Trusted Autonomous Systems Defence Cooperative Research Centre

Submission in response to "An AI Action Plan for all Australians: A call for views"

8 December 2020

Submitter: Professor Jason Scholz, Chief Executive Officer, TAS DCRC

Organisation: Trusted Autonomous Systems Defence Cooperative Research Centre

Address: PO Box 59, Toowong *(Yuggera land)* QLD 4066

Date: 8 December 2020

About TAS-DCRC

Trusted Autonomous Systems Defence Cooperative Research Centre (TAS-DCRC) is Australia's first Defence cooperative research centre, and is uniquely equipped to deliver world-leading autonomous and robotic technologies to enable trusted and effective cooperation between humans and machines. Our aim is to improve the competitiveness, productivity and sustainability of Australian industry.

Supporting Australia's defence capability

TAS-DCRC, together with its participants and the Department of Defence, is focused on developing the capacity of Australia's defence industry to acquire, deploy and sustain the most advanced autonomous and robotic technology through:

- delivering world-leading autonomous and robotic Defence technologies
- building innovative IP through targeted research and technology programs
- assisting Australian industry to develop new, improved and competitive autonomy technologies
- evaluating the utility of autonomous systems through capability demonstrations.

Supporting assurance and accreditation of autonomous systems

In addition to specific industry-led projects, TAS-DCRC is undertaking two 'common-good' activities that have broader, non-defence applications. Through these activities TAS-DCRC will:

- foster ethical and legal research including value-sensitive design
- develop policy pathways for projects and participants
- support development of Queensland air, land and marine ranges for trusted trials, test and evaluation
- establish independent, world-class certification pathways for global industry.

How we work

Trusted Autonomous Systems fosters collaboration between Australia's defence industry and research organisations and aims to increase small and medium enterprise (SME) participation in its collaborative research to improve capabilities of Australia's defence industry. Established under the Next Generation Technologies Fund, with $50 million invested over seven years, and a $15 million co-investment from the Queensland State government, TAS-DCRC aims to deliver trustworthy smart-machine technologies for new defence capabilities based on advanced human-machine teaming.

For additional information on TAS-DCRC, click here.

8 December 2020

Emerging Technologies Section
Digital Economy & Technology Division
Department of Industry, Science Energy and Resources

VIA EMAIL: artificial.intelligence@industry.gov.au

SUBMISSION IN RESPONSE TO AI ACTION PLAN FOR ALL AUSTRALIANS DISCUSSION PAPER

Dear Sir/Madam,

TAS-DCRC welcomes the opportunity to make this submission in response to the AI Action Plan for All Australians ('the discussion paper'). Our submission will mainly focus on non-Defence related AI applications and issues, in line with the scope of the discussion paper.

AI is the key underpinning of autonomous systems. TAS-DCRC is uniquely placed to provide a submission, noting our depth of experience and leadership in this field, and our multi-disciplinary team including scientists, ethicists and philosophers, lawyers, and academics. We have a broad focus on many facets of Artificial Intelligence (AI), and our work covers ethical, assurance, technical, and practical perspectives.

Through our common good activities, specifically *Activity 2: Assurance of Autonomy*, we are actively working to provide a better understanding for Australians about the assurance and accreditation pathways for autonomous systems, in the maritime, air, and land domains, and are working with the regulators and key stakeholders to improve those paths. Our goal is to improve the innovation pipeline, making it faster to design, build, test, assure and certify autonomous systems, while maintaining warranted trust and safety. This work will unlock the many safety, environmental and efficiency benefits autonomous systems can bring, boost Australian jobs, and cement Queensland's status as the 'Drone State' of Australia. We aim to attract more international participants to Australia to use our country's world-class large-scale test ranges, and to generate business for the numerous AI entrepreneurs in Queensland and Australia more broadly.

We welcome the significant investment the Commonwealth and State governments are making in AI throughout Australia, together with the opportunities it brings to all sectors of the Australian economy. We aim to assist the Australian government in further shaping the discussion paper, and broader AI conversation moving forward, to ensure that this record funding is placed where needed most, and to give Australia the best chance of harnessing the benefits of AI, including building sovereign capability and Australian jobs. We are well placed to run pilot programmes to deliver key initiatives, which have been identified as priorities in the discussion paper, as highlighted throughout our submission. Key to success will be a strong focus on ethics and trust, together with engagement and representation of Australia's diverse cultural groups, including first nation communities.

Yours Sincerely

Professor Jason Scholz | CEO

Trusted Autonomous Systems Defence CRC

PO BOX 59, Toowong, Qld 4066, Australia

jason.scholz@tasdcrc.com.au

Table of Contents

About TAS-DCRC **2**
Supporting Australia's defence capability 2
Supporting assurance and accreditation of autonomous systems 2
How we work 2

Executive Summary **5**

Summary of submission **6**

Introduction **8**

One: An AI Action Plan for All Australians **9**

Two: Business **12**

Three: Research **19**

Four: Society **22**

Conclusion **25**

Executive Summary

Artificial Intelligence (AI) is permeating every area of society, but how well the technology is trusted, how well it is applied to solve national problems, and how fairly the benefits of the technology are distributed, remain contentious. The government must act in order to ensure development and use of AI is consistent with Australian ethics and values. To best exploit our unique resources and talents, together with our position in the Southeast Asian region, the government must invest now in building educational pathways and skills relevant to AI, and supporting business, together with developing appropriate AI standards and frameworks. If Australia seeks to compete on the world stage in this domain, it must commit time and resources to facilitate development, testing and evaluation of these technologies. This submission examines three broad themes:

Australian Designed, Australian Built, Australian Owned

The government should ensure the development of Australian sovereign capability regarding AI expertise, data sets, algorithms and systems by:

- Identifying the data sets required by developers of AI-enabled systems or algorithms and then scoping how best to source and facilitate access to Australian organisations through a digital AI marketplace;
- Reducing barriers to entry for small and medium sized AI-centric businesses and prioritising implementation measures;
- Supporting development of guidelines, regulations and regulatory structures, education, with 'good data' to create a world leading Australian AI industry; and
- Implementing incentives, grants and funding to support industry-wide adoption of technical infrastructure for AI development, test, evaluation and delivery through, for example, funding low-risk digital and physical test environments in Australia.

