

towards our intelligent future



Te ara mō tātou **atamai** o āpōpō

AN AI ROADMAP FOR NEW ZEALAND

TE HUARAHI ATAMAI IAHIKO Ō AOTEAROA

FIGURE 1: Towards Our Intelligent Future

towards our intelligent future



#intelligentfutureNZ



PART A:
An AI Roadmap
for New Zealand

- 1. Understanding and Explaining AI
- 2. Recent AI Developments
- 3. An AI Roadmap For New Zealand

PART B:
Enabling AI
Adoption in
New Zealand

- 4. Research
- 5. Trusted, Accessible Data
- 6. Talent, Skills and Capability
- 7. AI Ethics, Regulation and Society
- 8. AI Readiness and Investment

About the AI Forum of New Zealand

THE ARTIFICIAL INTELLIGENCE FORUM OF NEW ZEALAND IS A NON-GOVERNMENT ASSOCIATION WITH A MISSION TO HARNESS THE POTENTIAL OF ARTIFICIAL INTELLIGENCE (AI) TO HELP BRING ABOUT A PROSPEROUS AND INCLUSIVE FUTURE NEW ZEALAND.

The rapid development of AI technologies presents major opportunities and challenges for our country: from creating world leading AI businesses, nurturing a pool of talented AI engineers and applying AI technologies to our agriculture, government, manufacturing and service industries to holding a meaningful national debate on the broader implications for society, New Zealand needs to actively engage with AI now in order to secure our future prosperity.

The Forum brings together citizens, business, academia and the government to connect, promote and advance the AI ecosystem to help ensure a prosperous New Zealand.

Acknowledgments

The AI Forum of New Zealand would like to acknowledge the following contributors:

Ben Reid (AI Forum NZ), **Matt Boyd** (Adapt Research), **Anna** and **Kelly Pendergrast** (Antistatic), **Tony Greaves** (Eltanin), **Monica Collier**, **Meredith Jeory** and **Liam Landon** (IDC New Zealand), for compiling local New Zealand case study interviews and research. **Hayden Glass** (Sapere). **Louise Taylor** and **Matt Bartlett** (Simpson Grierson), **Frith Tweedie** (EY Law) for contributing to sections on AI, Law and Regulation. **Graeme Muller**, **Chantal Thomas**, **Frances Barrett**, **Andrea Molloy** (NZTech) and **Martyn Rivett** (Propeller).

Steven McCleery and **Kylie Ormsby** at MOSHA for design and digital publishing.

Megan Tapsell (ANZ) for supporting our Te Reo translations.

Partners

The many staff from our project partners and AI Forum working groups and members who took time to provide in depth reviews and feedback throughout the process.

Thank you to the many organisations that supported the research financially and by providing case studies.

Please download a free digital copy of the e-report from the AI Forum website, www.aiforum.org.nz

Completing our name

Te Kāhui Atamai Iahiko o Aotearoa

Since our last report was released in 2018, we have formally adopted a name in Te Reo Māori: **Te Kāhui Atamai Iahiko o Aotearoa**.



Te Kāhui Atamai Iahiko o Aotearoa



The coining of the name was led by Dr Valance Smith, Assistant Pro Vice Chancellor of Māori Advancement at the Auckland University of Technology, with support from Tina Wilson of New Zealand Trade and Enterprise (NZTE) and Megan Tapsell of ANZ, who is also an advisor to the AI Forum's Māori Engagement Subcommittee.

Developing the name was a journey, in large part because there was no translation of the concept of Artificial Intelligence to use as a starting point. In Te Reo, words aren't always direct English translations, but rather aim to encapsulate the essence of the message and be consistent with Te Ao Māori – the Māori worldview. Other translations were considered for the word 'artificial', but the team felt that Electrical Intelligence best described the concept of AI, and has a better chance of adapting as the way AI is defined changes over time.

We hope that the words **Atamai Iahiko**, which share initials with Artificial Intelligence, are used widely across Aotearoa.

In addition, Dr Smith also translated the AI Forum vision, 'Harnessing the power of AI for a prosperous, inclusive and thriving future New Zealand.'

Tahuna te rama o te Atamai Iahiko kia mārama ai te ara o tōnui, o whakawhāiti, o momoho mā Aotearoa o āpōpō.

"Ignite the torch of AI to throw light upon the path of prosperity, of inclusiveness, of success for the Aotearoa of tomorrow."

Contents

**About the AI Forum
of New Zealand..... 02**

Acknowledgements 02
Partners..... 02

**Completing Our Name: Te Kāhui
Atamai Iahiko o Aotearoa.....03**

**Towards Our Intelligent Future:
An AI Forum Research Project 09**

Introduction: AI Forum of New Zealand..... 10
Foreword: Ministry of Business,
Innovation & Employment.....11

Key Highlights.....12

Executive Summary15

Towards Our Intelligent Future..... 15
New Zealand Shouldn't Wait for a
National AI Strategy 16
Aligning AI with New Zealand's Objectives.... 16
Introduction to AI in a Māori Context..... 19
Call to Action..... 20

PART A: An AI Roadmap for New Zealand

21

1. Understanding and Explaining AI 23


The Newest General Purpose Technology 23
Approaches to AI26
Narrow vs General AI.....26
What About the Robot Overlords? 26

CASE STUDY:
Financial Services, Xtracta



27

CASE STUDY:
Customer Services, Ambit




28

AI Realities..... 29
Everyday AI – What Does this Look Like
in the Real World?29
The Components of Machine Learning32
Machine Learning Considerations.....32
Neural Networks and Deep Learning.....33
Machine Learning Frameworks.....34
AI as a Service.....35

2. Recent AI Developments 36

International AI Developments 36
 Technological Advances36
 The Chips are up37
 Quantum AI Arriving Soon?.....38
 Future AI Trends38
 Vast AI Investment And Concentration
 of AI Superpowers39
 AI Investment by Governments.....41
 Explosive Market Growth42
 AI For Good43
 2030 Agenda for Sustainable Development.....43
Google: AI For Social Good 44
Microsoft: AI For Earth 45
 Solving Climate Change With AI.....46
 Movement Towards Regulation of AI46

CASE STUDY:
 Customer Service, FaceMe



49

CASE STUDY:
 Banking and Finance, ANZ




50

Crescendo if AI Ethics.....51
New Zealand AI Developments 56
 Policy Developments 56
 Ecosystem Advances.....58
 Investment 60
 Inactivity in the Last 12 Months61

3. An AI Roadmap for New Zealand 62

Placing AI in a New Zealand Context 62
 Wellbeing.....62
 Wellbeing Indicators.....64


CASE STUDY:
 Environmental, Orbica



65

Māori Wellbeing 67
Sustainable Development..... 67
 New Zealand Sustainability Goals67
Economic Growth..... 67
 International Economic Analysis.....69
 Impact by Region.....70
AI and New Zealand’s Economy 72
**Economic Complexity as an
 Indicator of Income Inequality 73**
AI and Public Policy74
 Strategic Responses and Decisions74
 OECD Recommendation on AI.....76

CASE STUDY:
 Public Sector, ACC



77

CASE STUDY:
 Social Enterprise, Sustainable Coastlines



78

First Steps for AI Policy in New Zealand.....79
 A Policy Map for AI in Aotearoa79
 What Does a Future Without AI
 Policy or Regulation Look Like?.....84
New Zealand Shouldn’t Wait for an AI Strategy..... 84
 Aligning AI to New Zealand’s
 Objectives: A Framework for Action85

Contents

Health 86
 Reducing Deaths from Breast Cancer..... 86

Transport..... 88
 Reducing Deaths from Road Accidents.....88

Conservation..... 89
 Improving Productivity for Wilding
 Conifer Management89

Call to Action – AI Roadmap Next Steps.....90

PART B:
 Enabling AI Adoption
 in New Zealand

91

4. Research..... 92

Case Study: Healthcare, Percision Driven Health...93

Case Study: Healthcare, Volpara Solutions.....93

Examples of AI Research in New Zealand..... 95

 Fundamental and Applied AI Research 95

 Government Initiatives 95


 Universities97

 Industry 99

**Where to Turn for AI Research Support
 in New Zealand**..... 100

How is New Zealand Tracking with AI Research?... 102

CASE STUDY:
 Agriculture, Robotics Plus

 **103**

CASE STUDY:
 Science, Agresearch

 **104**

New Zealand AI Research Opportunities..... 105

 AI and Health Research105

 AI and Agriculture Research.....105

 AI and Creative Industry Research.....105

AI Research as a Career..... 105

5. Trusted, Available Data..... 106

Data Governance 107

 Personal Data Rights.....107

 Data Privacy108

 Big Tech and Privacy.....108

China’s Social Credit Score..... 110

 Data Ownership.....110

 New Zealand Action..... 112

 A Social Contract for Data Use?..... 112

Māori Data 113

 Examples Of International Action on Data Rights
 and Ownership..... 113

CASE STUDY:
 Travel and Logistics, Air New Zealand

 **115**

 Data Trusts117

 Open Trusted Data Ecosystems117

 Data Trusts in New Zealand 118

Data as Infrastructure 119

 Data Standards..... 120

 Standards for Development of AI..... 121

 Data Standards: CDDQOS 121

 Data Standards In New Zealand 122

 Sufficient Data.....122

 Quality Data.....122

 Public Data Accessibility In New Zealand..... 123

Sia Data Exchange.....	124
Data Ventures.....	124

CASE STUDY:

Telecommunications, Spark



125

6. Talent, Skills and Capabilities127

Talent Demand and Supply 127

Global Demand for AI Talent Currently	
Outstrips Supply.....	127
Industry Demand for AI Talent Could Hamper	
Research Efforts at Universities.....	129
Few People Self-Report AI or Machine Learning	
Skills in New Zealand.....	129
There is a Lack of Diversity in the AI Talent Pool...129	
Increasing Education Opportunities and	
Facilitating a Talent Ecosystem May	
Help Ease the AI Talent Bottleneck.....	129
Andrew Ng's 5-Phases in His AI Playbook	130
New Zealand AI Talent Shortage?	130
The Struggle for Some New Zealand	
Organisations to Acquire AI Talent.....	131

Growing AI Talent in New Zealand132

Immigration.....	132
Skills Development.....	133
New Zealand Secondary Education	133
Edmund Hillary Fellowship	135
Shadowtech by Techwomen.....	136
AI for Education.....	136
A New Zealand Principal's Viewpoint.....137	
New Zealand Tertiary Education.....	138
Diversity.....	138
The Importance of Diversity.....	138
Self-Taught Developers	139
Non-Technical Skills	139

