

An Artificial Intelligence Agenda for New Zealand

lessons from the United States

Abstract

With recent advances in artificial intelligence (AI), policymakers are looking for tools to promote the associated benefits, as well as mitigate the attendant risks. Thanks to extensive experimentation by international partners, New Zealand is positioned to parse the lessons of those international policies and more quickly adapt those that are successful to its own context. This article identifies several key AI policies pursued by the United States that may be relevant in a New Zealand context and discusses broad lessons from the US experience to help inform successful AI policymaking in New Zealand. It also emphasises the unique bipartisan nature of AI policymaking in the United States to date, and the benefits of a consensus-building approach to AI policy.

Keywords artificial intelligence, ChatGPT, United States, international cooperation, data, public sector, risk management

Sam Mulopulos was a 2023 Ian Axford Fellow in Public Policy. He previously was the deputy staff director of the Homeland Security and Governmental Affairs Committee in the US Senate and the lead staffer for the bipartisan Senate Artificial Intelligence Caucus. He has also taught a course on AI ethics, law and policy as an adjunct professor at Georgetown University. The views expressed in this article are his own.

It is practically axiomatic that artificial intelligence (AI) will be one of history's most transformative technologies. With the launch of ChatGPT, and the proliferation of large language models (LLMs), widespread interest in AI has been piqued like never before. Yet, while AI stands to offer incredible benefits to society, it also poses unique and unprecedented risks. How can society wring those benefits – such as improved medical diagnosis – while simultaneously mitigating the risks, like algorithmic bias? This is the key question facing policymakers. Answering it in the rapidly changing current environment requires policymakers to be cognisant of the uses, and as well as the limitations, of existing tools.

Reducing the costs and maximising the benefits of AI systems is, at its core, a technological feat achieved through consistent and iterative research, testing and innovation. Around that 'core' exists public policy, which complements those

activities in two ways. First, AI policy should promote a social, economic and political context in which these technological feats are achieved easily and consistently, and then, when achieved, faithfully implemented. Second, policy should set legal constraints to mitigate harmful AI practices.

Since 2018, the United States has endeavored to complement AI innovation by enacting a multitude of AI-related policies (Institute for Human-Centered Artificial Intelligence, 2022, p.271). This includes a substantive government commission on AI, administrative guidance to industry, and new laws on topics like the regulation of the government's use of AI and the democratisation of access to AI

AI Caucus passed more than a dozen AI-related bills into law (Vincent, 2021).

In a recent nod to that success, members of Parliament in New Zealand formed the first cross-party AI caucus to bring multi-party perspectives to bear on AI policy. Led by Judith Collins, the cross-party AI caucus intends to serve as a forum to educate lawmakers about AI-related issues and promote similar cross-party collaboration in the policy process (Griffin, 2023). As New Zealand's lawmakers consider how to answer the key question posed by the proliferation of AI systems, the success and variety of AI-related policy initiatives in the United States may serve as a guide or source of inspiration for policymakers in Aotearoa.

As New Zealand's lawmakers consider how to answer the key question posed by the proliferation of AI systems, the success and variety of AI-related policy initiatives in the United States may serve as a guide or source of inspiration for policymakers in Aotearoa.

research tools. The development and adoption of these policies have effectively positioned the United States to tackle the AI moment.

Key to this success was a concerted effort in the US Congress to develop AI policy in a bipartisan way. Both the US Senate and US House of Representatives have bipartisan AI caucuses, which were created to convene legislators of both parties to educate policymakers about AI, and to collaboratively create and pass new AI laws. By bringing together various perspectives, the AI caucuses have challenged notions of legislative lethargy and built a high degree of consensus within the legislative branch around AI policy. The results of this emphasis on bipartisan consensus building are evident: in its first three years, from 2019 to 2022, the Senate

AI policy in New Zealand

As a general matter, New Zealand stands to punch significantly above its weight when it comes to AI. Between 2016 and 2021, New Zealand led the world with the greatest growth in AI-related hiring (Institute for Human-Centered Artificial Intelligence, 2022, p.143). Granted, that growth starts from a low baseline, but it is explosive and exciting, nonetheless. In 2017, New Zealand was ranked ninth (of 35) OECD countries for 'government AI readiness' (AI Forum New Zealand, 2019). This is even more impressive given that New Zealand did not release an AI strategy until 2021 (Ponti, 2021). Today, the University of Waikato, Auckland University of Technology and the University of Otago all have their own AI research centres (Hope, 2021; Auckland University of Technology,

n.d.; University of Otago, n.d.–b).