Australia as a World Leader in AI Ethics and Technology

Australia can and must become a world leader in AI ethics and technology. Australian leadership in the development of standards and ethics related to AI will ensure Australian ethical principles (as detailed in the Department of Industry, Science, Energy and Resources (DISER) AI Ethics Principles)[1] underpin global AI development. World-class educational pathways will generate and sustain an ethical AI-literate workforce and ensure that our expertise is desired worldwide and considered world-leading.

TAS-DCRC is well positioned to be a trusted Government partner to progress AI technology

The TAS-DCRC has a track record of converting government funding into tractable, practical outcomes through a fusion of industry leadership, academic rigour and government oversight to facilitate rapid and agile development of groundbreaking AI technologies. This capacity should be utilised to progress Australian AI development.

[1] https://www.industry.gov.au/data-and-publications/building-australias-artificial-intelligence-capability/ai-ethics-framework/ai-ethics-principles.

Summary of submission

Topic		Question	Recommendation
An AI Action Plan for all Australians	1.1	What is the role for government to support the uptake and use of AI technologies in Australia?	See section below
	1.2	What can be done to reduce barriers to AI adoption in Australia?	See section below
	1.3	Do we have the right vision for AI in Australia?	See section below
Business	2.1	How can we identify and unlock the value of uniquely Australian datasets?	(1) The government should fund a data scoping exercise, and support a pilot programme to deliver a data package for a high priority area
	2.2	How can we lower the barriers to entry for businesses and government developing, piloting or assessing the value of AI while ensuring appropriate consumer safeguards?	(2) The government should consider funding: a. a scoping exercise to develop a comprehensive understanding of the barriers to entry and how best to address them b. a project to develop a national risk management framework and corresponding standard for AI systems, to ensure a well-balanced approach to consumer protection c. a project to prioritise and deliver solutions to address the key barriers, including funding a pilot programme to improve digital infrastructure
	2.3	How can government help ensure that AI research, including international collaboration, is undertaken safely, ethically and responsibly?	(3) The government should invest in guidelines and regulations, education, good data, and appropriate regulatory structures (4) The government should consider the type of organisation that can best convert funding into trusted output in this field
Research	3.1	What are the problems Australia is facing where the development and application of AI could provide long-term solutions and how could these be prioritised?	(5) The government should create a register of national research priority areas

	3.2	How can Australia best coordinate its national research effort around areas of national priority?	(6) As per Recommendation (5), assign these priorities uniquely to research providers (could be derived from a competitive process) to ensure academic centres have and encourage depth with reduced duplication; incentivise IP transfer of research by individuals for industry exploitation (especially start-ups and SMEs)
	3.3	How can we better support industry-researcher engagement?	(7) The government should increase financial incentives, reduce bureaucratic requirements (e.g. remove royalties for Commonwealth IP), improve intellectual property rights and incentives for individuals to encourage 'ownership' and entrepreneurship, and improve accreditation frameworks
People		Not addressed	
Society	4.1	Is there more the government can do to support responsible and human centred development and use of AI in Australia?	(8) Australia should ensure that Australian ethics are imbued into the development, deployment and adaptation of AI technologies that will have an effect on the Australian people
	4.2	What approach should Australia take internationally to steward its values and commitment to the responsible and ethical use of the AI? How can Australia support its partners and neighbours in their efforts to make the most of AI?	(9) Australia should commit to the development of a comprehensive framework for the ethical design, implementation and assessment of AI
	4.3	What security issues associated with AI systems should be considered?	(10) The government should create infrastructure to support Australian sovereign capability to develop and share AI data sets and algorithms to ameliorate the risks associated with the use of foreign algorithms and AI systems
	4.4	How can Australia support its partners and neighbours in their efforts to make the most of AI?	(11) Australia should take a leading role in the global development of AI standards, including creating an institution-level database/marketplace to share information about project needs (industry) and research capabilities (academia and industry) to facilitate both research and commercial collaborations

Introduction

AI is creating new and exciting opportunities for Australian prosperity. The TAS-DCRC, through its research projects and common good activities, is an innovator, facilitator, and connector, as well as an ethical steward of this technology in both Defence and civilian contexts, and is perfectly placed to help private and public organisations in Australia get the most from AI.

AI and autonomous systems have an uncapped potential to deliver far reaching safety, environmental, and efficiency benefits, including removing humans from "dull, dirty, dangerous and distant" environments and jobs. The challenge for the Australian government is identifying where to focus funding for AI projects and how best to facilitate the use of that funding to ensure the greatest possible benefit is realised, and shared across the breadth of the Australian community.