7. AI, Ethics, Regulation And Society 141

AI and Ethics141

Implementing Ethical AI142

Fairness and Accountability.....	142
Transparency and Explainability.....	142
Safety.....	142

AI, Law and Regulation.....144

AI and Freedom of Speech146

How Youtube Uses AI Technology to	
Detect Violent Extremism	146
The Christchurch Call	
– Stopping Extremist Content Online.....	147

AI and Privacy.....148

AI and Digital Divides.....148

AI and Inequality..... 150

In Summary 150

8. AI Readiness and Investment 151

Are New Zealand Organisations Prepared for AI?.. 151

AI Adoption in New Zealand.....152

Māori Business and Organisation Perspectives.....154

Strategy, Maturity and Change..... 155

Why Do AI Projects Fail?.....	157
-------------------------------	-----

Coordinating Organisations Through Partnerships...157

Conclusion..... 158

Supporting Materials/Appendices.. 159

Understanding and Explaining AI:

Machine Learning 159

Algorithms.....	159
Learning Styles.....	159
Data.....	160
Models.....	162

Mathematical Methods for Machine Learning 162

Mathematical Methods for	
Supervised Learning.....	162
Mathematical Methods for	
Unsupervised Learning	163

The Research Team..... 164

Methodology..... 165

Corrections policy..... 165

Licence 165

References..... 166

Contents

Figures

FIGURE 1: New Zealand's AI Roadmap	01
FIGURE 2: A framework for using AI to achieve New Zealand's objectives.....	17
FIGURE 3: Explaining AI – A Taxonomy of Concepts and Terms	22
FIGURE 4: AI Function Areas	27
FIGURE 5: Machine Learning	30
FIGURE 6: Total estimate equity investments in AI startups, by startup location 2011-17 and first semester 2018	38
FIGURE 7: AI Market estimates (Billion US Dollars).....	40
FIGURE 8: Regulation of Autonomous Vehicles Worldwide.....	47
FIGURE 9: The Partnership on AI illustrates the power of a diverse partnership with commercial and non-commercial members.....	53
FIGURE 10: New Zealand Treasury Living Standards Framework	62
FIGURE 11: Illustrates how GeoAI can identify and classify buildings, such as houses and garages from imagery.....	63
FIGURE 12: S-curves: AI adoption in New Zealand – Rapid, Medium or Slow?.....	66
FIGURE 13: Previous technologies have followed a range of 's-shaped' adoption curves, with the internet showing very rapid take-off, and TV a dramatic acceleration phase	67
FIGURE 14: Two phases of economic gains from AI: first increased productivity, then increased consumption	68
FIGURE 15: New Zealand a Frontrunner or a Follower in AI Adoption?.....	69
FIGURE 16: New Zealand employment by sector.....	70
FIGURE 17: "Heatmap" analysis of 18 different AI strategies across 8 areas of public policy.....	72
FIGURE 18: ACC's Cover Decision Algorithm.....	75
FIGURE 19: Aotearoa AI Policy Framework.....	78
FIGURE 20: A framework for using AI to achieve New Zealand's objectives.....	83
FIGURE 21: Health.....	85
FIGURE 22: Transport.....	86
FIGURE 23: Conservation	87
FIGURE 24: The New Zealand research funding landscape.....	98
FIGURE 25: New Zealand spend on R&D as a percentage of GDP compared to selected OECD countries	100
FIGURE 26: The current landscape of data governance	105
FIGURE 27: Transparent Data Use Dial	110
FIGURE 28: A new framework for data governance – Data Trusts, Data Coops and Data Commons.....	116
FIGURE 29: Data as the "Infrastructure of the Future".....	117
FIGURE 30: Characteristics of data when it is treated as infrastructure.....	118
FIGURE 31: The number of machine learning engineers has increased ten-fold.....	126
FIGURE 32: Lack of data science talent is 2nd greatest barrier.....	126
FIGURE 33: Elements of AI course created in Finland and freely available worldwide.....	131
FIGURE 34: The numbers game: broad interest in AI fundamentals at school will produce AI professionals.....	132
FIGURE 35: How did you learn to code? Hacker Rank Skills Report 2018	137
FIGURE 36: Floridi et al: Five AI Ethical Principles	140
FIGURE 37: High-level timeline for Government action toward digital inclusion	147
FIGURE 38: IDC research New Zealand / Australia comparison	151
FIGURE 39: AI Maturity Model with Indicative Return on Investment.....	153
FIGURE 40: Main steps for training and deploying machine learning models	160

Tables

TABLE 1: Future Today Institute AI Trends 2019.....	36
TABLE 2: Announced investments in ten national/regional AI strategies	39
TABLE 3: Climate change solution domains, along with areas of ML that are relevant to each.....	44
TABLE 4: Declarations of AI Principles	50
TABLE 5: AI Forum working groups agenda 2018-2019	57
TABLE 6: Recent Venture Capital Investment in New Zealand AI	58
TABLE 7: Four Stages of AI Maturity.....	154
TABLE 8: Mathematical Methods.....	158
TABLE 9: Identifying AI Opportunities: Data Inputs and Machine Learning Outputs	159

Towards Our Intelligent Future

AN AI FORUM RESEARCH PROJECT

The AI Forum of New Zealand would like to extend our sincere gratitude for the generosity of all the Project Partners and Supporters who have made this report possible.

Principal Partners



MINISTRY OF BUSINESS,
INNOVATION & EMPLOYMENT
HĪKINA WHAKATUTUKI



Spark^{nz}

Health Sector Partner



Agriculture Sector Partner



Supporters



Research Partners



Introduction

AI Forum of New Zealand Te Kāhui Atamai Iahiko o Aotearoa

KIA ORA KOUTOU.

It's hard to believe that the AI Forum has only been in existence for just over two years. Starting from what felt like a national vacuum, within a year we published the 2018 report *Artificial Intelligence: Shaping a Future New Zealand* and placed a stake in the ground for all New Zealanders to engage with this transformational technology. Since then the Forum has continued to build a national focus on what many believe is the most significant technological and social development for generations: the large scale adoption of Artificial Intelligence.

In the last year the AI Forum has tripled its membership to 144 organisations and is now well represented across New Zealand. Universities, many government departments, global technology businesses and a wide range of corporates, SMEs and highly talented best-in-class startups are contributing to the conversation. We have established a set of working groups to monitor and advance the debate – including the development of practical ethical and legal principles to safeguard the rights and ambitions of all New Zealanders in the face of rapidly advancing technology.

Such representation coupled with a national spirit of collaboration is now materialising in our second piece of landmark research; *Towards Our Intelligent Future*. It is the result of over 9 months of parallel research work streams with our members and partners.

With a topic so large and impactful, the research has been broken up into two parts. This not only ensures that the key messages do not get lost in the detail, but it also aids the reader in consuming such information and encourages an interactive national debate on its findings. It builds upon our previous work with the aim of providing a more detailed, practical guide for Aotearoa New Zealand to achieve positive AI-enabled outcomes – for wellbeing, for sustainability and for the economy.

In addition to exploring the many business and productivity aspects of AI, this research emphasises the potential contributions AI could make towards solving some of our world's greatest



STU CHRISTIE, Chairperson
AI Forum of New Zealand
Te Kāhui Atamai Iahiko o Aotearoa



contemporary challenges: climate change and clean energy production, biodiversity loss, healthcare, global inequality and education. It also contributes a working policy framework which provides line of sight for how AI can help support wellbeing outcomes for all New Zealanders.

We believe that there is still much to be done. As a trusted independent body, the AI Forum will continue to contribute towards the development of a cohesive national AI strategy, help moderate the national conversation and build the next generation capability for New Zealand's future state. Our near term aspiration is for Aotearoa's AI community to actively engage with confidence with their international peers and continue to develop a global narrative of cooperation, safety and opportunity.

A project of this size and scale would not have been possible without the support of many organisations. In particular, the AI Forum would like to thank its Principal partners, notably the Ministry of Business, Innovation & Employment, ANZ, Google, IAG, Microsoft and Spark. The AI Forum would also like to recognise and thank the many universities, technology firms and industry organisations that helped enable this project.

Ngā mihi nui

Foreword

Ministry of Business, Innovation & Employment

**WE ARE A COUNTRY OF INNOVATORS, A COUNTRY THAT SEIZES
NEW OPPORTUNITIES IN ORDER TO GROW AND TO PROSPER.**

In order to build trust in anything new or disruptive we must begin with a conversation. Since the AI Forum's initial report was released in 2018, we have moved that conversation on to discussing how we as a country can harness the innovation potential of Artificial Intelligence (AI).

We are continuing to build our understanding of the impacts AI can have on our people, regions, businesses and environment. While caution is still needed as with any new disruptive technology, we are looking towards the realisation that applying this new technology will drive economic growth, empower workers and enhance the wellbeing of New Zealanders.

I am pleased to say that many New Zealand companies are already at the forefront of AI adoption and application. We have innovative companies working both in New Zealand and overseas helping firms to utilise the power of predictive analytics for competitive advantage, and we have companies that are leading the way on industrial robotics and digital humans.

While it is essential we are having these discussions on how AI will have many positive benefits to our everyday lives, we must also make sure that we are encouraging all New Zealanders to grasp opportunities to develop their technological skills.

The Ministry of Business, Innovation and Employment is proud to have supported the AI Forum to undertake further research into AI. Through embracing AI we can ensure that our people and businesses are prosperous and adaptable while supporting a sustainable economy and environment. I look forward to the ongoing conversation.



CAROLYN TREMAIN,
Chief Executive
*Ministry of Business,
Innovation & Employment*



**MINISTRY OF BUSINESS,
INNOVATION & EMPLOYMENT**
HĀKINA WHAKATUTUKI



Key Highlights

THERE IS NO SHORTAGE OF MAJOR CHALLENGES TO WHICH AI TECHNOLOGIES CAN BE APPLIED.

AI can be targeted towards improving New Zealand's wellbeing outcomes and solving complex social and environmental challenges. Relatively small investments in data and AI solutions could yield substantial gains and provide the pull through for the AI skills and capabilities which New Zealanders will need in the future.

WELLBEING

01

ALL POLICY MAKERS IN NEW ZEALAND NEED TO HAVE A PERSPECTIVE ON AI AND HOW IT IMPACTS DECISIONS IN THEIR DOMAIN.

