated to identify and mitigate the negative impacts of monopolies and/or unfair markets with limited competition. AI foundation model markets have been identified as a key concern for the Competition and Market Authority (CMA, 2023) who will be making recommendations for the UK Government about the matter in 2024. The same risk applies here. This risk is partly about sovereign capability and partly about an efficient, fair and competitive marketplace for AI foundation models.

A recent study by the Brookings Institute (Vipra and Korinek, 2023) draw similar conclusions to the CMA (2023). It finds markets for cutting-edge AI foundation models exhibit strong tendencies towards concentration and monopolistic structures. It advises that 'regulators need to ensure users experience reasonable pricing, high quality standards (including safety, privacy, non-discrimination, reliability, and interoperability standards), as well as disclosure and equal access rights'. Lastly, the report also notes that AI foundation models have the potential for novel, wide-ranging and deep economic impacts because they have such widespread application across the entire economy.

Reliability and stability

Technology platforms used by millions of Australian citizens and business can change abruptly. The greater the dependence on these platforms, the greater are the impacts. For example, social media platform Twitter has undergone substantial changes and is now rebranded as 'X' with changes to content moderation and pricing structures. This has impacted governments, companies, citizens and science communities that used Twitter extensively (Stokel-Walker, 2022). We've also seen technology platforms threaten to switch-off search in Australia (Judd, 2021) and impose temporary bans on news pages in Australia (Easton, 2021) over disagreements with the Australian Government about proposed media laws. Sudden changes to technology platforms used, and relied upon, by a nation's citizens can have a significant socio-economic impact. Governments may have little control or influence over the changes. Similar risks would apply to overseas owned and managed AI foundation models that have widespread application in Australia.

Cultural relevance and sensitivity

AI foundation models, and particularly generative AI foundation models, use and generate text, audio, video and imagery. The nature of the outputs they generate will be influenced by the inputs. This means the important and nuanced cultural norms and expectations of countries might be missed in material generated by the AI. Foundation AI models trained on offshore datasets could also potentially give inaccurate advice with respect to local laws, requirements and processes. The German, UK and Japanese projects to develop sovereign AI capability identify the need for culturally relevant content as a key reason (GOV.UK, 2023; LDP, 2023; AKI, 2023).

Confidentiality, privacy and security of information

Foundation AI models tend to become much more useful when users upload their own information (documents, images, videos, audio *etc.*). However, this creates a significant risk of private, sensitive and confidential information falling into the wrong hands. Whilst the providers of the AI foundation models are working on security measures there will be some uncertainty about the trustworthiness and reliability of their systems. At the current time much public sector use of AI foundation models is not possible because Australian governments cannot securely upload information (DTA, 2023). Future concerns may emerge about the extent of Australian industry information being fed into foundation AI models that could be used by other industries or governments.

Geopolitical risks

Social media platforms have been used by political organisations and nation states seeking to mislead, manipulate and/or disrupt normal functions of governments and society in target countries (Garrett, 2019; Gorodnichenko *et al.*, 2021). It's conceivable that AI foundation models relied upon by Australian citizens for wide ranging decisions and tasks could be subject to similar misuse. They're also at risk of adversarial prompts (Zou *et al.*, 2023) and data poisoning (*Ge et al.*, 2023) which are emerging techniques used by adversaries to manipulate responses from, and behaviours of, generative AI models. Furthermore, generative AI models could also be used to generate fake and highly personalised (and therefore realistic) audio, text, imagery and video designed to deceive individuals. Without sovereign capability governments may struggle to identify and mitigate misleading or manipulative content.

Ethics and content moderation

AI foundation models developed in other countries by private-sector corporations may not always meet the AI ethical standards of Australia (DISR, 2019). They may not have content moderation systems that sufficiently protect Australians from extremist material. And they may use training data and systems that fail to meet our ethical standards or laws. If AI foundation models were found to be ethically non-compliant or failing to meet laws or regulations, the Australian government may struggle to enforce these standards and laws given the decisions are being made in other countries.

Workforce impacts

Whilst many of the concerns raised about widespread job loss associated with AI (Frey and Osborne, 2017) did not happen (Coelli and Borland, 2019) AI foundation models might bring new challenges. They're likely to be unleashed into the services sector of the economy where the bulk of Australians work. Generative AI platforms such as ChatGPT have cognitive automation capabilities that will impact knowledge workers which represent over half of the workforce in advanced economies like Australia. Many industries including administrative services, banking, finance, software development, retail, health, energy and education could be impacted.

One recent study (Hui *et al.*, 2023) examined the short-term effect of ChatGPT on online freelancers in knowledge jobs considered risk-exposed such as writing-related services. The researchers found 'generative AI reduces overall demand for knowledge workers of all types, and may have the potential to narrow gaps among workers'. They also found that following the release of ChatGPT the number of jobs for freelancers on a major online platform decreased by 2% and total monthly compensation for those jobs decreased by 5.2%.

If foundation models impacting the Australian labour market are owned and managed offshore the workplace impacts might be difficult to mitigate; as with the impact of Uber on taxi drivers. Actions might be needed to ensure workers can adapt to, and benefit from, AI foundation models.

Policy levers

There's multiple policy levers which can be applied by governments to address sovereign capability issues relating to foundation AI models. These policy levers can help governments harness the opportunities and mitigate the risks of this technology. Here we identify a set of generic policy levers available to governments at all levels (Federal, State, Territory and Local) based on what we can see happening across the globe.