The document, *A Robotics Roadmap for Australia 2018*,[2] published by the Australian Centre for Robotic Vision, sought to highlight Australian industry use and development of robotics, and identify how Australia can support a vibrant robotics industry that supports automation across all sectors of the economy.[3] Version 2 of the Roadmap is currently under development, coordinated by Dr Sue Keay of the Queensland AI Hub. TAS-DCRC took a leading role in coordinating both the 'Defence Sector' and 'Trust and Safety' chapters. The 'Trust and Safety' chapter consolidates knowledge, ideas and insights from 18 stakeholder groups across the spectrum of industry, academia and government. Topics covered include 'safety and assurance'; 'what exists now'; 'emerging issues in assuring autonomous robots'; 'regulation in air, maritime and land'; 'robotics standards and regulation'; 'new ways of thinking about regulation' and 'designing democratically legitimate AI systems'. We would be happy to provide a copy of the chapter once it has been finalised.

The discussion paper seeks to identify existing programs or market solutions already addressing the highlighted action areas. TAS-DCRC is an exemplar of how a small agile team, with modest funding and a deep pool of expertise, is able to steward and facilitate game-changing technological, legal and ethical developments, putting Australian businesses at the forefront, to ensure the benefit will reach further. By partnering with industry leaders, Defence sponsors and academic institutions in a formal project management framework, with unique Intellectual Property (IP) rights preserving instruments for sovereign industry, we ensure academic deliverables are provided promptly and immediately applied in a practical way, to a degree not previously seen. The Defence CRC Guidelines and model implemented by TAS-DCRC could be applied in other non-Defence sectors, and is an efficient, successful method of turning modest investment into strong returns, with an ethical overlay capable of providing regulator, government and community trust.

[2] https://www.roboticvision.org/robotics-roadmap/.

[3] https://research.csiro.au/robotics/australian-robotics-roadmap/.

One: An AI Action Plan for All Australians

Q1.1: What is the role for government to support the uptake and use of AI technologies in Australia?

There are many reasons for why the government takes responsibility for addressing a challenge – it may be a question of common good, it may impact on justice and equitable access, or it may relate to national security. Where the market is unable to address the challenge itself, the government needs to take a leadership role in shaping policy, providing funding, and ensuring the protection of Australian interests and values.

The government must be proactively involved in the development of AI technology, and more specifically in the shaping of national policy, standards and guidelines. To achieve this, the government must invest funding and incentivise where appropriate. We see the priorities for government action as follows:

- The Australian government must invest in:
    - targeted research and supporting research institutions
    - digital infrastructure
    - AI assurance
- The government must develop a suitable regulatory framework and support regulators to enact that framework
- The government must support a high-technology ecosystem, centered on research, development and collaboration

If the Australian government wishes for Australia to remain competitive on the world stage and create genuine sovereign AI capability, it must commit significant resources and time to the creation and generation of a specific high-technology ecosystem, centred around research, development, and collaboration. Failure to wholly commit to this path may leave Australia with critical dependencies on foreign nations.

Q1.2: What can be done to reduce barriers to AI adoption in Australia?

Key barriers to AI adoption include:

- lack of education regarding the potential of AI for solving contemporary problems;
- lack of awareness of and access to Australian AI capabilities;
- difficulty and cost of developing, building, testing, assuring and accrediting AI products to get them to market;
- lack of a trusted safety rating and accreditation framework (for example an ISO standard);
- expense of purchasing products, caused by the high resource cost incurred by developers; and
- uncertainty regarding liability for developers and users of AI products.

This submission provides strategies for reducing these barriers.

Ethics must be addressed

There should be expectations that AI developers and providers have ethical and legal frameworks in place so that businesses, organisations and individuals can understand how they have ethically de-risked their operations and have more confidence that they can be trusted. For example, recently, before deciding whether to work with a global AI provider, we posed a series of questions to them to help us identify whether they are operating within ethical, legal and regulatory bounds. These questions included:

1. What responsibility does Company X take for being ethical and law-abiding in the development of AI technologies[4] How is responsibility codified into documentation, culture, technologies and process?
2. What frameworks, standards, policies, laws, best-practice guidelines etc. does Company X use to provide normative guidance, rules and methods or processes pertaining to ethics, law and regulation of AI to stakeholders (including staff, customers/clients, other businesses, governments and organisations)?[5]
3. What (if any) ethical AI services does Company X offer? (See Google[6])
4. What limits (if any) does Company X place on the type of data, models or algorithms used or technologies or implementations they help develop? What organisational structures are in place to ensure awareness and compliance with these limits?
5. What limits (if any) does Company X place on the type of customers, clients or projects they accept for AI work?
6. What oversight mechanisms does Company X have within the organisation to assist identification and management of ethical issues? (See the Facebook Oversight Board for Content Decision[7] and Microsoft's Office of Responsible AI (ORA))[8]
7. What research and development is Company X doing in ethical AI[9]

The TASDCRC (with Defence Science and Technology Group and Plan Jericho Air Force) is co-author of the DST technical report 'A Method for Ethical AI in Defence' that recommends an evidence-based and pragmatic taxonomy for reducing ethical risks in Defence AI projects—copy available upon request. This method is being evaluated via TASDCRC projects and case studies within Defence.

Q1.3: Do we have the right vision for AI in Australia?

While the vision set out in the discussion paper is a good start, we consider a stronger focus is needed on:

1. embedding ethics and facilitating trust;
2. consulting and collaborating with Australia's diverse cultural groups, including first nations communities;[10]

[4] E.g. ensuring security, differential privacy, reliability, auditability.
[5] See Microsoft Responsible Innovation: https://docs.microsoft.com/en-us/azure/architecture/guide/responsible-innovation/.
[6] https://www.wired.com/story/google-help-others-tricky-ethics-ai/.
[7] https://about.fb.com/news/2019/06/global-feedback-on-oversight-board/.
[8] https://www.microsoft.com/en-us/ai/responsible-ai.
[9] See Microsoft's AI, Ethics, and Effects in Engineering and Research (AETHER) Committee and working groups
[10] The 'Deadly Innovation' Program run by Duncan Kerslake at the Department of State Development, Tourism and Innovation Queensland, may be a good stakeholder to engage with on this point.