However, New Zealand's approach to AI is not without challenges. One of these challenges is access to talent. Managing the benefits with the risks of AI requires innovation, which itself requires a pool of talented and creative AI scientists. However, recent research found that of 22,000 AI experts worldwide, only 85 are in New Zealand (Gagne, 2018).

Another challenge is a lack of uptake: according to the most recent available data, from 2018, only 36% of New Zealand firms said their boards were thinking about AI (AI Forum New Zealand, 2018, p.79). Because of New Zealand's small scale, delayed deployment of AI where there is a strong business case risks a chicken-and-egg spiral. Attenuated deployment results in even less data, and of a poorer quality to boot. This means AI systems will fail to live up to their useful potential. When deciding to invest in new technology – whether systems produced domestically or purchased from abroad – businesses will steer clear if, as a result of the issues mentioned above, they find AI systems unhelpful; in that case the economy overall might not achieve the efficiency and productivity gains AI offers. Both of these challenges are examples of problems that public policy can help solve by investing in training and education for new talent and setting the contours of the market in a way which promotes responsible acquisition and deployment of AI systems.

In response to New Zealand's AI challenges, Boyd and Wilson (2017) argue that New Zealand should tackle AI issues by: researching the risk and impact of AI; informing and engaging the public; producing clear recommendations; and taking a global lead. In some of these areas there has been action. The AI Forum has led efforts to inform, educate and research. The New Zealand Institute of Directors produced a survey of emerging AI challenges (Institute of Directors and Chapman Tripp, 2016). And the New Zealand Law Foundation has funded a three-year project at the University of Otago, 'Artificial Intelligence and Law in New Zealand', to study the impacts of AI on employment and the administration of criminal justice (University of Otago, n.d.–a). But efforts to develop specific and

granular recommendations appear scant. Doing so requires that ‘existing policy ... be analysed, international policy co-opted as appropriate, and new policy ... developed’ (Boyd and Wilson, 2017, p.14). This article responds to that specific charge by enumerating some key successes of US AI policy and offers lessons for New Zealand. However, it is important to note that there are significant differences between the United States and New Zealand, especially in terms of the size and extent of available resources. In sharing the US AI policy experience, and suggesting lessons for New Zealand, this article recognises those differences. It suggests that New Zealand may be better positioned to adapt US AI policy to its own ends, rather than adopt it wholesale.

AI policy in the United States

The United States has proposed, enacted or implemented a cornucopia of AI laws and policies over the past several years. What follows are six policies which are attempts to answer, in their own ways, the question posed in the introduction about finding the good, and mitigating the bad, with respect to AI.

AI Risk Management Framework

In 2023, the National Institute for Standards and Technology (NIST) – since 1901 the United States’ federal standards and measurement entity – released the AI Risk Management Framework (RMF). The culmination of a lengthy, iterative and multi-stakeholder process, the RMF is a well-respected document which organisations can adopt to manage the risks involved in using AI systems (Kruger and Lee, 2023). While voluntary, the beauty of the RMF is that it is ‘rights-preserving, non-sector specific, and use-case agnostic’ (National Institute of Standards and Technology, 2023). The RMF achieves this by offering a model for recognising AI risks, and articulating the attributes of trustworthy AI (e.g., safe; secure and resilient; explainable and interpretable; privacy-enhanced; and fair, with harmful bias managed) (ibid.).