As a general purpose technology, AI and its many applications touch on public policy portfolios across all areas of government and society. Data science and AI can provide policy makers with unprecedented insights by modelling complex systems and scenarios and evaluating hard-to-measure policy outcomes.

PUBLIC POLICY

02

IT IS IMPORTANT THAT MORE NEW ZEALANDERS HAVE A HIGH LEVEL UNDERSTANDING OF AI.

This provides a clearer link from the technology to its many applications. New Zealand can join other countries which have begun training their general populations with free online courses.

PUBLIC UNDERSTANDING

03

TE AO MĀORI – A MĀORI WORLDVIEW – SHOULD BE INCORPORATED INTO THE PRACTICE AND DEVELOPMENT OF AI – including increased understanding of indigenous approaches to data sovereignty, wellbeing and business.

TE AO MĀORI

04

NEW ZEALAND CAN PLAY AN ACTIVE ROLE IN FACILITATING DIALOGUE BETWEEN DIVERSE CULTURAL AND ETHICAL TRADITIONS LOCALLY AND INTERNATIONALLY, INCLUDING UNDERSTANDING AI ETHICS FROM AN INDIGENOUS MĀORI VIEWPOINT.

There are now over 32 declarations of AI ethics principles around the world. New Zealand is uniquely placed to contribute to the global AI ethics conversation, including understanding AI ethics from an indigenous Māori viewpoint.

ETHICS

05

NEW ZEALAND'S DATA CAN BE A NATIONAL ASSET AND A TAONGA.

Data is the fuel which powers AI. Data trusts are legal structures that enable the sharing of data for collectively-defined outcomes. In a New Zealand context, pilot projects to explore data trusts in sectors including agriculture, tourism and health could be prioritised to provide solid underpinnings for successful AI adoption.

TRUSTED DATA

06

AI IS LIKELY TO BE A KEY ENABLER OF GOVERNMENT POLICIES TO LIFT THE PRODUCTIVITY OF KEY SECTORS.

AI is the newest general purpose technology – like electricity, the combustion engine or the internet before, AI has the potential to transform productivity, increase consumption and accelerate economic growth.

ECONOMIC GROWTH

NEW ZEALAND STRATEGIC INVESTMENT IN AI IS LAGGING BEHIND OUR PEERS.

Over twenty countries have now embarked on national AI strategies with significant per capita investment over five or more years. International private equity investment in AI has skyrocketed over the last five years, but not in New Zealand. Although a small number of companies are now reporting successful returns from AI, increased investment in AI will be needed to successfully harness the opportunities.

INVESTMENT

AI IS BEING APPLIED AROUND THE WORLD FOR SOCIAL, ENVIRONMENTAL AND HUMANITARIAN OUTCOMES.

Including: solving climate change, cleaning up the world's beaches, increasing accuracy of flood forecasting, monitoring global fishing activity, solving climate change, predicting and preventing famine and increasing cancer detection rates.

AI FOR GOOD

GLOBAL DEMAND FOR AI TALENT CURRENTLY OUTSTRIPS SUPPLY.

Overseas talent may be needed in the short term, but immigration is unlikely to be a long term solution to New Zealand's AI talent needs. Around the world, countries are investing in AI education at school and tertiary levels. In comparison, New Zealand is behind in providing resources specifically targeted to AI skills development and has yet to develop a programme of systematised AI education.

SKILLS AND TALENT

THERE IS A LACK OF DIVERSITY IN THE AI TALENT POOL WORLDWIDE.

In particular in New Zealand, there are very few women, Māori and Pasifika currently working in AI.

DIVERSITY

MANY EXISTING NEW ZEALAND LAWS APPLY TO AI, HOWEVER SOME NEW LAW MAY BE NEEDED.

Currently there is negligible AI-specific regulation worldwide, but AI has implications for laws relating to bias and discrimination, public sector use of AI and algorithms, copyright, employment, robo-advice, consumer protection, collusion, privacy, freedom of expression, autonomous vehicles and lethal autonomous weapons.

LAW AND REGULATION



Executive Summary

New Zealand shouldn't wait for an AI strategy to arrive – we need action now to apply AI to the wellbeing, sustainability and economic objectives we already have.

AI IS NOW BEING ADOPTED THROUGHOUT NEW ZEALAND, PARTICULARLY IN THE PRIVATE SECTOR, YET THERE IS SLOW MOMENTUM ON STRATEGIC NATIONAL AI POLICY AND INVESTMENT.

In 2018, the AI Forum published its first major research report, *Artificial Intelligence: Shaping a Future New Zealand*. The aim was to unite the diverse voices of an emerging national conversation about AI and how it could affect our country. Last year's analysis of the importance of AI for New Zealand identified six major recommendations, with the main one being the importance of forging a coordinated national AI strategy.



Since then, internationally, AI technologies have continued to evolve and diffuse at accelerated speed and scale.

Over twenty of New Zealand's peer nations have now defined and begun implementing coordinated national multi-year AI strategies focused on growing AI capabilities – in areas such as investment, skills, data, infrastructure, research – to accelerate AI adoption. Economic modelling has identified that early adopter nations stand to gain significantly more than those countries arriving late. There have also been considerable advances in how to use AI for good – social and environmental causes – as well as a rich debate about the ethical principles and regulation which will be needed to govern applications of this new technology.

One year after the release of the *Shaping a Future New Zealand* report, this deeper analysis finds that AI is now being adopted throughout NZ, particularly in the private sector, yet there is slow momentum on strategic AI policy and investment.

Progress towards the envisioned national AI Strategy has not advanced as rapidly as initially hoped. New Zealand's relative progress towards becoming an AI enabled nation increasingly appears too slow, when compared to international developments and domestic requirements.

Towards Our Intelligent Future

This new report, *Towards Our Intelligent Future*, builds upon our previous work with the aim of providing a more action-oriented, practical guide for Aotearoa New Zealand to achieve positive AI-enabled objectives. In particular, the report aims to:

- stimulate momentum towards national investment in AI to achieve high priority objectives.
- analyse key enablers to support AI adoption in Aotearoa.
- provide a first examination of Māori attitudes, perceptions and engagement with AI.
- provide case studies of AI in action across many facets of our economy and society.

But, above all, we have aimed to inform our ongoing national conversation.

We have structured the report into two parts which we summarise here.

Part A – An AI Roadmap For New Zealand looks to build a better understanding of what AI is, narrates important international and national developments which will influence AI uptake in the future and then proceeds to place AI firmly within a New Zealand context, aligning its applications to defined national objectives and presenting a framework for

national investment in data and AI, which can be considered as "the infrastructure of the future".

In **Chapter 1 – Understanding and Explaining AI** we provide a deeper overview of AI – the technologies, mathematical theories and above all the new applications that it now enables. We have provided a taxonomy of the main concepts involved with AI – not exhaustive and not without ongoing iteration – but a schema which supports non-technical readers to scaffold their understanding of the major concepts behind AI.

Over the last year the international backdrop against which New Zealand's AI capabilities will be measured have moved at an incredible pace. The AI Forum has been continuously tracking these developments – new technology announcements, training courses, economic impact reports, innumerable web articles and publications. We have summarised the major trends and events in **Chapter 2 – Recent AI Developments** – this provides an international and national perspective on which New Zealanders can build their understanding of AI opportunity.

And then in **Chapter 3 – An AI Roadmap For New Zealand** we place AI into a uniquely New Zealand context, illustrating clear line-of-sight from national AI investment to achieving New Zealand's national objectives for wellbeing, sustainability and for the economy.

New Zealand Shouldn't Wait for a National AI Strategy

The research identifies that there is a massive, immediate, opportunity for New Zealand to invest in AI that helps it achieve its existing strategic priorities: for example, managing the impact of climate change, achieving the "just transition" to renewable energy, increasing productivity growth and improving citizen wellbeing.

AI enabled solutions can be specifically harnessed towards the Wellbeing outcomes that the Government is now targeting, as well as the UN's 17 Sustainable Development Goals. The Government's new Industry Transformation Plans will also provide the economic objectives which AI can be targeted towards – so "AI for Wellbeing", "AI for Sustainability" and "AI for the Economy" are key themes.

Relatively small investments in data and AI solutions could yield substantial gains and provide the pull through for the AI skills and capabilities which New Zealanders will need in the future. As a small nation, New Zealand has limited resources and this makes it imperative that all AI investments are well targeted. Equally, as an agile nation, there are compelling reasons to remain optimistic for the opportunities to capitalise on focused AI applications – while also contributing significantly to international AI governance and policy development discussions.

Aligning AI with New Zealand's Objectives

In Chapter 3 we introduce a framework which supports more rapid momentum for New Zealand to invest in AI-enabled solutions to achieve objectives we already have.

The framework contains four layers:

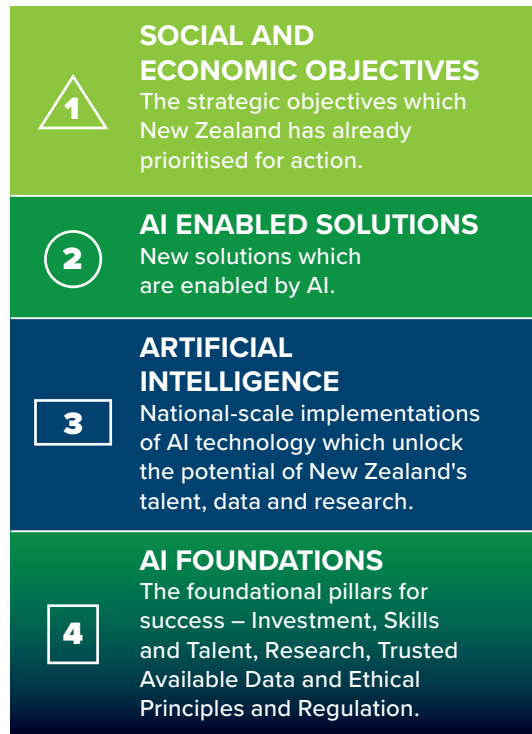



FIGURE 2: A framework for using AI to achieve New Zealand's objectives



In **Part B – Enabling AI Adoption In New Zealand** we turn our attention to the core foundational pillars of AI.

Chapter 4 – Research provides an outline of current AI research engagement across New Zealand’s research institutions, including research funding from government institutions and the private sector. Fundamental and applied AI research is a key driving force for successful adoption, providing support for the raw resources of a vibrant national AI ecosystem. New Zealand has diverse AI research happening across a multitude of domains from health to agriculture, evolutionary computation to public policy.