Build high performance computing infrastructure and democratise (and prioritise) access

The development of AI foundation models is limited by access to high performance computing. And that's mostly limited by access to AI accelerator chips – GPUs. There's not much Australian industry, governments, researchers or community can achieve in AI development without access to these GPUs. And they're under heavy demand. Microsoft signalled to investors that the lack of GPUs for AI is one of the key business risks in the near term (Microsoft, 2023). Many researchers in academia, industry, government and community can't get access to computational resources to design, build and fine tune AI foundation models (Ho *et al.*, 2021). In Japan it's seen by policy makers as the main limiting factor for realising the full benefits of AI (Kelly and Nussey, 2023). The UK Government recently announced an A$430 million project to build a supercomputer called 'Isambard-AI' at the University of Bristol to train and develop AI models. The UK Government media release says the 'new facility will serve as national resource for researchers and industry experts spearheading AI innovation and scientific discovery' (GOV.UK, 2023b). Investments in AI computing resources of this nature are being made because many scientists, researchers and innovators and start-ups don't have the computational resources to turn their AI idea into a product. Making this available could fuel the rate of innovation and see the emergence and scale-up of AI startups. Systems and criteria would be needed to prioritise access to these resources where social, economic and environmental benefits are greatest.

The OECD recently published a blueprint for building national compute capacity for AI (OECD, 2023b). This report identifies the lack of strategies by governments for improving computational resources as a 'blind spot' in national AI plans. The OECD (2023b) paper notes that the computational resources needed to train machine learning systems has increased by hundreds-of-thousands of times since 2012. Contemporary AI systems need specialised hardware, software and skilled workers who can use/make them. A lack of access to these things could thwart government and industry AI ambitions at the national level.

Negotiate bilateral or multi-lateral international collaborations to share AI expertise and resources

The vast bulk of AI research and product development is occurring outside Australia. The German whitepaper on AI foundation models estimates that 73% are from the US and 15% from China (AKI, 2023). However, there is broad recognition that a globally cooperative approach to AI development will benefit all countries. For example, the Global Partnership on AI (GPAI) now has 25 members including Australia, the US, Canada, the UK and 21 other countries of the Organisation for Economic Cooperation and Development (OECD, 2023b). This international

partnership is facilitated by the OECD and aims to catalyse AI development via knowledge sharing and ensure safe and ethical AI development across the globe. Researchers at Stanford University Zhang *et al.*, 2021) note that the National Security Commission on Artificial Intelligence (NSCAI) made a recommendation to the US Government to establish a Multilateral AI Research Institute (MAIRI). The MAIRI would promote international R&D cooperation on AI to boost innovation and economic prosperity. The Stanford University team (Zhang *et al.*, 2021) have proposed a blueprint for the establishment of the MAIRI and have called for congressional authorisation of funds for the National Science Foundation to lead the effort. Partnerships with overseas research agencies, universities, governments and companies will help Australia access AI expertise, data, knowledge and resources.

Increase workforce skills via training, education and improved access to national and global talent pools

Working with AI foundation models requires specialist skills in fields such as software engineering, natural language processing, computer vision, machine learning, mathematics, statistics and data science. However, there's a shortage of these skills. A recent survey by software firm SAS (2022) found that 63% of organisational decision makers in the UK don't have access to enough employees with AI or machine learning skills. The UK government (UK. GOV, 2021) estimates that UK universities are unlikely

to supply over 10,000 data scientists per year which falls short of the 215,000 jobs per year needing data science skills. The Australian Computer Society finds that nearly 70% of information technology jobs are in shortage, this includes jobs requiring AI skills. The skills necessary to make, fine-tune and operate AI foundation models can be developed via training and education initiatives or by ensuring the mobility of skilled workers. However, policy interventions in this space need to consider the science, technology, engineering and mathematics (STEM) skills paradox; where in markets with seemingly high demand many STEM graduates struggle to find work (Harris, 2014). Explanations include the challenges for STEM graduates to move to cities/regions with jobs, the highly specialised nature of STEM skills making an exact match hard for employers and job-seekers and the requirement for complementary skills such as communication skills, creativity, business skills *etc* without which the hire will not proceed. Despite these explanations the STEM skills paradox remains a paradox.

Develop resources, policies, regulations and testing systems to ensure the productive, safe and ethical development of AI foundation models

There's growing recognition worldwide that the rise of AI foundation models creates both opportunities and challenges for governments that may warrant new capabilities and new regulations (CMA, 2023; Brookings Institute, 2023; Ho *et al*., 2021; Bommasani *et al*., 2023). In 2023, the Australian government engaged in discussions with industry, community, and academic groups regarding the safe and ethical implementation of AI. On January 17, 2024, the government released a preliminary statement detailing strategies to ensure AI is securely and dependably used in critical, high-risk situations, while allowing for relatively unrestricted use of AI in scenarios deemed low-risk (DISR, 2024).

At a congressional hearing in the United States on 16 May 2023 OpenAI CEO Sam Altman urged the US Congress to regulate frontier AI models (Kang, 2023). In a recent paper (Anderljung *et al.*, 2023) published on arXiv OpenAI researchers state, 'advanced AI models hold the promise of tremendous benefits for humanity, but society needs to proactively manage the accompanying risks' and that 'highly capable foundation models that could possess dangerous capabilities sufficient to pose severe risks to public safety'. They propose a set of safety standards.

One approach for improving the safety of AI foundation models is a foundation model crash/safety testing capability. This would be a team of experts that attempts to 'break' or 'crash' an AI foundation model to ensure it's safe and ethical before released into society. This would be analogous to the way that the Australasian New Car Assessment Program (ANCAP) crash tests and rates cars to assess safety. AI foundation models could be tested for malicious or incorrect use that may breach AI ethics principles, safety or other community standards.