3. considering the national security and defence implications of AI, including the development and use of dual-use technologies; and
4. leveraging developments being made in industry, defence and academic contexts, to ensure maximum return for investment, prevent duplication, and ensure a focus on Australian sovereign capability.

Two: Business

Q2.1: How can we identify and unlock the value of uniquely Australian datasets?

1. The Government should fund a data scoping exercise, and support a pilot programme to deliver a data package for a high priority area

Access to good quality, ethically sourced, secure data sets is critical for the development of trusted AI-enabled systems in Australia. The Australian government has already made some progress in this space, as noted in the discussion paper, including the creation of the Australian Ocean Data Network. Coordinated early action now will remove a barrier to entry for SMEs developing AI technology and enable government control over access to 'data packages'. This will enable the application of best-practice ethical and security guidelines to data package access, and reduce the risk of compromised data sets being used by necessity because no alternative is available.

We recommend the government fund a scoping exercise to identify:

- what data sets are required by developers of AI-enabled systems now and expected to be required in the future;
- identify where that need is currently not being met;
- identify how best to source them; and
- how to facilitate access.

An example would be a set of robust data that provides a good, bias-free representation of relevant environmental contexts in which systems will be deployed. A further example would be a data set of oceanographic data for a location-specific volume of ocean, such as temperature, salinity, and seafloor depth. This would be partnered with high fidelity physical models and general data sets, and used to help train an autonomous system intended to operate an underwater autonomous vessel. Consideration also needs to be given to the security and ethical implications of releasing data sets, and the policies and procedures that need to be implemented as a result.

Where there is a clear need for location-specific data, such as digital twins for autonomous systems test ranges, the government should consider which organisations are already well-placed to collect that data, and fund a pilot programme. This should be undertaken concurrently with the scoping exercise, together with development of policies, procedures, and ethical guidelines.

TAS-DCRC is well-placed to conduct a pilot programme to gather data sets for QLD test ranges

The Queensland government, recognising the potential of the AI industry to bring jobs, business and investment to Queensland, has funded TAS-DCRC's A2 Activity: Assurance of autonomy, specifically the National Accreditation Support Facility Pathfinder (NASF-P) project. A key goal of NASF-P is identifying the barriers to uptake of autonomous technology, and working to improve them.

The NASF-P team have already identified lack of access to robust data, including location-specific data, as a major barrier to entry for SMEs developing AI technology such as autonomous vessels, and for potential customers due to the increased cost it represents.

The NASF-P project already has a draft project proposal prepared for improving digital infrastructure in Queensland, which focusses on facilitating faster development of autonomous systems and products, bringing developers to Queensland, and bringing clients to Queensland testing areas, by creating and then making available to approved users 'data packages' for these test ranges. This project would utilise local workers to assist in the collection of location-specific data, in order to facilitate job creation and ensure investment remains in the local economy.

The government could consider supporting the NASF-P digital infrastructure project as a pilot, and then, if successful, replicating the model in other priority locations around the county. A copy of the project proposal is available on request.

Q2.2: How can we lower the barriers to entry for businesses and government developing, piloting or assessing the value of AI while ensuring appropriate consumer safeguards?

2. The Government should consider funding:

a. a scoping exercise to develop a comprehensive understanding of the barriers to entry and how best to address them;

b. a project to develop a national risk management framework and corresponding standard for AI systems, to ensure a well-balanced approach to consumer protection; and

c. a project to prioritise and deliver solutions to address the key barriers, including funding a pilot programme to improve digital infrastructure.

Understanding the barriers

We know that there are significant barriers to entry for businesses and governments developing, piloting or assessing the value of AI technology, and that in considering how to address those barriers, it is necessary to understand the risks this technology raises and implement appropriate consumer safeguards. The government should fund a scoping exercise to support a better understanding of the key issues and priorities for this topic.

The main barriers we see for developing, piloting or assessing the value of AI, from an industry perspective, some of which were mentioned in the discussion paper, are:

1. an understanding of the regulatory structure of the environment the autonomous system will need to comply with and operate within;
2. lack of clear, documented and accepted assurance standards and methodologies;
3. lack of consistent risk management methodologies, frameworks and expectations;
4. lack of clear regulation, and regulator expectations;
5. lack of access to good quality data to support development of autonomous systems;
6. lack of affordable, accessible digital infrastructure for developers to use to develop and test algorithms, and to digitally test the whole-system operation;

7. lack of suitably experienced third-party surveyors, inspectors etc;
8. lack of access to affordable testing areas;
9. slow, expensive approval processes;
10. uncertainty regarding future regulatory requirements; and
11. accessing funding while allowing industry to maintain ownership of IP rather than ceding it to a government or academic institution, or a separate private company.

While the discussion paper identified some of the above listed barriers experienced by industry, it did not canvas many of the barriers experienced in a government context. The barriers experienced by government are more likely linked to lack of overall policy direction, uncertainty about risk management expectations, difficulty 'trusting' technology that is not well understood or in common use, and lack of clear frameworks, standards, regulations and legislation to rely upon in decision-making processes. The government also needs to understand community expectations in terms of the level of consumer safeguards that will be applied, so that they can ensure they are neither over, nor under-regulating, or being insufficiently or excessively cautious in decision making.