Alongside this work, the RMF provides procedures for organisations to operationalise the conceptual model. This includes testing, evaluation, verification

and validation processes. The RMF is now the third in a series of similar tools released by NIST, and if history is any guide, the RMF will find widespread adoption, even outside the United States. NIST’s Cybersecurity Framework and their Privacy Framework have been adopted by state and federal agencies, US firms, and, in a testament to the quality consensus they enjoy, foreign entities like the Bank of England, Siemens and Nippon Telephone & Telegraph (Kerry, 2023).

National Security Commission on Artificial Intelligence

Where the RMF is a culmination of the first wave of US AI policymaking, the origins of this policymaking lie back in 2018, when

Although at its core focused on national security, the NSCAI was not shy about branching out into other AI-related topics. Its work laid significant groundwork for how policymakers think about the opportunities and challenges raised by the use of AI, especially the ways in which policymakers can best align AI with the values of open societies. Because the commission’s proposals were detailed enough to specify not just what the policy ought to be, but *how* it would work in an administrative sense, it was easy for Congress to take many of the NSCAI’s proposals and swiftly turn them into law. The NSCAI is a testament to the success of short-term, hyper-detailed and government-backed efforts to convene

By democratising access ... tools beyond large technology firms and the national labs, more researchers will be able to contribute to the innovations which will improve the quality of AI systems and reduce their associated risks

the US Congress passed legislation to establish a new independent commission ‘to consider the methods and means necessary to advance the development of artificial intelligence, machine learning, and associated technologies to comprehensively address the national security and defense needs of the United States’ (National Security Commission on Artificial Intelligence, n.d.). Chaired by former Google chairman Eric Schmidt and former deputy secretary of defense Robert Work, the National Security Commission on Artificial Intelligence (NSCAI) was a uniquely successful government commission. Despite being a temporary commission, in its few years of operation the NSCAI produced a number of exceptionally granular reports proposing dozens of new policies related to the AI aspects of national security and US strategic competitiveness.

experts to propose solutions to specific emerging technology problems.

National Artificial Intelligence Research Resource

A good example of the NSCAI’s lasting success is the creation of the National Artificial Intelligence Research Resource (NAIRR). Originally proposed by the NSCAI, and championed by the Senate AI Caucus, the first step to create the NAIRR was signed into law on New Year’s Day 2021. The NAIRR is envisioned as a cyberinfrastructure of AI research tools available to AI scientists across the United States (National Artificial Intelligence Research Resource Task Force, n.d.). This new cyberinfrastructure will include shared access to computational resources, test beds, quality data sets, and other tools researchers need to pursue AI innovations. By democratising access to these tools

beyond large technology firms and the national labs, more researchers will be able to contribute to the innovations which will improve the quality of AI systems and reduce their associated risks (Office of Congresswoman Anna Eshoo, 2022).

Pursuant to the underlying law, the National Science Foundation (NSF) and the White House Office of Science and Technology Policy convened a task force of experts from government, industry and academia to develop a detailed roadmap for creating the NAIRR. This roadmap involves parsing the ways in which

developing trustworthy AI (National Artificial Intelligence Initiative Office, n.d.).

This includes harnessing trustworthy AI research efforts across government into a cohesive whole. For example, the Initiative prioritises the work being done pursuant to a five-year, \$360 million investment in 18 new national AI research institutes across the United States (National Science Foundation, 2023). It also includes the research funded by the NSF's Program on Fairness in AI, which partners with industry to support interdisciplinary work on AI transparency, explainability and

The [Artificial Intelligence Training for the Acquisition Workforce Act in 2022] requires US procurement officers to be trained in the capabilities and risks associated with AI.

government can create the NAIRR administratively, and where additional legal authorities and funding are needed from Congress (National Artificial Intelligence Research Resource Task Force, 2023). Congress is now considering these findings with an eye to establishment of the NAIRR.