Beyond safety there's many considerations about ethics, useability, productivity and competitive markets. Researchers at Stanford University (Bommasani *et al.*, 2023) note that despite the impending widespread deployment of foundation AI models 'we currently lack a clear understanding of how they work, when they fail, and what they are even capable of due to their emergent properties'. They argue for an interdisciplinary effort into all technical, social, economic, cultural and scientific dimensions of AI foundation models to inform effective policies and regulations.

Identify, validate and make available datasets which could be used to train AI foundation models

The past decade has seen an explosion of open data schemes by governments worldwide (Attard *et al.*, 2014). The provision of open data has been shown to generate substantial economic benefit (Petrović *et al.*, 2022). Businesses use these data to improve processes and create products and services (Hardy and Maurushat, 2017). Startup companies can create business models and scale-up with open data (Zeleti et al., 2016). Open data assets are becoming more valuable in the era of AI and, especially, foundation AI models (Gao and Janssen, 2020). Data enable the training of machine learning algorithms which can be applied more broadly across industry. And the quality of data will underpin the quality of the AI models and their usefulness to industry and society. Government holds vast quantities of data on wide ranging topics that could be used to train or fine tune AI models. Data security, privacy and confidentiality requirements must be upheld. And metadata will be required by model developers. But datasets will enable model development and novel solutions to wide ranging problems. There is much government is already doing via open data initiatives, and can do more, to make datasets available to Australia's AI ecosystem. Data quality, accessibility and reliability are key to successful open data programs. The provision of open data to the AI community is a central mechanism of the National Artificial Intelligence Resource and CREATE AI Act (Ho *et al.*, 2021).

Invest in building, adapting (fine-tune) and applying AI foundation models to improve government functions

The Australian Productivity commission identifies the building of foundation AI models as one of the policy interventions governments can make to achieve productivity benefits from AI (PC, 2024). AI foundation models have the potential to boost public sector productivity and improve the quality of government services and functions. However, government use of commercially available, and/or open access, AI foundation models is limited due to security, safety and ethical considerations (DTA, 2023). Governments may seek to develop specialised in-house AI foundation models to meet these requirements. Governments may also need AI foundation modelling capabilities not currently on-offer in the marketplace. Government investment in AI can catalyse growth of the domestic AI industry and workforce by purchasing a significant quantity of local skills, products and services. Following are some possible AI foundation models that could improve the functions of government:

- *Service delivery* – The interaction of citizens with government for wide ranging services (*e.g.* tax, welfare, immigration, health ...) could be made more efficient and effective using text-enabled or voice-enabled chatbots supported by AI foundation models (Darmawanto and Syahwami, 2023). Customer service staff could interact with a chatbot and moderate content or, in some cases, the chatbot could interact with the customer directly. Whilst online chatbots are already being used (*e.g.* the Services Australia digital assistant) next generation foundation AI models could substantially improve the scope of service delivery interactions that can be automated.
- *Policy formulation* – Governments prepare policy documents and reports on a regular basis on wide ranging topics. Foundation AI models could be used with human oversight to speed up this process and improve the quality of these documents. For example, researchers from the University of Salford in Manchester UK demonstrated the effectiveness of ChatGPT for writing company policies about quality management, data protection, equality, diversity and inclusion (Kutar and Fletcher, 2023).
- *Intelligence gathering and analysis* – With high levels of data security AI foundation models could be used to screen, interpret and inform decision makers about wide ranging risks (terrorism, crime, armed conflict ...). Often intelligence analysts are confronted with vast volumes of data and some type of automated digital approach is the only means via which the data deluge can be analysed. According to media reports the United States Central Intelligence Agency is building its own Chat-GPT style AI to support the work of its intelligence analysts working with open-source information (Martin and Manson, 2023).
- *Healthcare services* – AI foundation models could help with diagnosis, prognosis and treatment decisions in Australia's health system which is over two-thirds in the public domain. Researchers in the field of digital health are advocating for the careful development of such AI foundation models could also be used at the population level to screen for risks and design preventative health programs. systems as they can improve patient outcomes and improve cost-effectiveness (Fries *et al.*, 2022; Moor *et al.*, 2023).
- *Natural hazard and environmental forecasting* – AI foundation models could improve the spatial granularity and timeliness of forecasts for bushfires, floods, cyclone paths and heatwaves enabling improved protection, response and strategic planning. An example of such a model is the IBM and NASA geospatial AI foundation model trained on Earth Observation data. It can be used for tracking changes in land use, monitoring natural disasters and predicting crop yields (Blumenfeld, 2023). Similar AI foundation models could be adapted/built for Australian conditions.
- *Research and development* – AI foundation models could be designed to improve the process of reviewing the literature, designing experiments and interpreting the results (and the implications) for Australian R&D activities. The are initiatives across the globe underway to develop AI capabilities for science at the Argonne Laboratory in the US (Stevens, 2020), the Turing Institute in the UK (Gil et al. 2020) and CSIRO (2023) in Australia.
- *Other areas* – Government functions in education, transport, law enforcement, defence, resources, agriculture, environment, biosecurity and other sectors could be improved by AI foundation models.

A healthcare example

Medical researchers have been quick to identify the potential for AI foundation models to transform and improve health services (Fries *et al.*, 2022; Moor *et al.*, 2023; Harrer, 2023). Historically, AI models for healthcare have been designed to solve a specific problem or perform a specific task. However, AI foundation models with multi-modality capabilities can take wide ranging data as inputs and perform wide ranging downstream tasks. An example of how an AI foundation model for healthcare could work, adapted from Fries *et al.* (2022) and Moor *et al.* (2023), is as follows.

It's likely that most of the tasks performed by the AI foundation model would augment (not automate) what was previously done by a human in the healthcare system. And all would require oversight to ensure ethical and safe outcomes. It's possible that a medical AI foundation model could provide another layer of protection for patients to help clinicians with accurate diagnoses. The model could also flag risks otherwise unseen, help patients understand conditions and make better decisions about treatment. Clearly data confidentiality and privacy would be of utmost importance.