Supporting risk-based decision making

In our experience, there is a lack of technology-specific understanding; i.e. how does an autonomous vessel work, and a lack of risk-based understanding; i.e. what are the risks, what are the mitigation measures, and where is the point the risk is appropriately mitigated, including in relation to appropriately protecting consumers and the wider Australian public. This risks either that no action is taken or that no decisions are made as agreement cannot be reached on required risk mitigation, or overregulating the number of mitigations to the point that it is not feasible to conduct the development or operation of an autonomous system because it is excessively expensive or slow to comply with the mitigations required.

There needs to be a specific risk management framework developed for AI systems, supported by best practice and/or standards. This framework would support assessment and mitigation of common risks to achieve an acceptable level of safety. This body of work would reduce a barrier for both industry and government developers and users of AI technology, as well as support the regulators; for example the Australian Maritime Safety Authority and the Civil Aviation Safety Authority.

Prioritising solutions and supporting a pilot programme to improve digital infrastructure

Once there is a clear understanding of the barriers to entry for business and government in developing, piloting or assessing the value of AI, and an appropriate risk management framework in place, an effort should be undertaken to prioritise the order of addressing those barriers. The prioritisation should consider the barriers that are having the biggest impact, and the barriers that can be addressed the quickest.

One of the key barriers identified by numerous TAS-DCRC stakeholders, primarily SMEs, is the lack of affordable, accessible digital infrastructure that developers can use to develop and test algorithms, for example machine learning algorithms, using real-world, location-specific

data simulations, which are secure; subject to strict ethical use guidelines, and whose output is trusted by regulators.

The government should invest in a pilot programme that provides secure soverign digital infrastructure to support Australian AI development, test and evalution and marketing

Large corporations, such as Google, IBM, or Fujitsu have a (currently) nearly insurmountable commercial advantage over start-up and small to medium enterprises (SMEs), because their existing software and cloud infrastructure is advanced and highly automated in its support for software development and, specifically, machine learning. The cost and complexity in establishing such infrastructure and emerging new processes is high, creating a disproportionate barrier to entry for SMEs as well as dependence on non-Australian digital infrastructure.

Our experience working with developers of major game-changing technology highlights:

- the myriad independent systems in use by SMEs, which can be hard to integrate into industry (or Defence) applications, reducing software reuse and increasing down-stream integration cost and complexity;
- the low maturity level of machine learning technical infrastructure and processes (typically "level 0" or lower);[11]
- the linear waterfall approach to software development, which hinders speed to deployment and transitioning to Agile methodology is challenging;[12]
- the lack of responsible innovation practices (e.g. absence of ethical, legal or regulatory standards) in machine learning development, test and evaluation,[13] (this includes Primes);[14]
- the difficulty of harmonising process-driven agile software methodology with data-driven machine learning developments;[15]
- the lack of automated testing, which leads to slow, inefficient and more error prone products;
- the independent development of simulations for machine learning training, causing significant duplication of effort across firms and inconsistent quality; and
- the labour-intensive nature of research translation from universities into technology development.[16]

[11] See generally, Google ML pipeline models: https://cloud.google.com/solutions/machine-learning/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning.
[12] Fernando Almeida, 'Challenges in Migration from Waterfall to Agile Environments' (2017) 5(3) *World Journal of Computer Application and Technology* 39–49.
[13] See CSIRO: https://www.csiro.au/en/Research/LWF/Areas/FSPs/Responsible-Innovation.
[14] The term "Prime" refers to the defence primary contractors – companies that deliver many of the major Defence projects in Australia: https://www.defenceindustry.gov.au/about-defence-industry.
[15] See for example: Google ML pipeline models: https://cloud.google.com/solutions/machine-learning/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning.
[16] Universities also lack significant AI infrastructure, the US and Australia, see generally, https://hai.stanford.edu/blog/national-research-cloud-ensuring-continuation-american-innovation.

Government investment in the right third party to develop soverign digital infrastructure would address the gaps we have identified in the development of trusted autonomous software for industry in a way that leverages the technologies, knowledge and practices of large international firms. This will remove unnecessary barriers to entry for qualified Australian SMEs and speed the path to market. This pilot programme will serve as a leading industry testbed, growing capability and assisting industry and government to evaluate candidate technologies and methodologies.

Sovereign digital infrastructure would provide requisite functionality, a machine learning pipeline, and simulation environment with common tools. This would be hosted on a secure Australian cloud environment, and be accessible to approved Australian SMEs as a common-good service. The government should also invest in a best-practice machine learning pipeline for responsible, secure development and operation.

The secure, shared system will allow and encourage mutual awareness of data and software components across industry and Goverment, reducing duplication and encouraging reuse. It will also facilitate capture of IP for the benefit of industry.

The key safety regulators, AMSA and CASA, need support to enable improved regulation of emerging technology, in a way that facilitates innovation, harnesses the opportunities for increased safety, environmental outcomes and efficiency that AI offers, and properly balances risk with regulatory impact.

Australia's regulatory structures and regulators are not well placed to adapt to emerging technology, as evidenced in the barriers seen in maritime regulation, and by the difficulty of creating a suitable approach to small drones and an appropriate assurance framework for larger drones in the air domain. Expecting traditional regulators to cope with an influx of new technology, with little government policy direction, limited additional resources to upskill a workforce or IT system, and no prerogative to reform decision-making and management structures to achieve a more agile agency, is bound to be problematic. The government needs to deliver a consistent policy on the prioritisation of improving the regulation of emerging technology and the appropriate risk appetite, and ensure regulators are well placed to respond accordingly.