In **Chapter 5 – Trusted, Available Data** we examine the role of data for successful AI adoption. Trusted data can be a national asset and a taonga. New Zealand will require large quantities of reliable, trusted data, together with data standards and data governance systems. Difficult contemporary challenges relating to protections for personal data will need to be resolved. Maintaining privacy has broader benefits for society over and above protections for individuals. Meanwhile market opportunities from sharing commercial data for collectively-defined outcomes should be directly addressed. We explore the concepts of data trusts, data coops and data commons and how they can be applied to practical solutions, providing frameworks for data governance which enables collective value to be delivered while ensuring that data rights are respected.

In a New Zealand context, pilot projects to explore data trusts in agriculture, tourism and health could be prioritised to provide solid underpinnings for successful AI adoption in these sectors.

Chapter 6 – Talent, Skills and Capabilities looks at the imperative to nurture and grow a skilled, capable and diverse AI talent pool against a background of international competition. Accelerating investment into AI skills development in school, tertiary and continuous lifetime learning will be essential. Incorporating AI into the New Zealand Digital Technologies/Hangarau Matahiko curriculum roll-out is a key opportunity to be seized. Meanwhile continued monitoring that immigration settings are sufficiently balanced to solve near term skills shortages is needed. Finally, there is a desire to ensure people from all backgrounds are present within Aotearoa’s AI talent pool – we

find that while there appears to be relatively high ethnic diversity generally, there are very few women, Māori and Pasifika currently in AI in New Zealand.

Chapter 7 – AI Ethics, Regulation and Society examines more deeply the new challenges in ethics and the law which AI raises. Existing legislation will need to be reviewed to ensure it is fit for purpose in the new context – and similarly society will need to revisit debates around free speech, privacy, inclusion and protecting democracy in light of potentially harmful applications of AI technologies. Here in New Zealand, it will be particularly essential to include the Te Ao Māori (Māori world view) perspective into any ethical AI guidelines proposed. A vocal and vibrant conversation on AI Ethics has reverberated around the world and in its progressive tradition New Zealand is contributing strongly about the role of data and AI within our own – and international – society.

And finally in **Chapter 8 – AI Readiness and Investment** we turn our attention to how ready New Zealand organisations are to begin adopting AI. Research from our partner IDC indicates that although 44 percent of New Zealand organisations are planning to adopt AI within 2-5 years, only 1.2 percent have actually adopted AI today. As illustrated by the case studies contained within this report, a number of New Zealand organisations are now actively working with AI and beginning to yield results from their investments. However the depth and breadth of AI across various industry sectors is still a small part of the economy. Tools such as AI maturity models will be helpful for organisations to understand where to invest – and have patience for the returns to arrive.

The report concludes with a call to action for New Zealand’s government, private sector, iwi and tertiary institutions to collaborate and advance the settings for AI success in New Zealand – working together in partnership *Towards Our Intelligent Future*.

Introduction to AI in a Māori context

AS AI BECOMES WIDESPREAD ACROSS BUSINESS AND SOCIETY, IT WILL HAVE PARTICULAR IMPACTS FOR MĀORI. EXPERTS, ACTIVISTS AND MĀORI BUSINESS LEADERS ARE BEGINNING TO EXPLORE THE OPPORTUNITIES AND CHALLENGES THAT AI PRESENTS.

A Māori worldview can be incorporated into the practice and development of AI, and innovators like Te Aroha Grace, Innovation Officer at Ngāti Whātua Ōrākei, have discussed how AI may support indigenous mana' and provide opportunities to develop a unique, culturally informed AI ecosystem in New Zealand. Currently the Māori AI landscape is in its early stages, but exciting projects are emerging.

Work has begun on understanding indigenous approaches to AI. Recently, a New Zealand delegation from Te Hiku Media and Dragonfly Data Science attended a workshop on indigenous protocols and AI in Hawai'i.² There are more developed conversations regarding Māori data sovereignty with numerous publications and a data hui exploring the issue. As AI is developed and implemented in New Zealand, it is important that the Treaty of Waitangi partnership is honoured, with Māori being included in any decision making on AI policy and regulation. These reports provide further investigation of indigenous data sovereignty, wellbeing frameworks and business/organisation approaches.



Call to Action

NEW ZEALAND SHOULDN'T WAIT FOR AN AI STRATEGY TO ARRIVE. WE NEED ACTION NOW TO APPLY AI TO THE WELLBEING, SUSTAINABILITY AND ECONOMIC GOALS WE ALREADY HAVE.

Our previous recommendations in last year's *Shaping a Future New Zealand* report remain more relevant than ever, but are still in the early stages of being absorbed and responded to.

Thus we have consciously refrained from promoting another set of recommendations in this new report.

We hope instead to stimulate momentum with the materials and tools contained herein, enabling greater understanding, practical engagement and action.

We make the following suggestions to start work immediately:

1. AI FOR WELLBEING

The role of AI should be consciously considered for all of New Zealand's priority policy objectives.

2. AI DRIVES PRODUCTIVITY GROWTH

AI should feature in every one of the Industry Development Plans for economic growth.

3. FLAGSHIP AI PROJECTS

New Zealand needs to collectively and quickly identify a manageable number of high-value flagship opportunities to apply AI – for example in health, transport and the environment.

4. PARTNERING ON AI INVESTMENT

Leverage public-private-iwi-academic partnerships to invest in these flagship projects.


5. CREATE DEMAND FOR FOUNDATIONAL AI CAPABILITIES

Use these AI investments to accelerate the pull-through of core AI capabilities needed for New Zealand – this will inform the development of a pragmatic strategic approach – rather than waiting for a formal national AI strategy to emerge fully-formed.

Towards Our Intelligent Future **Te ara mō tātou atamai o āpōpō**

PART A:

An AI Roadmap
for New Zealand



1. Understanding and Explaining AI

“Everything should be made as simple as possible, but not simpler.”

ALBERT EINSTEIN (attributed).

IT IS IMPORTANT THAT MORE NEW ZEALANDERS HAVE A HIGH LEVEL UNDERSTANDING OF AI SO THAT THERE IS A CLEARER LINK FROM THE TECHNOLOGY TO ITS MANY POTENTIAL APPLICATIONS.

Understanding and explaining Artificial Intelligence (AI) for a mainstream audience remains a challenge. Despite the considerable potential of AI, the lack of general understanding of the fundamentals of AI is an obstacle when looking for AI enabled solutions to tangible problems.

Here we aim to demystify AI and help readers understand how these complex but exciting technologies may serve their business, organisation or community.

The Newest General Purpose Technology

AI can be considered as a **General Purpose Technology (GPT)** which, like electricity, the combustion engine or the internet before, has the potential to transform productivity and accelerate economic growth. Furthermore, in the new socioeconomic vocabulary of the early twenty first century, AI has the potential to improve wellbeing across the whole of society.

AI technologies already permeate many aspects of our lives, both directly and indirectly. In recent years, a combination of huge datasets, increased computing power and the availability of advanced learning algorithms have enabled rapid advances in what AI can do.³ This includes considerable breakthroughs in the ability for AI programs to recognise images and sounds, solve complex problems, translate languages, power autonomous robots and analyse the sentiment of people’s speech. As a result of these advances, tasks that previously required

human cognition to carry out are becoming widely automated through businesses and organisations.

Despite its potential and wide usage, it can be difficult to succinctly define what AI is. AI is a “suitcase term” which can be used to refer to everything from neural networks to autonomous robots in sci-fi thrillers to the search engine you use to look up pictures of cats.

In its recently adopted Recommendation on Artificial Intelligence, the Organisation for Economic Cooperation and Development (OECD) defines an **AI system** as:

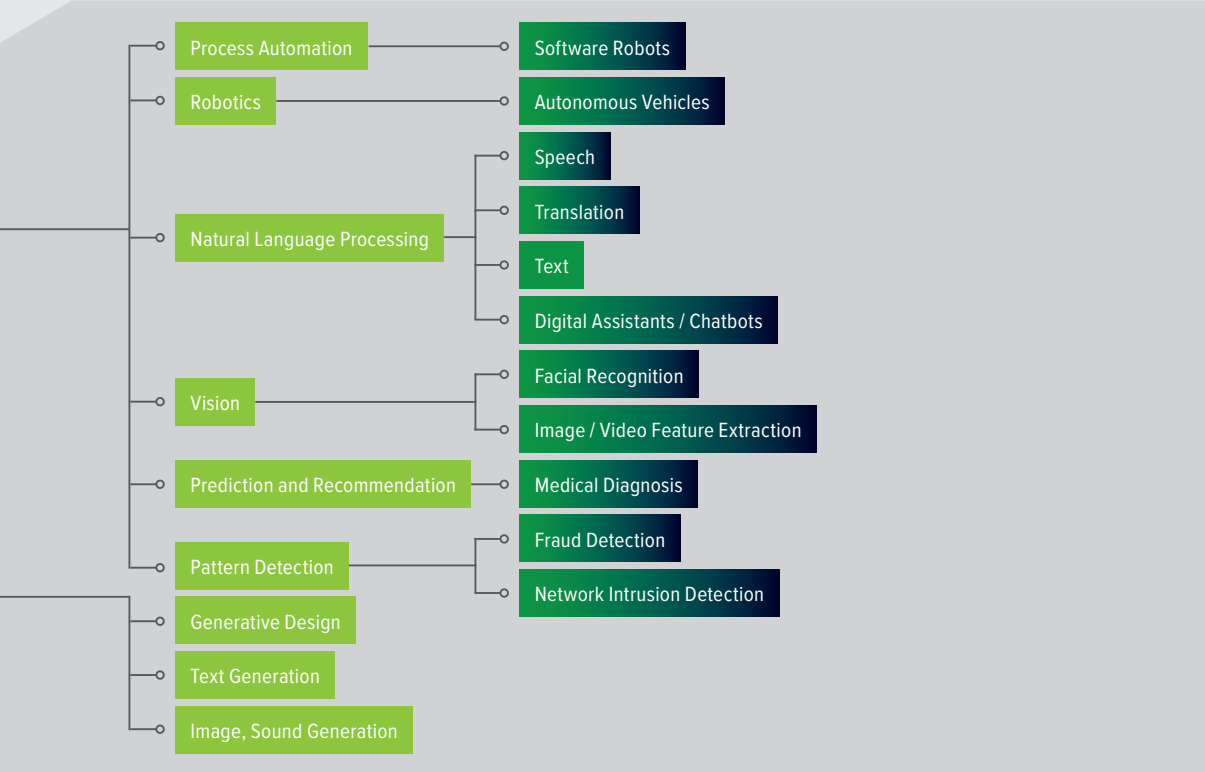
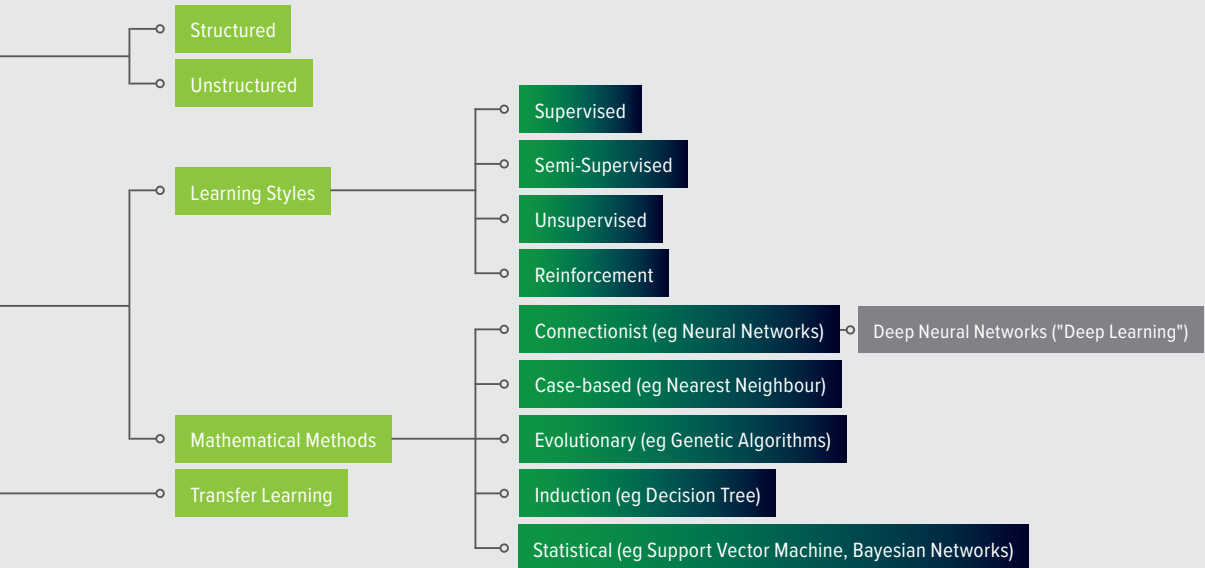
“(An AI system is) a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.”