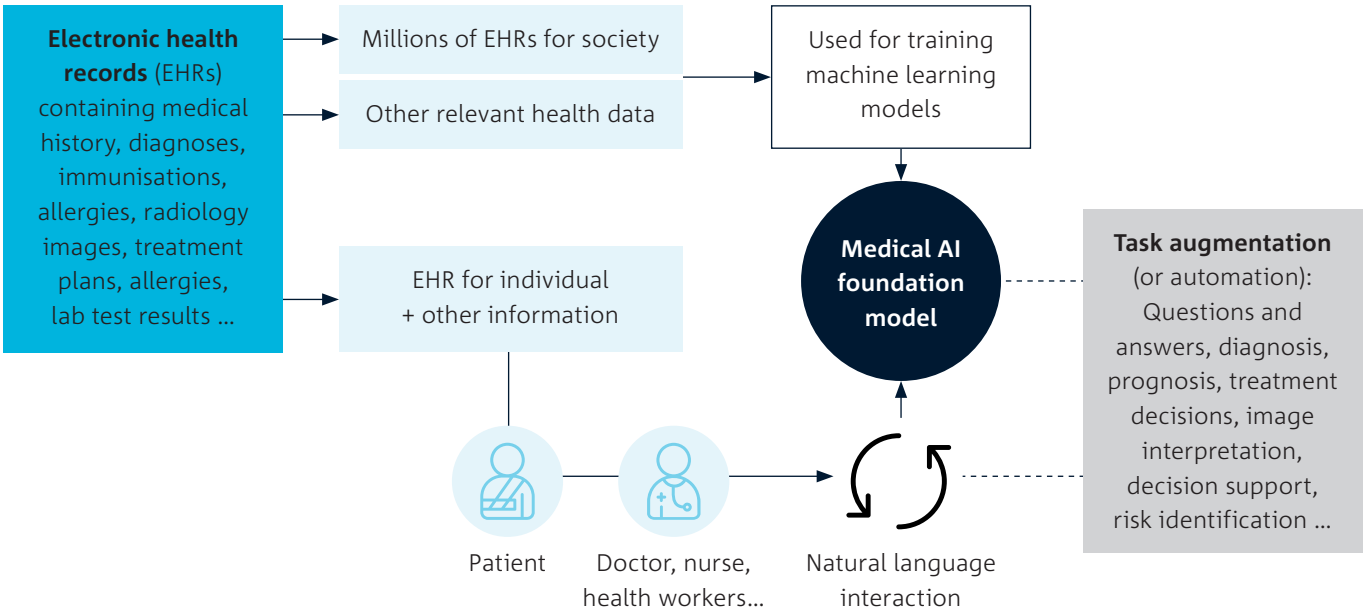

The benefits of such a system could be substantial. It could improve diagnosis and treatment decisions. A research paper in the Medical Journal of Australia (Scott and Crock, 2020) suggests that 14% of clinical encounters involve some type of diagnostic error. The article finds that Australia has 141,000 diagnostic errors each year of which 21,000 are associated with serious harm and 2,000 – 4,000 with death. The primary cause is 'cognitive errors' by clinicians involving the use of mental shortcuts or rules of thumb and overconfidence bias.

Another benefit from a foundation AI healthcare model could arise from bringing forward cancer detection. An early-stage detection is associated with improved survival and recovery chances. Given that 165,000 Australians were diagnosed with cancer in 2023 many people could benefit (AIHW, 2023). Governments, research institutes and healthcare service providers United Kingdom are actively investing in the use of AI to bring forward cancer diagnosis and they're achieving positive results (Hunter *et al.*, 2022).

This is an example for healthcare. But AI foundation models could be used to improve the efficiency and effectiveness of workflows in wide-ranging Federal, State/Territory and Local government functions and service delivery systems.

A weather and climate forecasting example

AI approaches are increasingly being used for climate, weather and natural hazard forecasting. Reviews of this field find AI improves the computational efficiency and overall quality and accuracy of forecasts (Dewitte *et al*., 2021). This creates significant social, economic and environmental benefit. For example, researchers from CSIRO developed a deep learning model for forecasting global monthly sea surface temperature anomalies associated with El Niño-Southern Oscillation (ENSO). The model was based on 70 years of climate data and accurately predicts El Niño and La Niña events 18 months in advance (Taylor and Feng, 2022). Long range predictions of these events are valuable as they improve planning and preparedness lead-times for sectors such as agriculture, healthcare, water and energy.

The next frontier could involve the creation of AI foundation models for climate, weather and natural hazard forecasting. Researchers from IBM, NASA, the University of Alabama and the Karlsruhe Institute of Technology in Germany review AI foundation models based on the transformer architecture for generalized weather and climate forecasting problems (Mukkavilli *et al*., 2023). They argue that such models could be used for 'downscaling' climate data to finer spatial resolution, predicting hurricanes and forecasting weather conditions conducive to bushfires. They also argue that the technology has progressed sufficiently to enable the creation of a fully generalizable AI foundation model for weather and climate forecasting.

Efforts to develop such models are underway. A research team from the University of California in Los Angeles, Microsoft and AI company Scaled Foundations recently published on ClimaX; a pre-trained AI model based on the transformer architecture that can be 'fine-tuned to address a breadth of climate and weather tasks, including those that involve atmospheric variables and spatio-temporal scales unseen during pretraining' (Nguyen *et al*., 2023). On 31 October 2023, the United Kingdom Met Office announced that it would be partnering with the Alan Turing Institute to build an AI model for fast and accurate prediction of significant weather events to save lives and protect infrastructure (Met Office, 2023). This project will use the UK's new supercomputer at the University of Bristol (GOV.UK, 2023b). On 26 September 2023 NASA announced research underway with Oak Ridge National Laboratory, IBM, NVIDIA and US universities to create an AI foundation model for weather and climate (Ramachandran and Lee, 2023).