Case Study – innovation measures to address barriers

The Queensland government, which is investing heavily in positioning Queensland as the Smart Drone State, recognised early on the above-mentioned issues, and has funded the TASD CRC *A2 Activity: Assurance of Autonomy*, specifically the National Accreditation Support Facility Pathfinder (NASF-P) project, to investigate setting up a third-party entity that could better link operators with the regulator. This model has the potential to alleviate many of the barriers currently in place in the maritime, air and land domains, which would facilitate innovation while retaining a focus on trust and safety.

The government should support the development of similar proposals or organisations in other industry sectors, noting their potential to dramatically improve uptake of emerging technology in Australia.

Q2.3: How can government help ensure that AI research, including international collaboration, is undertaken safely, ethically and responsibly?

3. The Government should invest in guidelines and regulations, education, good data, and appropriate regulatory structures

There is currently a lack of suitable guidelines and regulations governing AI research, including the development of AI-enabled systems and the data they are trained on. This needs to be addressed in a manner consistent with best practices internationally, reflecting agreed ethical concepts, as a matter of priority.

More investment in transdisciplinary tertiary education in the area of trust and safety of AI-enabled systems is required to support the development of a new professional field. This will ensure an open sharing of knowledge and ideas, reduce duplication, and promote shared understanding of the different contexts of design, testing, certification and use. The rapidly increasing market for this type of tertiary education provides significant opportunity for Australia to position itself at the leading edge of offerings and iterative course development.

Early investment in high quality data sets is also essential. Without these, we risk AI-enabled systems either not being developed to their full potential, or being delivered using compromised data that presents security, ethical and legal issues. These issues were canvassed in the discussion paper, and we support the discussion paper in highlighting the risk around using data that embeds historical biases.

The government should consider establishing a new organisation to develop best practices and guidelines for responsible AI development in accordance with the recommendations made in the AI Standards Roadmap released by Standards Australia.[19]

Australia has the capacity to become a world leader in AI, and the decisive action by the government will help Australia to achieve that potential. Identifying the specific areas that best align to Australia's needs, resources, and values, and investing and subsidising where needed, will help to ensure these areas are prioritised by industry. Investing in education and organisations dedicated to those specific areas will also be critical.

4. The Government should consider the type of organisation that can best convert funding into trusted output in this field

The government should consider the benefits of the model used successfully by the TAS-DCRC, noting the structure, persons employed, strategic and legislative governance, and

[19] Final Report: An Artificial Intelligence Standards Roadmap: Making Australia's Voice Heard, Standards Australia: https://www.standards.org.au/getmedia/ede81912-55a2-4d8e-849f-9844993c3b9d/O_1515-An-Artificial-Intelligence-Standards-Roadmap-soft_1.pdf.aspx.

Defence linkages, have contributed to an environment that fosters collaboration and ethical development of AI related technology among diverse stakeholders.

Case study: a Collaborative Research Centre (CRC), stewarding CRC-Projects

The TAS-DCRC model highlights the immediate value and application of academic projects to live real-world projects. By partnering industry leaders, Defence sponsors and academic institutions within a formal project management framework, it ensures academic deliverables are provided promptly and immediately applied in a practical way. A key feature of the Defence CRC and CRC-P models is that IP created by individual projects remains vested in the industry lead(s) for that project. The TAS-DCRC ensures all legal terms are satisfied up-front to avoid the problem of having to disentangle valuable IP at the end of a project. Our model would translate well in other industry sectors, and is an efficient and successful method of turning modest investment into strong returns, with an ethical overlay capable of providing regulator, government and community trust.

Three: Research

Q3.1: What are the problems Australia is facing where the development and application of AI could provide long-term solutions and how could these be prioritised?

As identified in the discussion paper, there are significant opportunities to address national problems using AI. For example, AI can improve safety, security, efficiency, and equity in transport, agriculture, mining, business, defence and health. A prime example is the use of assisted driving technologies to improve road safety and reduce the instances of motor vehicle accidents. The impact on civil liability regimes, insurance, risk management, and health frameworks that result will see a significant reduction in the negative consequences flowing from motor vehicle related damage to persons and property.

Some of the key problems we have identified in an Australia context, where the development and application of AI could provide long-term solutions, include:

1. high levels of road fatalities and injuries;
2. humans are employed in dull, dangerous and dirty jobs;
3. lack of access to services for remote and regional communities;
4. congestion and high infrastructure costs of transport;
5. high environmental impact of shipping and transport more broadly;
6. the capacity of Defence personnel to protect Australia, and Australian servicemen and servicewomen, against increasingly sophisticated adversaries; and
7. the safety of Defence personnel and civilians in conflict zones.

In addition to this problem-driven approach to AI, what has proven to be an international attractor of research talent and reputation is to create and solve a World-class challenge. Google's Deep Mind Go challenge is a case in point; and more recently the natural language capabilities of the Generative Pre-trained Transformer GPT-3. Perhaps Australia might take a global leadership initiative and develop an Artificial General Intelligence (AGI) challenge supported by government-funded world-class compute infrastructure, open to Australian universities and industries.

5. The Government should create a register of national research priority areas

Australian resources can be allocated more efficiently if the government establishes a register of national priority areas to direct effort toward areas of greatest impact. A comprehensive taxonomy will give a clearer indication of the direction Australia needs to take to create robust capabilities in the near-to-medium term. Australia also needs to make more use of the research centres at its disposal through increased engagement and funding. The highly agile and responsive nature of CRC structures enable them to coordinate fast-tracked deliverable solutions to problems in AI research and development that traditional research pathways are unable to deliver. Making increased use of these facilities will be key to Australia's future strength as a leader of AI technology on the world geopolitical stage.