OECD, 2019

Whether through image analysis, natural language processing, process automation, autonomous vehicles or making better predictions, there is an infinite variety of use cases for which people can deploy AI.

FIGURE 3: Explaining AI – A Taxonomy of Concepts and Terms





APPROACHES TO AI

The theory and application of AI as we know it today has been around for decades. The recent explosion of interest in AI – particularly machine learning – is largely due to three factors: the availability of large datasets (including large labelled datasets), increases in parallel computer processing power, and the availability of improved machine learning algorithms that anyone can use, often through open source sharing. This has allowed increasingly complex problems to be solved, more quickly and at lower cost.

The family of technologies currently referred to as AI include many different approaches. Historically, rules-based or symbolic AI (“Good Old Fashioned AI” – GOFAI) such as Expert Systems were the leading paradigm for AI. Other popular AI approaches include planning and scheduling, knowledge representation and reasoning. However, many of the most interesting applications of AI today result from approaches that use machine learning.⁴

Machine learning is a subset of artificial intelligence which often uses statistical techniques to give computers the ability to “learn” with data – creating models to process, interpret, and respond – without being explicitly programmed with a predefined set of rules.

There are many machine learning techniques, each one with its own specific capabilities and applications. We discuss the techniques and technologies behind Machine Learning in more detail in the Appendix.

NARROW VS GENERAL AI

The machine learning AI of today isn’t very multitasking. Machine learning algorithms are trained to do a specific task, for example recognising objects in photographs or translating text from one language to another.

Artificial General Intelligence (AGI or Strong AI) is the as-yet hypothetical existence of intelligent machines which can successfully perform any intellectual task that a biological human can. While much current AI research is actively working towards general intelligence, for now AGI resides within the realm of science fiction.

Narrow AI (or Weak AI) is AI that is focused on one narrow task. All currently existing AI systems are Narrow AI.

What about the robot overlords?

THE AI WE KNOW TODAY IS VERY DIFFERENT FROM THE HUMANOID ROBOTS OR COMPUTER PROGRAMS THAT POPULATE MANY BOOKS AND MOVIES.

There are some people who think that AGI is possible in this century. One 2013 survey of AI experts⁵ reported the median estimate of respondents was that there is a 50 percent chance that high level machine intelligence will be developed around 2040-2050, rising to a nine in ten chance by 2075,⁶ although this survey has itself been called into question.

There are others who believe that AGI could be an existential threat for humanity.⁷ For the moment at least, AGI is confined to the future and the world of science fiction.

While considering the implications of AGI is a useful thought experiment and raises important questions as to how we develop, deploy and regulate AI systems today, we choose to focus our attention on the challenges and opportunities of narrow AI systems in wide use today.



CASE STUDY FINANCIAL SERVICES

Xtracta: Intelligent Process Automation (IPA)

Medium and large enterprises rank automation in the top three capabilities they expect from their digital workplace environments. In the past, robotic process automation (RPA) has helped increase accuracy and efficiency in a number of business tasks. However, Intelligent Process Automation (IPA) is a game changer that combines cognitive technologies and advanced analytics with RPA.

New Zealand company Xtracta leverages this technology to provide intelligent automated data capture and data entry services across multiple industries. Xtracta's products can read and process documents such as receipts, invoices, contracts, and other paper based or electronic documents.

The system uses several types of AI, including optical character recognition, natural language processing, image processing and Xtracta's own algorithms. The IPA is trained on photos of invoices and other documents, and has access to the company's accumulated knowledge pool. When combined with customer specific datasets, this allows the technology to read documents and provide real time, cloud based document classification, data extraction and data validation.

IPA is a step forward from existing RPA. An IPA system can be trained to identify non-standard data, and provide confidence levels for its assessment of difficult to interpret inputs.



Xtracta licenses vendors, such as MYOB and ABM[®] to integrate the technology into their own products. Clients can also build in processes for data validation and to teach the AI to recognise potential fraudulent activity.

The technology also helps employees to work smarter. For example, a major healthcare provider which builds and operates retirement villages in New Zealand and Australia has automated the processing of approximately 15,000 invoices per month. The healthcare provider uses Xtracta's API together with their financial system to digitise physical invoices. Estimates suggest savings of at least ten hours manual work per day through the automated processing.

Enterprises can address how RPA will augment their business and coexist with human labour. Enterprises should consider implementing a single RPA process within a single business unit and then build a road map that allows for scaling.

Xtracta
your data engine

Learn more at www.xtracta.com

CASE STUDY CUSTOMER SERVICE

AMBIT: Chatbots



New Zealand scale-up Ambit deploys digital employees on its conversational AI platform, and delivers a solution within weeks. Digital employees operate through a conversational interface and reduce the effort and time for a customer to search a website, make enquiries on the phone, or queue at a store. Ambit's platform is suitable for any industry and customers include telecommunications, utilities, financial, professional and government businesses.

Ambit's platform uses natural language understanding (NLU), machine learning, and contains a series of language models that ensure a digital employee is on message, relevant to the audience, and provide a near human interaction.

For customers seeking digital employees, an engagement with Ambit includes two parts. First, the customer works with the Ambit team (or a partner) in conversation design workshop. This lets the customer build ideas about what they want the conversation experience to be like and how that might look across the entire business. These ideas are then validated. Once outputs are developed Ambit shifts into the design, build and integrate phases.

Ambit is on a growth trajectory and continues to build out the capability of its platform. This year, it will add customer experience management functionality, 'automated' conversation design and deeper language processing capability.

Using data obtained from the chatbot and other conversation channels, the platform will provide actionable insight across customers businesses.

To be ready for digital employees within your organisation, it is recommended that your organisation:

- has digital customer engagement as a key business strategy
- can expose data sets essential for the digital employee to personalise and contextualise experiences
- can deploy security and privacy measures to maintain trust in your digital employee.

Ambit's single biggest challenge now is future growth, scaling and tackling the international market.



Learn more at www.ambit-ai.com

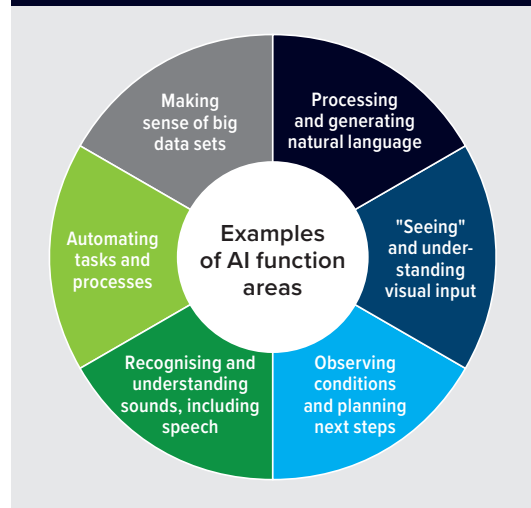
AI Realities

To understand how AI can and will affect our everyday lives, it's helpful to understand some of the key functions that AI can perform and the resulting products and tools. Example function areas are shown in the diagram opposite.

EVERYDAY AI – WHAT DOES THIS LOOK LIKE IN THE REAL WORLD?

The use cases detailed below show some of the existing applications of AI. The section also describes the ways that AI has been mobilised to help increase analytic capabilities, optimise business and create new opportunities. Of course, not all opportunities are equally positive, so the examples also note how AI can have unintended consequences or be damaging to individuals or the public good.

FIGURE 4: AI Function Areas



AI USE CASES: A snapshot of what we are seeing now

SUPPORT CHATBOT



Chatbots are programs that conduct conversation with users via text or voice. They are designed to simulate conversation, so users can interact with them as they would a human. Chatbots are increasingly being used to realise positive social outcomes. This includes helping to find answers to simple / common questions about a person's legal rights, providing mental health support or companionship through responsive chat. Support chatbots are often used to provide assistance to people who cannot afford the services of professionals such as a lawyer or counsellor, or do not feel comfortable sharing their problems with a human. This support is presently limited, however and cannot understand context or answer complicated questions.