The National Artificial Intelligence Initiative

In the same legislation which charted a path towards the NAIRR, Congress also established another AI Caucus priority, the National Artificial Intelligence Initiative. The law codifies, and expands upon, the American AI Initiative first launched by President Donald Trump in 2019, a clear indication that the United States has treated AI as a presidential priority across administrations of different parties (Parker, 2020). Seeking to put the federal AI research house in order, the National Initiative organises the AI research and development (R&D) enterprise in the United States to be coordinated, strategic, and focused on building talent and

accountability (National Science Foundation, 2021). The law also established a new National Artificial Intelligence Initiative Office in the White House to serve the 'technical, programmatic, and administrative' needs of the Initiative, as well as new advisory committees to formally incorporate stakeholder engagement (Harris, 2021).

AI in Government Act

Alongside the focus on AI research and innovation, the United States has also begun to explore AI regulation. Enacted in 2021 as a Senate AI Caucus initiative, the AI in Government Act sets standards for the federal government's own use of AI. A case of leading by example, the AI in Government Act signals to companies and countries that the United States is serious about creating sensible rules for AI systems in order to reap their benefits and reduce their risks.

The law requires the White House, through the Office of Management and Budget (OMB), to issue each agency strategic guidance for how the agency can

best adopt safe and trustworthy AI. This includes recommendations for ways agencies can remove barriers to the adoption of trustworthy AI, and best practices for addressing any bias or algorithmic unfairness (Schatz, 2019). To promote accountability, the Act requires agencies to respond with their own public plans detailing how the agency will comply with the guidance and the law (ibid.). Because technical AI expertise can vary greatly across government agencies, the law establishes an AI Center of Excellence within government to act as a kind of 'help desk' for agencies as they deploy and manage AI systems (Senate Homeland Security and Governmental Affairs Committee, 2020).

Unfortunately, implementation of the Act has been uneven across government. The Center of Excellence has been stood up, and some agencies have taken serious steps to use AI responsibly to deliver services to citizens. But, at present, OMB has not issued the required guidance to all agencies, creating a patchwork of AI usage across the US government (Senate Homeland Security and Governmental Affairs Committee, 2022). However, this unevenness can be expected to be smoothed when OMB eventually issues the required guidance.

Artificial Intelligence Training for the Acquisition Workforce Act

Since human talent is critical to ensure the responsible use of AI systems, Congress passed the Artificial Intelligence Training for the Acquisition Workforce Act in 2022. The law requires US procurement officers to be trained in the capabilities and risks associated with AI. Since the US government is a major purchaser of AI systems, procurement officials must know what to look for when making a purchase. The training required includes analysis of the benefits and risks of AI, the science behind the technology, and how AI can be made safe and trustworthy (Senate Homeland Security and Governmental Affairs Committee, 2021).

Without this workforce training, agencies may come to rely heavily on government contractors to make decisions about AI acquisition and use. This reliance risks atrophying an agency's in-house

expertise and may render the agency overly dependent on outside knowledge to make procurement decisions (ibid.). Relatedly, the United States has passed other laws to promote training and development of the AI workforce, mostly in a military and national security context (Heinrich, 2020). The Artificial Intelligence Training for the Acquisition Workforce Act stands out for its efforts to ensure AI expertise across governmental agencies.

Informing an AI agenda for New Zealand

With this survey of US AI policy accomplishments in hand, what lessons can be extracted to inform a New Zealand AI agenda? Of course, the two countries are different and not every policy adopted by the United States would be relevant or appropriate. The United States' significantly larger fiscal resources and expansive civil service mean it will be difficult, and in some cases inappropriate, for New Zealand to merely *adopt* US policies. Rather, New Zealand is better positioned to *adapt* US policies for its own context, drawing broad inspiration and thinking from abroad for its own AI efforts instead of rote application. For example, the non-regulatory NIST AI Risk Management Framework may be more appropriate for adoption given that firms outside the United States already comply with other NIST frameworks because of their quality. On the other hand, New Zealand may be better positioned to adapt the United States' work on AI training for the acquisition workforce for its own ends, recognising the differences in size and methods of each country's procurement systems.