AI will dramatically impact Australia's defence capability

The development of AI will dramatically impact Australia's defence capability. The small population, coupled with the large land mass and extensive coastline, place Australia at a disadvantage compared to other nations. Australia could leverage the increased force projection that AI related technologies will provide in future. Accordingly, the incorporation of AI into Australia's defence strategy will likely prove crucial moving forward. The decline in Australian heavy industry and manufacturing over the past several decades has left Australia dependent on neighbouring countries for production and supply in these sectors. This situation leaves Australia potentially vulnerable in a time of increasing tensions in the geopolitical arena, along with the possibility of renewed great power conflict. Building sovereign capability in AI-related technologies is therefore a critical outcome.

Q3.2: How can Australia best coordinate its national research effort around areas of national priority?

6. *As per Recommendation (5), the Government should assign these priorities uniquely to research providers (potentially through a competitive process), to ensure academic centres have, and encourage, depth of research with reduced duplication; incentivise IP transfer of research by individuals for industry exploitation (especially start-ups and SMEs)*

The Australian government should direct funding to the high priority areas as determined by the register outlined in Recommendation 5. Additionally, another way to support the development of Australian AI capability is to embed researchers into industry (and vice-versa) so as to promote the continued growth of smart industry and innovation. Empowering industry through partnering with researchers will generate strong innovation pathways for Australia's future participation in the global marketplace.

The TASD-CRC has experience working with overseas bodies seeking to invest significant funding into research in Australia. These bodies report considerable duplication of AI research capabilities across academic institutions and significant difficulties in finding world-class small cells. The government may be able to help by assigning priorities/areas of AI to specific research providers. This might be achieved through a competitive process, with reviews to ensure academic centres have, and are, developing and encouraging depth and reduced duplication.

Q3.3: How can we better support industry-researcher engagement?

7. *The Government should increase financial incentives, reduce bureaucratic requirements (e.g. remove royalties for Commonwealth IP), improve IP rights and incentives for individuals to encourage 'ownership' and entrepreneurship, and improve accreditation frameworks*

TAS-DCRC agrees with the discussion paper observation that collaboration between public research organisations and businesses is below the OECD average. This should be addressed by leveraging industry-researcher engagement, supported by concessional tax incentives to facilitate entry into the market. A significant factor in achieving this would be to address the barriers caused by overly bureaucratic requirements. Continued government support of CRCs

will improve Australian industry-researcher engagement across all levels of AI related industry. This will assist cutting edge academic research to permeate industrial entrepreneurship and initiatives, for a national commercial benefit. Developing a better understanding of the barriers to assurance and accreditation of AI technologies would also assist in this process. Australia needs bipartisan support to ensure continued technological development, or there is a risk that progress will be impeded at this crucial point in Australia's history.

Government might better incentivise IP development in Australian industry by transferring Commonwealth background IP direct to industry via non-exclusive licenses without creating future financial obligations on industry. The Commonwealth should cease its policy of seeking to pre-load financial impediments on industry for its background IP, long before it has been developed by industry into a commercialisable form. Examples of this unwelcome Commonwealth practice include royalties, and other financial provisions – all of which create disincentives for external investors in those companies. Rather, we recommend the Commonwealth give away its background IP in the national interest, to contribute whole-heartedly to the success of new sovereign industry developments, and then use the taxation system to reap the revenue from industry following successful commercialisation.

Academic institutions should not be encouraged to hold and exploit IP. The role of IP exploitation should rest with industry – not academia. This should be incentivised to pass the IP on, in the national interest. Further they should be encouraged to empower researchers in their individual ownership of that IP if they are seeking to exploit it through a new start-up firm or SME.

Australia has many, many small companies, and few make it to 'medium' size in this industry. Incentives to build AI industries at medium scale should be specifically encouraged. The development of supporting national infrastructure described elsewhere in this submission explains this further.

Four: Society

Q4.1 Is there more the government can do to support responsible and human centred development and use of AI in Australia?

8. Australia should ensure that Australian ethics are imbued into the development of AI technologies that will have an effect on the Australian people

In order for AI to become ubiquitous in society, the public must be able to trust in the design, implementation and use of AI technologies. The DISER AI Ethics Framework provides a foundation and first step for the establishment of trust.[20] Australia should promote and subsidise case studies using the DISER framework and adapt and develop the framework into pragmatic and affordable tools. Uptake and compliance with this framework can be assisted through research, and via education and training.[21]

Additionally, the Australian government should create a funding pathway through grants, scholarships and research opportunities to enable Australian ethicists and Australian AI researchers to further Australia's theoretical, conceptual and applied expertise in the ethical use of AI. These research opportunities should cover both normative and descriptive areas. That is, exploring how decisions ought to be made and how humans actually make ethical and moral decisions using real technologies *in situ* and what socio-technical factors affect ethical decision making.

Funding should extend from purely research opportunities into secondary and tertiary study pathways (including apprenticeships, mentorships and placements with Industry) to help create the next generation of AI-literate Australians. This should be undertaken in conjunction with training for public service officers, who will inevitably have to create policies, enact legislation, determine non-compliance and enforce penalties related to the failure to meet requisite ethical standards.

As described in the discussion paper, businesses need clearer guidance on how ethics frameworks should be applied. The government should fund programs to build better decision support tools that augment and enhance human ethical reasoning, reflection and considerations to align with society's expectations of human accountability. The government can then use this information to develop (or adopt) and recommend pragmatic tools, such as the Open Data Institute's 'Data Ethics Canvas'[22] or the Markkula Centre for Applied Ethics

[20] https://www.industry.gov.au/data-and-publications/building-australias-artificial-intelligence-capability/ai-ethics-framework/ai-ethics-principles.