Examples

- **OSCAR** from Air New Zealand assists customers with commonly asked queries.
- **Rentbot** from New Zealand based CitizenAI provides users with answers to tenancy related questions through Facebook messenger.
- **Woebot** provides Cognitive Behavioural Therapy (CBT) lessons and mood tracking.
- **Clearhead** is a mental wellness portal which uses a chatbot to screen users for depression and anxiety, recommending mental health resources and booking professional help.

PREDICTIVE MAINTENANCE



Predictive analytics helps companies understand when a piece of equipment is likely to require maintenance, repair or replacement. Predictive maintenance systems can use machine learning or deep learning to analyse data via Internet of Things (IoT) sensors, audio or image data, to detect anomalies or make forecasts on upcoming repair needs.

Examples

- GE Digital created **Predix** to help oil and gas businesses create automated analytics models for predictive maintenance of industrial equipment.
- IBM's **Predictive Maintenance and Quality** allows predictive maintenance implementations for elevators and hydrants among other applications.

ENVIRONMENTAL MONITORING



AI assisted environmental monitoring expands the capacity for organisations and companies to review, monitor, and understand environmental data. These systems combine video, photographic, or satellite monitoring systems with AI analysis to perform tasks like categorising animal species or recognising changes in an environment.

Examples

- **The Cacophony Project** is a New Zealand based organisation that observes and automatically identifies predators using machine learning algorithms to help protect our native birds.
- **TrailGuard AI** uses motion activated cameras and AI technology to detect poachers entering African wildlife reserves.

ASSISTED DIAGNOSIS



Assisted diagnosis brings AI into a medical context to improve the speed or accuracy of medical diagnoses, or to help doctors consider all the possibilities of a set of symptoms. Assisted diagnosis models can analyse large volumes of medical records to discover patterns that can help doctors to quickly and accurately identify illnesses, or to review X-rays or mammograms. While this application is in its infancy, some studies have shown that AI models can be as accurate as doctors in making some diagnoses. Due to the huge volumes of data in the health profession, AI could have a multitude of uses. There are also ethical questions raised regarding the use of patient data and where it is appropriate for AI to play a role in medical diagnoses.

Examples

- **Doctors at the University of California, San Diego** have developed systems that can analyse eye scans for lesions, hemorrhages and other signs of diabetic blindness.
- AI company **Corti SA** listens to medical hotline calls, and uses machine learning to help detect a caller's cardiac arrest likelihood based on caller vocabulary, tone and background noise.

CONTENT PRODUCTION



There have been recent advances in AI models that can generate written content that is indistinguishable from a human.⁹ Current AI writing technologies can draft email subject lines, finish the sentences of a human writer, and produce full passages of text based on a short sample. Some news organisations are testing tools to help journalists draft articles. Other current and potential uses include drafting news articles, song lyrics, advertising copy, and assisting authors with their writing. Concerns have been raised about potential misuse of models that can convincingly write like a human, including to quickly spread disinformation online.

Examples

- **OpenAI's GPT-2 language model** uses unsupervised learning to produce coherent paragraphs and carry out rudimentary comprehension based on a small text sample.
- **Forbes** is testing a tool that drafts articles for journalists to review and edit.
- **Gmail's Smart Compose** feature suggests text to finish a user's sentences, as well as suggestions for subject lines based on the content of an email.

FACIAL RECOGNITION FOR SECURITY



AI facial recognition is a subset of biometric authentication that identifies people by measuring the unique shape and structure of their faces. It can work using 3D mapping, 2D feature recognition or skin texture analysis. This technology can be used either on a single person or on dynamic camera footage to identify or verify people. It has a range of security uses, including verifying a person's identity at airport customs checkpoints, as an alternative to locks and keys for people's homes, unlocking a smartphone, or recognising the movements of people in a city and identifying potentially problematic behaviour. The use of this technology has raised ethical discussions regarding acceptable uses and surveillance. Some cities are banning the use of the technology.

Examples

- **Many iPhones** allow their owners to unlock the phone using their facial biometric data.
- **Some airports** use facial recognition technology for customs processes, taking a photo of a traveller and checking it against their passport.
- **Yancheng prison** in China is reportedly using facial and movement recognition technology to produce a daily report about prisoners.

AUTONOMOUS ROBOTS



Autonomous robots perform tasks or behaviors with a high degree of autonomy, responding to environmental conditions and adapting accordingly. Autonomous cars, delivery and security robots all use AI to sense and respond to their surroundings, performing tasks safely without human intervention.

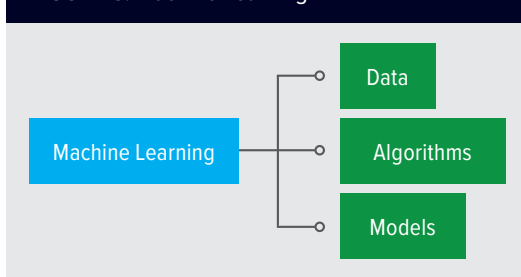
Examples

- **Self-driving** cars are being tested in multiple locations in the United States and also here in New Zealand.
- **NAVii** and other retail robots travel the aisles of large stores. Scanning shelves, they use machine learning and computer vision to identify out-of-stock and misplaced items.

THE COMPONENTS OF MACHINE LEARNING

With a task or desired outcome in mind and an appropriate set of **data**, a data scientist can decide which mathematical methods to apply. The data is fed into a piece of software that executes a training **algorithm**. This is a computational process that implements one or more mathematical methods, and can refine its process as new data is provided. After running the algorithm with training data, the data scientist has a **model** which can be used to make predictions on new data. The model is tested and refined, before being implemented in the field.

FIGURE 5: Machine Learning



MACHINE LEARNING CONSIDERATIONS

While based on mathematics, it's important to remember that a machine learning model is built on human decisions. Humans choose which data and mathematical methods to use and on what basis to make decisions about model life cycle, such as updates. All of these factors have a considerable impact on the outputs of a model, as does the quality and content of the data itself.

Considerations When Building Models

In *Weapons of Math Destruction*, Cathy O'Neil explains how models can encode human prejudice, misunderstanding, and bias into the software systems that increasingly manage our lives.¹⁰ Here are six important areas to be considered when building and implementing machine learning in a company or organisation. The importance of these considerations varies depending on the type of decision being made and the potential impact an error will have on people:

- 1) **Relevance of the data:** sometimes there won't be sufficient or appropriate data available. In these cases, it is possible to use proxy or stand-in data, or historical data to draw conclusions. However, each approach will impact on the accuracy or fairness of predictions.
- 2) **Explainability of the model:** some organisations like Google, IBM and Microsoft are including transparency or explainability into their ethical principles. According to the AI ethics guidelines global inventory,¹¹ AI ethics principles documents from over 65 governments and organisations contain some statement about transparency/explainability. Some mathematical methods, like decision trees, are more easily explainable than others, which may alter the type of machine learning an organisation needs to use.
- 3) **Size and diversity of dataset:** available training datasets are not always big enough to produce accurate results. Additionally, training data may not reflect the diversity of real world applications, subsequently affecting the accuracy of the machine learning model. For instance, facial recognition algorithms trained using mostly caucasian faces may have trouble accurately recognising darker skin tones when used in a live application.
- 4) **Impact of errors:** the type of mathematical method used may require a compromise between speed and accuracy. The impact of algorithmic errors can be minor, for example, a Netflix viewer recommended a film they don't enjoy. Other errors could have a major impact, for example, a prisoner may be denied parole. Organisations should consider the impact of errors when making choices regarding machine learning implementation.
- 5) **Privacy:** data can contain values that can identify specific individuals.¹² There are specific legal and ethical considerations when dealing with the collection and analysis of this kind of data, which is called "personal information" in New Zealand. This includes compliance with the Privacy Act 1993 (currently being updated), which governs the collection, use, storage and disclosure of personal information.


NEURAL NETWORKS AND DEEP LEARNING

Artificial neural networks are mathematical methods that are loosely inspired by the structure of biological neural networks. Computer scientists often refer to the layers of connected nodes as artificial neurons, as they approximate the function of neurons in a brain. Neural networks consist of input and output layers of computing units, which transform the inputs into something that the output layer can use.¹³ Neural networks often have multiple additional layers between the input and output layers: these are known as deep neural networks. In **deep neural networks**, (or simply “deep learning”), the middle layers learn not from features in the training examples, but from the outputs of the preceding layers.¹⁴

Deep learning is well suited to complex tasks like recognising speech or identifying objects in a photograph. Deep learning transfers the responsibility for feature extraction from the programmer to the program.¹⁵ For example, devising a complete set of classification guidelines to identify both a blurry close-up of a tabby's nose and a wide shot of a hairless cat's tail as 'cat'

would be impossible for a human programmer. With deep learning, each layer of processing can solve one part of the problem before passing it on to the next layer.¹⁶ The programmer no longer needs to define every element that makes a cat a cat. The deep learning program devises its own process in order to deliver correct results.

Recent excitement about AI is largely due to the possibilities of deep learning like the ability to understand speech, recognise objects, and make judgements about emotion. Already, deep learning programs have revealed vast new AI capabilities, enabling autonomous vehicles to avoid hazards, and helping Siri and Alexa to respond to spoken commands. Deep learning does have its limitations, primarily requiring huge data sets for training. It can also require massive processing power, and its multiple hidden layers mean that it can be hard to see or explain exactly how the program works. However, it's clear that this area of AI will continue to create vast new opportunities for innovation and optimisation.



MACHINE LEARNING FRAMEWORKS

The wide variety of open-source frameworks and libraries for developing AI make it increasingly easy and accessible to design custom AI solutions. Some of the most popular frameworks include Tensorflow, Microsoft CNTK, Caffe, Keras, Apache MXNet, Skymind, Deeplearning4j and PyTorch. Many of the libraries use Python, a popular and open source programming language. These libraries are contributed to by some of the world's largest technology companies and can be downloaded for free. Brief descriptions of some of the most popular open-source frameworks and libraries are provided below:



TensorFlow – TensorFlow is an end-to-end open source platform for machine learning, created by Google and used to design, build, and train deep learning models. It has a comprehensive, flexible ecosystem of tools, libraries and community resources that lets researchers push the state-of-the-art in machine learning (ML) and developers easily build and deploy ML powered applications.



Microsoft Cognitive Toolkit – Microsoft Cognitive Toolkit (CNTK) is an open-source toolkit for commercial-grade distributed deep learning. It describes neural networks as a series of computational steps via a directed graph. CNTK allows the user to easily realise and combine a variety of popular model types.