However, there are still general themes which are relevant for New Zealand. These are the importance of accurate statutory definitions, the need for public sector leadership on AI, and the use of expert commissions to drive substantive change.

The need for accurate definitions

A good definition is critical to the success of any legislation. A broad definition over-includes categories of objects or behaviours which policymakers do not intend to cover, just as a narrow definition is under-inclusive. With technology, the challenge is even more difficult. Given the slow-

moving nature of government, law and policy always risk falling behind the pace of technological development. Definitions of terms like 'artificial intelligence', 'artificial intelligence system', 'deepfake', 'technical standards' and 'explainability' all must be precise enough to cover the technological concept as it exists but be flexible enough to cover some degree of future change.

Given the effort US policymakers have put into developing robust, realistic and accurate definitions, policymakers in New Zealand can utilise their work when developing law and policy of their own.

its safe use. Imagine if the government set regulations for the private use of aircraft but did not hold itself to the same standard for its own use of aeroplanes? As policymakers design private sector regulatory frameworks for artificial intelligence, they can build confidence in, and experiment with, rules by applying those rules to the public sector first. This is the key lesson of the AI in Government Act, which has set rules for the government's use of AI well before questions of private sector regulation were on the table. Now that private sector regulation is a topic of

As policymakers design private sector regulatory frameworks for artificial intelligence, they can build confidence in, and experiment with, rules by applying those rules to the public sector first.

This is not to say that New Zealand must wholesale import every AI-related definition written by the United States, but that these definitions serve as a well-vetted source of inspiration. Moreover, since AI has been deployed in a multitude of ways and means in the United States, policymakers can be sure that these definitions cover the gamut of use cases and technological specifications. Drawing inspiration from existing definitions also ensures that policymakers do not need to constantly reinvent the wheel when it comes to the basic building blocks of lawmaking.

The importance of public sector AI leadership

The public sector stands to be uniquely effective at accelerating the deployment of responsible AI. Most fundamentally, it is easier for the government to regulate itself than it is to craft regulation – with its plethora of use cases and stakeholders – for private industry. Moreover, trust in AI – as with any technology – is best built when the government can be a leader in

public debate, policymakers benefit from the regulatory test drive offered by the law.

Another example of the public sector's AI leadership potential is its ability to guide adoption of the technology for high-value use cases. This is not to say that private industry cannot also play this role; rather, that the public sector should be bold enough to assume that it (especially when working *with* industry) can also do innovative things with respect to technology. The public sector is ripe with use cases where AI can greatly improve the delivery of services to citizens. Deploying responsible AI to solve those problems offers dual benefit too – greater societal trust in an emerging technology application and in the public sector itself. One high-value use case is in the healthcare sector in New Zealand, where the public sector can leverage its data to improve the speed and quality of care (Callaghan Innovation, n.d.).

By setting standard rules, R&D priorities and talent training across government, the United States has provided the tools to incentivise agency leadership

to take the initiative when it comes to public sector use of AI. The New Zealand system might benefit from a similar approach, albeit scaled appropriately to the New Zealand context. A New Zealand approach might even go beyond this to specify the high-value AI projects deserving of time, funds and institutional resources. Nothing breeds success like success; enumerating and focusing attention on the deployment of AI systems in these high-value public sector use cases will only make it easier to accelerate the adoption of responsible AI across government, and the economy.

Australia, the European Union, Japan, Singapore and the United Kingdom have all been active with respect to AI policy, and research into their experiences would further hone what policies and initiatives might be most valuable for a New Zealand AI agenda.