Whilst models and data from other countries can be useful in Australia, they have substantial limitations. As per the UK example, Australian climate and weather modelling needs to be designed specifically for Australia's unique environmental and socio-economic conditions. An AI foundation model for weather and climate forecasting in Australia could involve the components shown below.

Datasets → **AI model development** → **Applications**

Fine scale spatial and temporal climate and weather data relating to Australia and our world region from the Bureau of Meteorology.

Other fine scale weather and climate from organisations the National Oceanic and Atmospheric Administration (NOAA) and the European Centre for Medium-Range Weather Forecasts (ECMWF).

Social, economic and environmental data relating to the impacts of severe climate and weather events in Australia (e.g. forest fuel loads, land use patterns, infrastructure ...)

Use data to train and develop a generalisable AI weather and climate foundation model designed for Australia.

Create a multi-model interface for researchers, scientists, climatologists, emergency management organisations, insurance organisations, meteorologists, media and the public.

For faster, more accurate and more spatially and temporally detailed forecasts on...

- Flood events
- Bushfire conditions
- Cyclones
- Droughts
- Heatwaves
- Hailstorms
- Electric storms
- Extreme rain and winds
- Coastal inundation
- Sea level rise.

Other examples

Due to the general-purpose nature of AI models, they can be put to productive use in practically any industry sector. We give more detailed descriptions of the first two possible applications (as above) because they've been well documented in the research literature. However, countless more AI foundation models will emerge to support wide-ranging tasks and decisions in other industries. Here we briefly describe a few more potential industry applications of AI foundation models as examples of a much longer list.

The Construction Industry – Cost estimation, checking blueprints and speeding up development assessment

AI foundation models could be used to identify structural, safety and other non-compliance issues in building plans (blueprints). They could be used for cost estimation and forecasting the risk of cost-blowouts associated with large projects. AI foundation models could be used to speed up and simplify development assessment and impact assessment. They could be used for supply chain management and material selection and optimisation. These applications would depend on feeding large volumes of Australian-sourced training data into the machine learning algorithms. The datasets could include development applications lodged with local governments, design blueprints, material testing data and project plans. A review paper on the use of Generative Pre-Trained Transformer AI foundation models for the construction industry was published last year (Saka et al., 2023).

The banking and finance industry – improving security, risk appraisal and customer service

The banking and finance sector is already making significant use of machine learning to monitor transactions and detect suspicious activity. However, there is scope to do more with AI foundation models. For example, in Italy the 'UniTTab' is using deep learning with the transformer architecture for fraud detection, predicting product churn (customers leaving the bank) and to classify the behaviour of clients (e.g. loan default predictions). The AI foundation model they're developing is multi-modal in that it handles numerical, textual, categorical and date/time data. The data they've used comes from 'an important Italian bank' and includes 1.5 billion transactions from 3 million anonymised clients during 2020 to 2022. The model was found to be 96% accurate for fraud detection, 95% accurate for loan default prediction and 91% accurate for predicting customer churn. The researchers conclude that this is a 'first step towards the creation of foundation models for transactional time series data' for the banking sector (Cucchiara, 2023; Cucchiara et al., 2023).

Improved electricity-network planning and management

The electricity network is a complex and interconnected system of generation, storage, transmission and end-users. This system produces vast amounts of data from IoT devices and sensors, and these data can be used to train AI foundation models to inform planning and operational decisions. AI foundation models could be used for more accurate demand forecasting, energy storage optimisation, grid optimisation, predictive maintenance, regulatory compliance, market analysis (price forecasting) and planning for the replacement of carbon-emitting infrastructure with renewable energy systems.

Traffic management and road safety

At federal, state/territory and local levels Australian governments hold vast quantities of information about traffic flows and road safety. AI foundation models could be trained on these data to learn about ways to reduce trip durations, avoid traffic jams and improve safety for vehicle occupants, cyclists and pedestrians. The datasets that could be useful in training AI foundation models for traffic management include traffic flow data, accident reports, road infrastructure data, public transport data, freight and logistics data, pedestrian and cyclist data and vehicle data.

And more – where there's a dataset there's a way

Wherever there are large, fit-for-purpose and quality datasets, with metadata, there's scope to develop or fine tune AI foundation models. Data is an underlying requirement for AI foundation model development. In Australia most industry sectors have these datasets; and oftentimes they're under-exploited. There's an opportunity to build AI foundation models with intuitive (natural-language) interfaces to support wide-ranging tasks and decisions in these industries. Using Australian data will make the model more useful and responsive to Australian application.

Conclusion

The rise of AI foundation models represents a paradigm shift, and new era, for the field of AI. Multi-modal foundation models can handle diverse information input and output formats including text, audio, video, imagery and robotics. This makes them powerful problem solvers in wide ranging tasks in practically all industry sectors. Against a backdrop of advancing digital technology foundation AI has the potential to transform industries, particularly those in the services-sector, and boost productivity. Companies and industries which develop and adopt foundation AI technologies may achieve a significant competitive advantage over others.

The vast bulk of AI foundation models being used by Australian industry are developed offshore. This is an excellent opportunity. These models can help workers, companies, research institutes and governments be more productive and achieve more. Access to the model interface, software or API can often be purchased cost effectively. Some models are open source and freely available.