PyTorch – PyTorch an open-source deep learning framework, developed primarily by Facebook's AI research group and used for applications such as deep learning and natural language processing. PyTorch is known for providing tensor computations with strong GPU acceleration support and building deep neural networks on a tape-based autograd systems.



Apache MXNet – Apache MXNet is a modern open-source deep learning framework used to train, and deploy deep neural networks. It is scalable, allowing for fast model training, and supports a flexible programming model and multiple languages The MXNet library is portable and can scale to multiple GPUs and multiple machines.



Eclipse Deeplearning4j – DL4J is a commercial-grade, distributed deep learning library written for Java and Scala. Integrated with Hadoop and Apache Spark, DL4J brings AI to business environments for use on distributed GPUs and CPUs. Deeplearning4j was originally written by Skymind in 2014 and in 2017 it joined the Eclipse foundation for open source software.



Keras – Keras is an open-source neural-network library. Designed to enable fast experimentation with deep neural networks, it is capable of running on top of TensorFlow, Microsoft Cognitive Toolkit, Theano, or PlaidML. In addition to standard neural networks, Keras has support for convolutional and recurrent neural networks.



Caffe – Caffe (Convolutional Architecture for Fast Feature Embedding) is a deep learning framework, originally developed at the University of California, Berkeley. Caffe was mainly built to support convolutional neural networks and is a popular framework for those working on computer vision, image processing, and feed forward networks.

AI AS A SERVICE

AI solutions can either be built in-house using custom solutions or frameworks like those mentioned above, or deployed using ‘infrastructure as a service’ technologies provided by large AI companies. Key players in this space include Google Cloud, Microsoft Azure, Amazon AWS, and IBM Watson.

Subscription to these services allows organisations to upload and manage their data in the cloud, and deploy a range of ready-to-use AI solutions such as machine vision or natural language processing across the data. These services can provide popular open source deep learning frameworks, efficient AI development tools, and powerful servers.¹⁷ Many New Zealand organisations currently use these cloud-based solutions.

As noted above, many enterprise software companies also offer AI as a service solutions. Salesforce Einstein and SAP Leonardo offer AI services that come pre-integrated with the companies’ other services. In addition, many enterprise database servers are increasingly incorporating embedded machine learning – this includes Oracle, Microsoft SQL Server and MarkLogic.

Deloitte Insights notes that established vendors are not the only organisations supplying AI tools and

“The cloud is fundamental to the AI model in two ways. Firstly, the data sets these companies are using would not be accessible if it was not for the cloud. Secondly, only the cloud can enable businesses to cope with the phenomenal scale required by providing such data-intensive services to multiple clients at an affordable cost.”

STEVE SINGER (Talend).¹⁸



services, and it’s not yet clear what the dominant vendor model will be. There are currently a number of highly innovative startups developing diverse tools and solutions, and this proliferation will likely continue until we see consolidation of AI tools and infrastructure markets, and standards begin to emerge.¹⁹

While global cloud providers and international companies comprise a majority of the AI solution ecosystem, New Zealand also has some home-grown companies and organisations providing high performance computing and cloud data storage solutions.

New Zealand based Catalyst provides open source solutions and cloud data storage.²⁰ Having data hosted locally can reduce data transfer costs and help protect data sovereignty. Companies like Dragonfly Data Science use Catalyst’s cloud services for their Kōkako product which distinguishes spoken Māori from English.²¹

New Zealand also has some onshore hubs for high-performance computing. For example, NIWA has invested \$18 million in supercomputing technology,²² which helps ensure that researchers around the country can harness this computing power to benefit all New Zealanders.²³ The University of Canterbury also has a supercomputing facility²⁴ and both NIWA and UC are part of the New Zealand e-Science Infrastructure (NeSI) collaboration.²⁵

2. Recent AI Developments

“We’re at the end of the age of discovery and there’s likely to be an age of implementation.”

KAI-FU LEE, Chinese AI Investor²⁶

DURING THE LAST YEAR, AI ACTIVITY HAS CONTINUED TO ACCELERATE.

As AI technology continues to mature it has begun to generate returns on investment outside of global technology companies.

In parallel, an AI for Good movement to stimulate innovation in social causes has gained significant momentum with a number of proof points demonstrating uses for AI in humanitarian, health, environmental and social justice outcomes. There has also been a global crescendo of discussion on the many facets of AI ethics.

International AI Developments

TECHNOLOGICAL ADVANCES

Just a few years ago, the skills and experience required to design and develop machine learning algorithms were at PhD level or above. However, as discussed earlier, the increased availability of open source, ready to use machine learning frameworks and toolkits have significantly reduced the barrier to entry for developers. In addition, cloud-based AI developer tools are maturing at pace, which is further improving productivity and output for AI developers.

Major technology companies are now working on tools which will bring basic AI capabilities to a mass market, general developer audience without the need for years of training to use them.

“The latest AI ‘boom’ could be generally summarized as the big leap of functional applications thanks to the explosion in big data, computing power and continued advancing algorithms. Now comes a time when AI applications make a real impact on the economy,”²⁷ says Baidu’s Haifeng Wang.

EXAMPLES OF RECENT AI PLATFORM ADVANCES:

Cloud Vision²⁸ and Cloud Video Intelligence²⁹

Google Cloud offers two computer vision products that use machine learning to help understand and analyse images and videos with strong prediction accuracy.

Cloud AutoML is a suite of machine learning products that enables developers with limited expertise to train models specific to their business needs. Currently, AutoML supports translation, natural language, and vision use cases leveraging transfer learning and neural architecture search technology.

Azure Cognitive Services³⁰

Cloud hosted APIs for application developers to easily add AI capabilities such as speech recognition, voice synthesis, vision, language understanding, translation, knowledge and search.

Azure Machine Learning Service³¹

The Microsoft Azure Machine Learning Service now provides hosted Jupyter notebooks and allows developers to build, train and deploy machine learning models without having to touch any code.

AWS Textract³²


Is a cloud service which automatically lifts text and data from scanned documents without necessitating manual reviews, custom coding or experience with machine learning solutions. Textract contextualises the information it is reading based on its format and the fields presented.

Advances in reinforcement and transfer learning

Advances in new machine learning techniques including reinforcement learning and transfer learning have enabled machines to learn directly from their environment and take knowledge gained from one problem domain and apply it to a different but related domain. Recent examples include advances in gaming from DeepMind and OpenAI.



Alphabet subsidiary Deepmind made headlines using reinforcement learning to teach machines to play ever more complex games better than humans. Firstly, the deceptively complex board game Go (AlphaGo – 2015). Then, the single player strategy game, StarCraft II (2018). DeepMind produced an AI system (AlphaStar) that beat the world's best StarCraft II players – this was significant since the skills needed to win include game theory, long term planning with imperfect information and a vast action space. In 2019, Deepmind announced it has trained machine agents to successfully cooperate with both artificial and human teammates, resulting in human level performance in the canonical 3D first person multiplayer game, Quake III Arena Capture the Flag.³³



OpenAI Five Finals



OpenAI

Meanwhile, the team at OpenAI also hit the news when it released OpenAI Five. The machine learning project was trained to play the complex multiplayer game Dota 2, the first AI to beat the world champions in an eSports game after defeating the reigning Dota 2 world champions at the OpenAI Five Finals in April, 2019.³⁴


The Chips are Up

DURING THE PAST TWO YEARS, THE MARKET FOR DEEP LEARNING HARDWARE HAS EXPERIENCED A DRAMATIC PERIOD OF GROWTH AND EVOLUTION, LED BY MANUFACTURERS NVIDIA AND INTEL.

Google's Tensor Processing Unit (TPU) is an example of a processor specifically built for deep learning, designed to work with the TensorFlow system. For example, TPUs were used in the famous Go match between DeepMind AlphaGo and world champion Lee Sedol.

Industry analysts Tractica predict a ramp-up in deep learning chipset volumes in 2019-2020, when winners begin to emerge.³⁵ They forecast the market for deep learning chipsets will increase from US\$5.1 billion in 2018 to US\$72.6 billion in 2025. This is mostly made up of the edge computing market, where AI computation is done "at the edge" on a device as well as the enterprise data centre environments which are most common today.

Already high-end smart phones from Apple and Huawei incorporate dedicated machine-learning processors ("neural processors") that enable sophisticated models to be executed on the device, ie. at the edge of customer/user interaction.



Tensor Processing Unit 3rd Generation
PHOTO: Zinskauf, shared under CC BY-SA 4.0

In multiplayer gaming, the major AI research challenges of mastering strategy, tactical understanding and team play have seen major advances in recent years. While these developments show impressive advances in AI capabilities, they also highlight the significant limitations of machine learning systems.

QUANTUM AI ARRIVING SOON?

Looking even further forward, global researchers are now focused on using advances in quantum computing to accelerate AI. Quantum computing is the use of quantum mechanical phenomena, such as superposition and entanglement, to perform computation. Quantum computers can solve problems that are computationally too difficult for a classical computer. In the quantum universe, bytes can exist

in two states (qubits) at once, allowing computations to be performed in parallel. Therefore, two qubits are able to hold four values at the same time: 00, 01, 10, 11.

Quantum computers promise not only to be more powerful than anything built to date, they require special algorithms capable of doing new things. Labs such as Google AI³⁶ and Qualcomm³⁷ are developing novel quantum chips and algorithms to dramatically accelerate computational tasks used for machine learning.

FUTURE AI TRENDS

The Future Today Institute has published a comprehensive 2019 *Tech Trends Report*,³⁸ detailing 44 international trends in AI which provides a glimpse of more future innovations soon to arrive..

1	Consumer-Grade AI Applications	23	Real-Time Machine Learning
2	Ubiquitous Digital Assistants	24	Natural Language Understanding (NLU)
3	A Bigger Role For Ambient Interfaces	25	Machine Reading Comprehension (MRC)
4	Deep Linking Everywhere	26	Natural Language Generation (NLG)
5	Proliferation of Franken-algorithms	27	Generative Algorithms For Voice, Sound and Video
6	Deployable AI Versions of You	28	Real-Time Context in Machine Learning
7	Ongoing Bias in AI	29	General Reinforcement Learning Algorithm
8	AI Bias Leads To Societal Problems	30	Machine Image Completion
9	Making AI Explain Itself	31	Hybrid Human-Computer Vision Analysis
10	Accountability and Trust	32	Predictive Machine Vision
11	AI Hiding its Own Data	33	Much Faster Deep Learning
12	Undocumented AI Accidents on the Rise	34	Reinforcement Learning and Hierarchical RL
13	The AI Cloud	35	Continuous Learning
14	Serverless Computing	36	Multitask Learning
15	New Kinds of Liability Insurance for AI	37	Generative Adversarial Networks (GANs)
16	Generating Virtual Environments From Short Videos	38	New Generative Modeling Techniques
17	AI Spoofing	39	Capsule Networks
18	Ambient Surveillance	40	Probabilistic Programming Languages
19	Proprietary, Home Grown AI Languages	41	Automated Machine Learning (AutoML)
20	AI Chipsets	42	Customized Machine Learning
21	Marketplaces For AI Algorithms	43	AI For the Creative Process
22	Even More Consolidation in AI	44	Bots

SOURCE: Future Today Institute, shared under CC BY-NC-SA 4.0 International License.