Lastly, increased R&D investments are another example of public sector AI leadership. Like the National Artificial Intelligence Initiative and the NAIRR, public investments in R&D both promote continued improvement of AI systems and support the local talent essential for AI's success. Already Callaghan Innovation's technology incubator programme helps domestic tech startups access up to \$1 million in funding with an eye to developing New Zealand tech talent (Callaghan Innovation, 2021). Increasing the amount of financing or expanding the number of startups that can participate may be one way to leverage public funding for domestic AI talent training and retention. Further research should be conducted to understand the effectiveness of the incubator programme, and similar initiatives, as well as the role they play in developing and retaining domestic tech talent. One reform, as the AI Forum notes,

is that New Zealand's technology R&D efforts would benefit from a 'centrally coordinated framework to encourage open, collaborative AI research' (AI Forum New Zealand, 2018, p.40). Given the recentness of US efforts to do the same via the National Artificial Intelligence Initiative, New Zealand policymakers may be keen to see what elements of this are translatable to New Zealand's smaller scale while still retaining value.

The value of expert commissions

Given the volume of NSCAI recommendations which have been

enacted as law or administratively adopted as policy, the use of an independent, expert commission to drive substantive AI policy development has been a hallmark of the United States' AI policy success. Expert commissions make up for the shortfalls of in-house public sector AI expertise and can ensure a variety of stakeholder viewpoints. The NSCAI, for example, comprised commissioners from government, industry and academia, and a staff of experienced civil servants. The commission's recommendations were also exceptionally detailed, making it easy to convert them into law and policy.

Importantly, the NSCAI was also time limited. Rather than establishing an indefinite new bureaucracy, Congress wanted a commission to investigate AI challenges over the near-term from 2018 to 2021 (National Security Commission on Artificial Intelligence, n.d.). The short duration had a tangible benefit. By making

the commission's work a sprint, policymakers were more inclined to respond quickly to its recommendations, rather than let them languish in a long queue of other proposed bills.

Like the United States with the NSCAI, and even NIST, New Zealand has already empowered existing actors to study AI problems. The Productivity Commission has done serious work on the public's perception of automation, finding that New Zealanders earning less than \$50,000 per annum had more negative views towards automation – seeing it as a threat to jobs – than those earning over \$50,000 (Heatley, 2020). The Productivity Commission has also investigated the labour market impacts of new technologies like AI (Productivity Commission, 2020). And the AI Forum has released a swathe of recommendations, though its report on the matter may now be dated given recent advancements in AI technology (AI Forum New Zealand, 2019).

One AI Forum recommendation is for the establishment of an AI ethics and society working group. Policymakers seeking to focus expertise to concretely solve AI problems in New Zealand may want to create such a working group. Giving it a short duration, like the remit of the NSCAI, may increase the impact of the working group. This would ensure that such efforts augment, rather than supersede, existing independent investigation and policy development efforts, such as those of the Productivity Commission and the AI Forum. It is also worth noting that four years ago – the AI Forum recommendations were released in 2019 – is ancient history in AI terms. Recommendations can quickly grow stale, or not fit for purpose, in the presently fast-changing environment. This suggests that policymakers should balance a short duration working group with the need for dynamism. Rather than letting recommendations become dated as snapshots in time, regularising limited duration working groups every few years may be one way to keep ideas fresh and relevant while also ensuring that those new ideas are taken seriously and implemented.

Conclusion

By their nature, law and policy move more slowly than the pace of technological

change. But for small, advanced economies, like New Zealand, this lag can be an advantage. As other countries have rushed with responses to the AI moment, New Zealand can benefit from their experimentation by adopting the policies that work and discarding the approaches that have been less successful.

Ultimately, the challenges facing Washington and Wellington are not all that different. How the United States has responded so far – whether it be coordinated R&D efforts, government leadership on new AI rules, or in-house talent development – stands to offer lessons, inspiration and guidance for policymakers in New Zealand. By enumerating some of the most significant US responses, and identifying common attributes of successful AI policymaking, this article hopes to be a resource for New Zealand policymakers interested in creative and comparative approaches to AI.

But the United States is only a single open society. Scholars and practitioners should conduct additional research into the experiences of other open societies to pull out similar lessons from their AI work. Australia, the European Union, Japan, Singapore and the United Kingdom have all been active with respect to AI policy, and research into their experiences would further hone what policies and initiatives might be most valuable for a New Zealand AI agenda.