But the productivity gains are not assured. There's much complexity about which AI tools to adopt, how to adopt them and how to ensure they're used ethically. However, dependence on overseas (or private tech-corporation) made and operated AI foundation models also creates sovereign capability risks and concerns about fair and open markets for product development. This has been observed in policy responses by the United States (Wald, 2023), United Kingdom (CMA, 2023), Japan (LDP, 2023) and Germany (AKI, 2023). Questions are being asked about:

- Investments in creating new AI foundation models, or fine-tune existing ones, to achieve social, economic and environmental benefits exceeding the costs.
- The extent to which the innovation ecosystem has a '*level playing field*' where all parties can access compute-power, datasets and skills to build/apply AI foundation models.
- Impacts on the competitive interests of workers and companies within sectors – *e.g.* consider the effect of Uber on taxi drivers/companies in Australia.
- The safety, security, reliability, cultural-fit, content-moderation and ethics of AI models developed offshore, especially those which have language and image outputs.

The same sorts of considerations are likely to apply in Australia. We're in the early phases of the AI foundation model journey. Things are moving fast. Much more powerful computers will be used to train and develop forthcoming models. More data will be fed into them, and the global research community will discover new and improved ways of making them.

This brief report provides a snapshot of AI foundation models and policy levers that can be applied to harness the benefits and mitigate the risks. We hope it will help flag some of the issues and help frame decisions by government, industry and community organisations.

References

AIHW (2023) Cancer data in Australia (last updated 31 Aug 2023). Australian Institute of Health and Welfare, Australian Government, Canberra.

AKI (2023) Large AI Models for Germany – Feasibility Study 2023. A report by Akademie für Künstliche Intelligenz AKI gGmbH at German AI Association for German Federal Ministry of Economics and Climate Protection, Berlin.

Alarcon N (2020) OpenAI Presents GPT-3, a 175 Billion Parameters Language Model. NVIDIA Developer Technical Blog, NVIDIA, Santa Clara, California.

Anderljung M, Barnhart J, Leung J et al. (2023) Frontier AI regulation: Managing emerging risks to public safety. arXiv:2307.03718v3. DOI: 10.48550/arXiv.2307.03718.

Arundel A, Kabla I. (1998) What percentage of innovations are patented? empirical estimates for European firms, Research Policy 27(2), DOI: 10.1016/S0048-7333(98)00033-X.

Attard J, Orlandi F, Scerri S, Auer S (2015). A systematic review of open government data initiatives. Government information quarterly, 32(4), 399-418.

Australian Treasury (2023) Intergenerational report 2023: Australia's future to 2063. Treasury, Australian Government, Canberra.

Blumenfield J (2023) NASA and IBM Openly Release Geospatial AI Foundation Model for NASA Earth Observation Data. Earth Data Open Access for Open Science. NASA, Washington DC.

Bommasani R, Hudson DA, Adeli E et al. (2023) On the Opportunities and Risks of Foundation Models. Center for Research on Foundation Models (CRFM), Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford University. arXiv:2108.07258v3.

Bousmalis K, Vezzani G, Rao D et al. (2023) RoboCat: A Self-Improving Foundation Agent for Robotic Manipulation. arXiv:2306.11706 [cs.RO]. DOI: 10.48550/arXiv.2306.11706.

Braue D (2023) Tech alone won't boost Australia's productivity - Intergenerational Report eyes mining to sustain long-term growth. 29 August, ICT News, Australian Computer Society, Sydney.

Brynjolfsson E, Li D, Raymond LR (2023) Generative AI at Work. National Bureau of Economic Research working paper 31161, DOI 10.3386/w31161.

Capello, R., Lenzi, C. and Perucca, G., 2022. The modern Solow paradox. In search for explanations. Structural Change and Economic Dynamics, 63, pp.166-180.

CMA (2023) AI Foundation Models: Initial Report. Competition and Markets Authority, UK Government, London.

Coelli MB, Borland J (2019) Behind the Headline Number: Why not to Rely on Frey and Osborne's Predictions of Potential Job Loss from Automation. Melbourne Institute Working Paper No. 10/19, October 2019, The University of Melbourne.

CPA (2023) Back to the future? AI uptake stalls among Asia-Pacific businesses. Media Release 29 August. Certified Practicing Accountants Australia, Canberra.

CSIRO (2023) AI at CSIRO - Applying AI to tackle Australia's greatest challenges. CSIRO website accessed 25 October 2023 www.csiro.au/ai, CSIRO, Canberra.

Czarnitzki, D., Fernández, G.P. and Rammer, C., 2023. Artificial intelligence and firm-level productivity. Journal of Economic Behavior & Organization, 211, pp.188-205.

Cucchiara, R (2023) What large language models like GPT can do for finance? Nature Comments 5 July 2023.

Cucchiara R, Garuti F, Luetto S, Sangineto E (2023) Deep Learning and Large Scale Models for Bank Transactions. 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy. Workshop proceedings (https://ceur-ws.org/).

Damioli, G., Van Roy, V. and Vertesy, D., 2021. The impact of artificial intelligence on labor productivity. Eurasian Business Review, 11, pp.1-25.

Darmawanto EN, Syahwami (2023). The Role of Artificial Intelligence in Government Services: A Systematic Literature Review. Open Access Indonesia Journal of Social Sciences. DOI:10.37275/oaijss.v6i3.163.

DCCEEW (2023) Australia's real solar potential, Media Release (19 July) by The Hon Chris Bowen MP Minister for Climate Change and Energy, Department for Climate Change Energy the Environment and Water, Australian Government Canberra.

DE (2022) University Research Commercialisation Action Plan. Department of Education, Australian Government.

Dell'Acqua F, McFowland E, Mollick E et al. (2023) Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality (September 15, 2023). Harvard Business School Technology & Operations Management Unit Working Paper No. 24-013, DOI: 10.2139/ssrn.4573321.