[21] For example, 'AI and the Law' offered by the University of Melbourne, Centre for Artificial Intelligence and Digital Ethics (CAID) in 2021. CAIDE is unique in bringing together a multi-disciplinary team to deliver an undergraduate course, conduct research, and facilitate opportunities for collaboration with different academic institutions https://handbook.unimelb.edu.au/2021/components/btrack-133.

[22] https://theodi.org/article/data-ethics-canvas/.

'Framework for Ethical Decision Making'.[23] This will assist AI developers and providers to ethically de-risk their projects through step-by-step workshops, documentation, project management planning, and ethical risk management, test and evaluation.

Q4.2: What approach should Australia take internationally to steward its values and commitment to the responsible and ethical use of the AI?

9. Australia should commit to the development of a comprehensive framework for the ethical design and implementation of AI

An Australian-centric ethical framework can be further supported by advancing Australian participation in policy and regulation development on the world stage. The recently created Global Partnership on Artificial Intelligence (GPAI) is the world's first multilateral and multi-stakeholder forum with a focus on AI. Comprising multiple participants,[24] 'GPAI working groups will collaborate across four themes: Responsible AI, Data governance, Future of work, and Innovation and commercialisation'.[25] Australian academics have been nominated to work on the taskforce to build an ethical framework for AI, including distinguished Professor Genevieve Bell, founder of 3Ai at ANU.[26]

Through the development of a robust Australian AI ethical framework, coupled with participation in trans-national, and supra-national discussion, Australia should position itself to become a world leader in the ongoing dialogue around the ethical development and deployment of AI related technologies into the future.

Q4.3: What security issues associated with AI systems should be considered?

10. The Government should create infrastructure to support Australian sovereign capability to develop and share AI datasets and algorithms to ameliorate the risks associated with the use of foreign algorithms and AI systems

With any new technology, the extent of ontological risks (unknown-unknowns) and epistemological risks (known-unknowns) are significantly larger compared to current or former technologies. These unknowns could leave Australia exposed to attack if not addressed. TAS-DCRC agrees with the discussion paper that through the use of foreign-sourced datasets or algorithms, the risk of malign actors infiltrating or influencing Australian AI organisations and systems is high.

[23] https://www.scu.edu/ethics/ethics-resources/ethical-decision-making/a-framework-for-ethical-decision-making/.
[24] Australia is a founding member of the GPAI, along with Canada, the European Union, France, Germany, Italy, Japan, New Zealand, the Republic of Korea, Singapore, Slovenia, the United Kingdom, and the United States of America. https://cecs.anu.edu.au/news/college-leaders-selected-international-ai-taskforce.
[25] https://cecs.anu.edu.au/news/college-leaders-selected-international-ai-taskforce.
[26] Dean of the ANU College of Engineering and Computer Science, Professor Elanor Huntington, and Director of the 3A Institute, Distinguished Professor Genevieve Bell, have been selected among Australia's first nominations to join a forum for the Global Partnership on Artificial Intelligence (GPAI).https://cecs.anu.edu.au/news/college-leaders-selected-international-ai-taskforce.

The Australian government should focus on the creation of infrastructure to support the development and sharing of sovereign capability in AI datasets, algorithms and associated disciplines. This could take the form of funding, subsidies, grants etc, detailed earlier in this submission. Encouraging partnerships between industry and academia to intensify Australian AI capability development, through organisations like TAS-DCRC, is essential.

Q4.4: How can Australia support its partners and neighbours in their efforts to make the most of AI?

11. Australia should take a leading role in the global development of AI standards, including creating an institution-level database/marketplace to share information about project needs (industry) and research capabilities (academia and industry) to facilitate both research and commercial collaborations

Australia is a relatively small nation in terms of population, however, our expertise allows us to punch well above our weight. The government should ensure that there is extensive Australian participation at international standards bodies to influence AI development into the future. Additionally, as AI technology expands, it is critical to ensure the Australian AI industry has the opportunity to contribute to government participation at international forums. A focal point between industry, academia, and government would be beneficial in facilitating these discussions. Organisations like TAS-DCRC operate as focal points between industry, academia and government, and are ideally placed to facilitate interactions of this nature.

The ability for Australia to influence global discussion is contingent on the continuance and growth of Australian expertise in AI, including in relation to AI ethics and the responsible use of AI. The actions required by government to ensure that Australia remains at the forefront in this area are discussed throughout this TAS-DCRC submission.

To support our international partners and global neighbours, the Australian government should create an institution level database to share unclassified information about research projects, facilitate collaboration, and improve the global use of AI in line with our Australian vision.

Conclusion

If Australia wishes to maintain or improve its position on the global stage, it must embrace the safe and appropriate integration of AI-related technologies into Australia society. AI, if appropriately exploited, can add to Australia's prosperity by creating more agile industry and business, enhancing our sovereign capabilities, and enriching education, health and the environment. AI and its related technologies will become ubiquitous and transformational if appropriately supported and regulated by progressive government policy.

The government must promote a diverse and robust research and education culture, and provide economic support and encouragement for new multi-disciplinary courses that will equip Australia's workforce of the future.

Our TAS-DCRC vision for Australia includes building robust AI frameworks, technologies and expertise, to strengthen Australian sovereign capability in the AI domain. The 11 recommendations contained within this submission provide a roadmap towards achieving this future.