Given that the artificial intelligence field is fast-moving, and the technology stands to bring tremendous change, there is an urgency for national governments to respond in ways which are thoughtful and action oriented. By learning from the experiences and efforts of international partners, policymakers can ingest myriad ideas and adapt them for their own purposes. This avoids wasting resources by reinventing the wheel of AI policy, while

also helping them find policies that are good enough for the present moment without having to spend years, perhaps fruitlessly, designing the perfect policy on their own.

If perhaps a latecomer to the AI moment, New Zealand has an exciting AI trajectory, especially given the multi-party attention that legislators have begun to give to the issue. As a result, New Zealand is well positioned to cherry-pick the AI policies from around the world to design responses which are most effective for its size and particular situation. And just like in the United States, maintaining a consensus-building posture with respect to AI will be key for New Zealand to successfully tackle the AI moment. By looking at one open society – the United States – and what its AI policy experience might offer New Zealand, this article intends to act as a catalyst for that effort.

References

- AI Forum New Zealand (2018) *Artificial Intelligence: shaping a future New Zealand*, Auckland: AI Forum New Zealand, https://aiforum.org.nz/wp-content/uploads/2018/07/AI-Report-2018_web-version.pdf
- AI Forum New Zealand (2019) *Towards Our Intelligence Future: an AI roadmap for New Zealand*, https://aiforum.org.nz/wp-content/uploads/2019/09/Towards-our-Intelligent-Future_v1.01.pdf
- Auckland University of Technology (n.d.) 'Centre for Artificial Intelligence Research', <https://www.aut.ac.nz/study/study-options/engineering-computer-and-mathematical-sciences/research/centre-for-artificial-intelligence-research-cair>
- Boyd, M. and N. Wilson (2017) 'Artificial intelligence: how might the New Zealand government respond', *Policy Quarterly*, 13 (4), pp.36–43
- Callaghan Innovation (2021) 'A vision for AI', <https://www.callaghaninnovation.govt.nz/customer-stories/vision-ai>
- Callaghan Innovation (n.d.) *Thinking Ahead: innovation through artificial intelligence*, Wellington: Callaghan Innovation, <https://www.callaghaninnovation.govt.nz/sites/all/files/ai-whitepaper.pdf>
- Gagne, J. (2018) 'Global AI talent pool report', <https://jfgagne.com/talent/>
- Griffin, P. (2023) 'Our politicians need to wake up to the power of AI', *Business Desk*, 4 May, <https://businessdesk.co.nz/article/opinion/our-politicians-need-to-wake-up-to-the-power-of-aim>
- Harris, L. (2021) *Artificial Intelligence: background, selected issues, and policy considerations*, Washington, DC: Congressional Research Service
- Heatley, D. (2020) *New Zealanders' Attitudes towards Robots and AI*, research note 2020/1, Wellington: New Zealand Productivity Commission
- Heinrich, M. (2020) 'Heinrich, Portman introduce bipartisan legislation to strengthen Department of Defense's artificial intelligence capacity', press release, 17 June, <https://www.heinrich.senate.gov/newsroom/press-releases/heinrich-portman-introduce-bipartisan-legislation-to-strengthen-department-of-defenses-artificial-intelligence-capacity>
- Hope, S. (2021) 'Waikato University launches new artificial intelligence research institute', Stuff, 28 April, <https://www.stuff.co.nz/waikato-times/news/300287255/waikato-university-launches-new-artificial-intelligence-research-institute>
- Institute for Human-Centered Artificial Intelligence (2022), Palo Alto: Stanford University
- Institute of Directors and Chapman Tripp (2016) 'Determining our future: artificial intelligence opportunities and challenges for New Zealand: a call to action', Auckland: Institute of Directors and Chapman Tripp
- Kerry, C. (2023) 'NIST's AI Risk Management Framework plants a flag in the AI debate', Brookings Institution, 15 February, <https://www.brookings.edu/articles/nists-ai-risk-management-framework-plants-a-flag-in-the-ai-debate/>
- Kruger, L. and M. Lee (2023) 'What you need to know about NIST's AI Risk Management Framework published in January 2023', Deloitte, 2 February, <https://ukfinancialservicesinsights.deloitte.com/post/102i6q5/what-you-need-to-know-about-nists-ai-risk-management-framework-published-in-janu>
- National Artificial Intelligence Initiative Office (n.d.) 'National Artificial Intelligence Initiative', <https://www.ai.gov/#NAII-NATIONAL-ARTIFICIAL-INTELLIGENCE-INITIATIVE>
- National Artificial Intelligence Research Resource Task Force (n.d.) 'About the task force', <https://www.ai.gov/nairrtf/>