Dewitte S, Cornelis JP, Müller R, Munteanu A. (2021). Artificial intelligence revolutionises weather forecast, climate monitoring and decadal prediction. Remote Sensing, 13(16), 3209.

DISR (2019) Australia's AI Ethics Principles – A voluntary framework. Department of Industry Science and Resources, Australian Government, Canberra.

DISR (2024) Safe and responsible AI in Australia consultation Australian Government's interim response. Department of Industry Science and Resources, Australian Government, Canberra.

DoD (2021) Fact sheet – Sovereign industrial capability priorities. Department of Defence, Australian Government, Canberra.

Ducci F (2020). Natural monopolies in digital platform markets, Cambridge University Press, London.

Easton W (2021) Changes to Sharing and Viewing News on Facebook in Australia. Meta Newsroom 17 February 2021, Meta, Sydney.

Egan R (2022) You might think solar panels have been perfected – but we can still make them even better and cheaper. The Conversation, 24 October.

Frey CB, Osborne MA (2023) The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change. DOI: 10.1016/j.techfore.2016.08.019.

Fries J, Steinberg E, Fleming S, Wornow M, Xu Y, Morse K, Dash D, Shah N (2022) How Foundation Models Can Advance AI in Healthcare. News, Stanford University Human-Centred Artificial Intellgience, California, United States.

Gao, Y., & Janssen, M. (2020). Generating value from government data using AI: An exploratory study. In Electronic Government: 19th IFIP WG 8.5 International Conference, EGOV 2020, Linköping, Sweden, August 31–September 2, 2020, Proceedings 19 (pp. 319-331). Springer International Publishing.

Garrett K (2019). Social media's contribution to political misperceptions in US Presidential elections. PLoS ONE, DOI: 10.1371/journal.pone.0213500.

Ge Y, Wang Q, Yu J, Shen C, Li Q (2023) Data Poisoning and Backdoor Attacks on Audio Intelligence Systems. IEEE Communications Magazine. DOI: 10.1109/MCOM.012.2200596.

Gil Y, King R, Kitano H (2020) AI scientist grand challenge: Artificial Intelligence Systems Capable of Nobel-Quality Discoveries. The Turing Institute, London, United Kingdom.

Gorodnichenko Y, Pham T, Talavera O (2021) Social media, sentiment and public opinions: Evidence from #Brexit and #USElection. European Economic Review. DOI: 10.1016/j.euroecorev.2021.103772.

GOV.UK (2021) Quantifying the UK Data Skills Gap - Full report. Policy Paper published on 18 May by the UK Government Department of Education, Training and Skills, London.

GOV.UK (2023a) Initial £100 million for expert taskforce to help UK build and adopt next generation of safe AI. Press Release from the United Kingdom Government on 24 April 2023. UK Government, London.

GOV.UK (2023b) Bristol set to host UK's most powerful supercomputer to turbocharge AI innovation. Press release 13 September 2023. UK Government, London.

Hajkowicz SA, Bratanova A, Schleiger E, Naughtin E (2023) Australia's artificial intelligence ecosystem report. A CSIRO report for the National Artificial Intelligence Centre, CSIRO, Canberra.

Hajkowicz SA, Whittle J (2023) Yes, AI could help us fix the productivity slump – but it can't fix everything. The Conversation 22 June.

Hardy, K., & Maurushat, A. (2017). Opening up government data for Big Data analysis and public benefit. Computer law & security review, 33(1), 30-37.

Harman, G. and Harman, K., 2004. Governments and universities as the main drivers of enhanced Australian university research commercialisation capability. Journal of Higher Education Policy and Management, 26(2), pp.153-169.

Harrer S (2023) Attention is not all you need: the complicated case of ethically using large language models in healthcare and medicine. eBioMedicine 90, DOI: doi.org/10.1016/j.ebiom.2023.104512.

Harris M (2014)The STEM shortage paradox. Physics World 27(10). DOI: 10.1088/2058-7058/27/10/40.

Ho E, King K, Wald R, Wan C (2021) Building a National AI Research Resource: A Blueprint for the national Research Cloud, Human Centred Artificial Intelligence, Stanford Law School, Stanford University, California.

Hui X, Reshef O, Zhou L (2023) The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market (July 31, 2023). Available at SSRN: https://ssrn.com/abstract=4527336 or http://dx.doi.org/10.2139/ssrn.4527336.

Hunter B, Hindocha S, Lee R (2022). The Role of Artificial Intelligence in Early Cancer Diagnosis. Cancers (Basel) DOI: 10.3390/cancers14061524.

Judd B (2021) Google is threatening to pull its search engine from Australia. So what does that mean for you? ABC News 25 January 2021.

Kalliamvakou, E (2022). Quantifying GitHub Copilot's impact on code quality. GitHub Blog 10 October.

Kelly T, Nussey S (2023) Japan needs computing power surge to stay in AI race, says government adviser. Reuters Technology News 20 July 2023.

Knight W (2023) OpenAI's CEO Says the Age of Giant AI Models Is Already Over. Wired 17 April.

Kutar M, Fletcher G (2023) We asked ChatGPT to write a company HR policy – and the results were promising. The Conversation 19 June.

LDP (2023) The AI White Paper: Japan's National Strategy in the New Era of AI. Liberal Democratic Party of Japan, Tokyo.

Martin P, Manson K (2023) CIA Builds Its Own Artificial Intelligence Tool in Rivalry With China. Bloomberg 27 September 2023.

McIntosh D (2018). We need to talk about data: how digital monopolies arise and why they have power and influence. Journal of Technology Law & Policy, 23, 185.

Met Office (2023) AI to take weather forecasting by storm. Press Release 31 October 2023. UK Met Office, United Kingdom Government, London.

Microsoft (2023) Annual report, shareholder letter (www.microsoft.com/investor/reports/ar23/).