- National Artificial Intelligence Research Resource Task Force (2023) *Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem: an implementation plan for a national artificial intelligence research resource*, Washington, DC: National Artificial Intelligence Research Resource Task Force
- National Institute of Standards and Technology (2023) *Artificial Intelligence Risk Management Framework*, Washington, DC: Department of Commerce
- National Science Foundation (2023) 'NSF announces seven new national artificial intelligence research institutes', press release, 4 May, <https://www.nsf.gov/cise/ai.jsp>
- National Science Foundation (2021) 'NSF program on fairness in artificial intelligence in collaboration with Amazon', <https://new.nsf.gov/funding/opportunities/nsf-program-fairness-artificial-intelligence>
- National Security Commission on Artificial Intelligence (n.d.) 'FAQs', <https://www.nscai.gov/about/faq/>
- Office of Congresswoman Anna Eshoo (2022) 'Eshoo, Gonzalez, Portman, Heinrich send bipartisan, bicameral letter supporting National AI Research Resource', press release, 27 January, <https://eshoo.house.gov/media/press-releases/eshoo-gonzalez-portman-heinrich-send-bipartisan-bicameral-letter-supporting>
- Parker, L. (2020) 'The American AI Initiative: the U.S. strategy for leadership in artificial intelligence', Organization for Economic Cooperation and Development, <https://oecd.ai/en/wonk/the-american-ai-initiative-the-u-s-strategy-for-leadership-in-artificial-intelligence>
- Ponti, J. (2021) 'New Zealand introduces national AI strategy', OpenGov, 18 May, <https://opengovasia.com/new-zealand-introduces-national-ai-strategy/>
- Productivity Commission (2020) *Technological Change and the Future of Work*, Wellington: New Zealand Productivity Commission
- Schatz, B. (2019) 'S.1363 – AI in Government Act', <https://www.congress.gov/bill/116th-congress/senate-bill/1363/text>
- Senate Homeland Security and Governmental Affairs Committee (2020) 'S.Rept. 116-226 – AI in Government Act of 2019', Washington, DC: Senate Homeland Security and Governmental Affairs Committee
- Senate Homeland Security and Governmental Affairs Committee (2021) 'S.Rept 117-40 – Artificial Intelligence Training for the Acquisition Workforce Act', Washington, DC: Senate Homeland Security and Governmental Affairs Committee, <https://www.congress.gov/congressional-report/117th-congress/senate-report/40/1?s=3&r=4&q=%7B%22search%22%3A%5B%22AI%22%2C%22Training%22%2C%22Act%22%5D%7D>
- Senate Homeland Security and Governmental Affairs Committee (2022) 'Portman presses OMB on implementation of AI in Government Act', press release, 22 December, https://www.hsgac.senate.gov/media/press/portman-presses-omb-on-implementation-of-ai-in-government-act_-/
- University of Otago (n.d.-a) 'Artificial intelligence and law in New Zealand', <https://www.cs.otago.ac.nz/research/ai/AI-Law/index.html>
- University of Otago (n.d.-b) 'Centre for Artificial Intelligence and Public Policy', <https://www.otago.ac.nz/caipp/index.html>
- Vincent, B. (2021) 'New session sparks new priorities for Senate AI Caucus', NextGov, 21 January, <https://www.nextgov.com/artificial-intelligence/2021/01/new-session-sparks-new-priorities-senate-ai-caucus/171557/>