Microsoft and the Tech Council of Australia (2023) Australia's Generative AI opportunity. A joint report by Microsoft and the Tech Council of Australia, Microsoft, Sydney.

Mukkavilli S K, Civitarese DS, Schmude J, Jakubik J, Jones A, Nguyen N, et al. (2023). AI Foundation Models for Weather and Climate: Applications, Design, and Implementation. arXiv preprint arXiv:2309.10808. DOI: 10.48550/arXiv.2309.10808.

Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J., Topol, E.J. and Rajpurkar, P., 2023. Foundation models for generalist medical artificial intelligence. Nature, 616(7956), pp.259-265.

Narechania TN (2021). Machine Learning as Natural Monopoly. Iowa Law Review 107, 1543.

Nguyen K, Hambur J (2023) Adoption of General-purpose Technologies (GPT) in Australia: The Role of Skills. Septemner Bulletin, Reserve Bank of Australia, Melbourne.

Nguyen T, Brandstetter J, Kapoor A, Gupta JK, Grover A (2023). ClimaX: A foundation model for weather and climate. arXiv preprint arXiv:2301.10343.

Noy, S and Zhang, W (2023) Experimental Evidence on the productivity effects of productivity and happiness. Github.

NSF (2024) Democratizing the future of AI R&D: NSF to launch National AI Research Resource pilot. National Science Foundation Press Release (24 January), United States Government, Washington DC.

OECD (2023a) Initial policy considerations for generative artificial intelligence, OECD AI Papers No. 1 September 2023. Organisation for Economic Cooperation and Development, Paris.

OECD (2023b) A blueprint for building national compute capacity for artificial intelligence. OECD digital economy papers, February 2023 No. 350. Organisation for Economic Cooperation and Development, Paris.

PC (2024) Making the most of the AI opportunity: Research paper 1, AI uptake, productivity, and the role of government. Productivity Commission, Australian Government, Canberra.

Petrović, N., Milić, P., & Prlinčević, B. (2022). Using Open Government Data for Economic Development. The European Journal of Applied Economics, 19(2), 129-141.

Ramachandran R, Lee T (2023) Building an AI Foundation Model for Weather and Climate. NASA Earth Data News, 26 September 2023, NASA, United States Government, Houston.

Sadler D (2023) Nearly 70% of Aussie IT jobs are in shortage - We're facing a skills challenge not seen in 60 years. Australian Computer Society ICT News 5 October 2023, Sydney.

Saka A, Taiwo R, Saka N, Salami BA, Ajayi S, Akande K, Kazemi H (2023) GPT models in construction industry: Opportunities, limitations, and a use case validation. Developments in the Built Environment. DOI: 10.1016/j.dibe.2023.100300.

SAS (2022) How to solve the data science skills shortage. A report for SAS prepared by Coleman Parkes Research, London, UK.

Scott IA, Crock C (2020) Diagnostic error: incidence, impacts, causes and preventive strategies. Medical Journal of Australia 213 (7). DOI: 10.5694/mja2.50771.

Stevens R (2020) AI for Science: Report on the Department of Energy (DOE) Town Halls on Artificial Intelligence for Science. Argonne National Lab. (ANL), Argonne, Illinois, United States. DOI: 10.2172/1604756.

Stokel-Walker C (2022) Twitter changed science — what happens now it's in turmoil? Nature News Feature (20 December).

Taylor J, Feng M. (2022). A deep learning model for forecasting global monthly mean sea surface temperature anomalies. Frontiers in Climate, 4, 178.

Tonkin C (2023) Westpac says AI boosted coding productivity 46 per cent increase in productivity, no drop in code quality. ICT News (6 June 2023), Australian Computer Society, Sydney.

Villa J, Troiano D(2020) Choosing Your Deep Learning Infrastructure; The Cloud vs. On-Prem Debate, Determined AI (30 July 30), Determined AI, Retrieved on 23 October 2023 from: https://www.determined.ai/blog/cloud-v-onprem.

Vipra J, Korinek A (2023) Market concentration implications of foundation models: The invisible hand of chatgpt. Brookings Institute Center on Regulation and Markets Working Paper #9, Washington DC.

Watt M (2003) The Commercialisation of Photovoltaics Research in Australia: A report for Science and Innovation Mapping Department of Education Science and Training. Centre for PV Engineering, University of NSW, Sydney.

White House (2023) Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (30 October 2023), The White House, United States Government.

Wood FQ (1992) The commercialisation of university research in Australia: issues and problems. Comparative Education, 28(3), pp.293-313.

Yang, C.H., 2022. How artificial intelligence technology affects productivity and employment: firm-level evidence from Taiwan. Research Policy, 51(6), p.104536.

Zeleti FA, Ojo A, Curry E (2016). Exploring the economic value of open government data. Government Information Quarterly, 33(3), 535-551.

Zhang D, Lawrence C, Sellitto M et al. (2021) Enhancing International Cooperation in AI Research: The Case for a Multilateral AI Research Institute. Human Centred Artificial Intelligence, Stanford University.

Zou A, Wang Z, Kolter JZ, Fredrikson M (2023) Universal and Transferable Adversarial Attacks on Aligned Language Models. arXiv:2307.15043. DOI: 10.48550/arXiv.2307.15043.

As Australia's national science agency, CSIRO is solving the greatest challenges through innovative science and technology.

CSIRO. Creating a better future for everyone.

Contact us
1300 363 400
csiro.au/contact
csiro.au

For further information
Data61

Stefan Hajkowicz
+61 7 3833 5540
stefan.hajkowicz@data61.csiro.au

Elanor Huntington
+61 2 6218 3458
elanor.huntington@csiro.au

B&M | 24-00012