VAST AI INVESTMENT AND CONCENTRATION OF AI SUPERPOWERS

Perhaps the most dominant AI related meme in 2018 was created by Chinese venture capitalist Kai-Fu Lee in his book *AI Superpowers: China, Silicon Valley and the New World Order*. Lee describes an unassailable duopoly between the USA and Chinese AI companies due to scale of AI investment in these two countries.

This is backed up by analysis in the West, for example, think tank The Future Today Institute states:

*“There are nine big tech companies – six American, and three Chinese – that are overwhelmingly responsible for the future of artificial intelligence. They are the G-MAFIA in the US: Google, Amazon, Microsoft, Apple, IBM and Facebook. In China it’s the BAT: Baidu, Alibaba and Tencent. Just nine companies are primarily responsible for the overwhelming majority of research, funding, government involvement and consumer-grade applications.”*³⁹

Lee also anticipates the bifurcation of the USA and Chinese AI ecosystems, with the growth of Chinese apps (WeChat, Alipay) becoming dominant in many emerging markets while the West consolidates onto existing US apps (Facebook, Google).⁴⁰




Cover of *AI Superpowers* by Kai Fu Lee

There are counterpoints to the AI Superpowers rhetoric, but the number and scale of AI investment deals in the USA and China certainly dwarfs that of the rest of the world, despite overall worldwide private equity investment growing an estimated 16 times between 2011-2017.



FIGURE 6: Total estimate equity investments in AI startups, by startup location 2011-17 and first semester 2018



SOURCE: OECD estimates, based on Crunchbase (July 2018).

The OECD reports that after five years of steady increases, private equity investment in AI began accelerating in 2016, with the amount of private equity invested actually doubling from 2016 to 2017. In total, it is estimated that more than US\$50 billion was invested in AI startups during the period 2011 through to mid 2018.⁴¹ This investment was dominated by the USA and China, with a dramatic upsurge in Chinese AI startup investment since 2016. From just three percent in 2015, Chinese companies attracted 36 percent of global AI private equity investment in 2017. This growth partly reflects the Chinese Government’s investment strategy to be a world leader in the field of AI.

According to research firm CBInsights, a total of 17 AI startups reached US\$1 billion+ valuations in 2018,⁴² up from nine the previous year. Almost all of these “Unicorn” firms are based in the USA or China.⁴³

More recent commentary has speculated on the probability that there are increasing signs of an AI investment bubble happening.⁴⁴ Furthermore, related to increasing US trade disputes with China, Chinese investments in US startups have been on the decline since peaking in 2016. Furthermore,

“I think the challenge from an NZ Inc. perspective is the amount of investment we are seeing offshore and the amount of investment we are seeing onshore is bifurcating – [other] countries are accelerating away in terms of the amount of investment and actually we’ve seen a fall in investment levels... so I would suggest a significant amount more money needs to be invested in the early stage technology space in order for us to just catch up on where other countries are going.”

RICHARD DELLABARCA – CEO of NZVIF⁴⁵

the US Department of Commerce has started work on drawing up controls on the export of sensitive AI technologies.^{45a} Foreign direct investment from China dropped to US\$5 billion in 2018 from US\$29 billion in 2017 and US\$46 billion the prior year.⁴⁶

For perspective, the US\$16.5 billion of global AI investment in 2017 is compared with US\$0.58 billion (NZ\$876 million) of total private equity and venture capital investment in New Zealand technology companies that year,⁴⁷ a tiny fraction of which were AI businesses.

AI INVESTMENT BY GOVERNMENTS

In addition to private sector AI investment, governments around the world are also digging deeply into their pockets. The Pentagon is expecting to spend US\$2 billion on next generation AI and the UK Government has pledged £1 billion to invest in similar projects.

At least 20 countries, including the USA and China have announced national AI strategies backed by significant public investments. A selection of these are summarised in the table below.

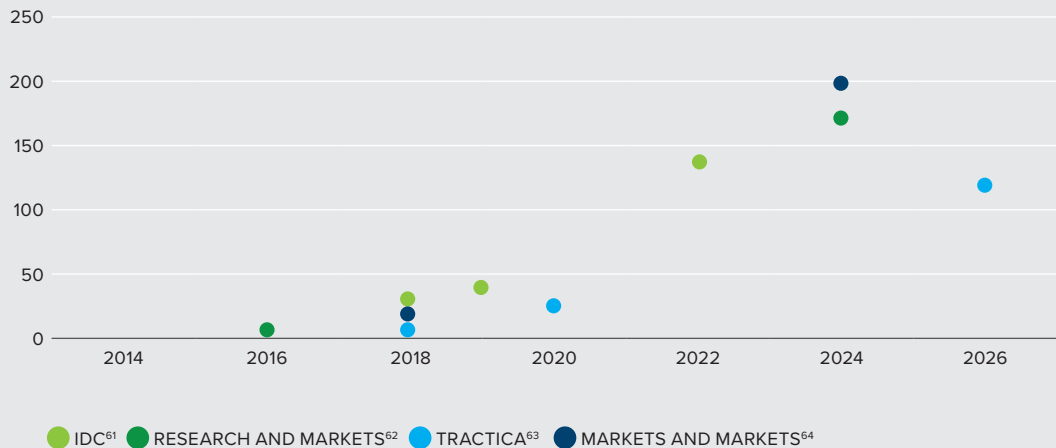
Given the imbalance between AI investment in US and China and the rest of the world, effort will be needed to conceive a global governance model for AI which ensures that the benefits and threats are better balanced internationally.

It is worth noting that these investments are for the initial AI strategies themselves, and several have received further ongoing funding and investment. For example, the Canadian Federal Government has also invested nearly CA\$1bn in superclusters heavily focused on AI – with 1 being purely focused on AI receiving CA\$270m. The Ontario Government has also invested around CA\$180m in a regional AI strategy.

TABLE 2: Announced investments in ten national/regional AI strategies

COUNTRY	DATE	INVESTMENT (\$USD)	POPULATION	INVESTMENT (\$USD (TOTAL)/CAPITA)	BASIS	NZ\$ PER CAPITA PER YEAR (USD TO NZD 1.50)
USA ⁴⁸	2019	2 billion+	329m	7.3	Defense \$2b over 5 years, Govt R&D \$2 billion in 2017 (assumed same figure moving forward) ⁴⁹	\$10.94
China ⁵⁰	2017	N/A	1420m			
Denmark ⁵¹	2018	114m (DKK 750m)	6m	19	PA 2019-2025	\$4.75
Singapore ⁵²	2017	111m (\$150m)	6m	18.5	Over five years	\$5.55
Taiwan ⁵³	2018	1.3bn (NT\$10 billion/year)	24m	56	Over four years	\$21.14
Australia ⁵⁴	2018	22m (AU\$29.9M)	25m	1	Over four years	\$0.33
Canada ⁵⁵	2017	95m (CA\$125)	37m	2.6	Five years	\$0.77
South Korea ⁵⁶	2018	1.9b (KRW 2.2 trillion)	51m	37	Five years	\$11.17
United Kingdom ⁵⁷	2018 (updated 2019)	1.65b (£1.3b)	67m	24.6	Not stated	\$7.40 (Assume five years)
France ⁵⁸	2018	1.75b (€1.5 b)	66m	26.5	Over five years	\$7.95
Germany ⁵⁹	2018	3.93b (€3b)	83m	47.3	Over 6 years 2019-2025	\$11.83
EU ⁶⁰	2018	1.75b (€1.5 b)	513m	3.4	Over 3 years 2018-2020	\$1.71

FIGURE 7: AI Market estimates (Billion US Dollars)



EXPLOSIVE MARKET GROWTH

There is a reason for the substantial investment noted above – markets for AI products and services are projected to continue to grow explosively in coming years. A number of market size estimates are shown in Figure 7, showing predicted exponential expansion in global spending on AI systems.

There is significant variation between these predictions as McKinsey Global Institute identifies (see quote).

Nonetheless, it is increasingly apparent that the global market opportunity for AI products and services is substantial, growing and available to those positioning and investing now.

“Analysts remain divided as to the potential of AI: some have formed a rosy consensus about AI’s potential while others remain cautious about its true economic benefit. This lack of agreement is visible in the large variance of current market forecasts, which range from \$644 million to \$126 billion by 2025. Given the size of investment being poured into AI, the low estimate would indicate that we are witnessing another phase in a boom-and bust cycle.”

MCKINSEY GLOBAL INSTITUTE⁶⁵

AI FOR GOOD

The application of AI technologies for social, environmental and humanitarian outcomes rather than pure financial profit has gained traction in the last year. Major technology firms Amazon, Google, Microsoft and SAP have all established grant programmes to support initiatives with some variation of AI for Good.

Recent examples of how AI for Good is being explored include:

- studying the health of native bird populations⁶⁶
- quicker diagnosis of sick cows on dairy farms⁶⁷
- improving the quality of bail decisions in the criminal justice system⁶⁸
- detecting more counterfeit goods crossing borders⁶⁹
- improved cancer detection⁷⁰

- increased accuracy and confidence of grading diabetic eye disease⁷¹
- cleaning up the world's beaches⁷²
- increased accuracy of flood forecasting⁷³

The 2019 CogX Festival of AI and Emerging Technology in London was presented in partnership with 2030Vision. This is a new initiative aiming to transform the use of technology through collaborative partnerships and innovative projects in order to support the delivery of the United Nations Sustainable Development Goals (SDGs – 17 Global Goals – see box) and unlock the commercial opportunities they offer.

This key message that AI is a fundamental new tool to help bring about solutions to the world's major development challenges. The SDGs provide a more tangible framework to aim for than the more abstract concept of "AI For Good".



2030 AGENDA FOR SUSTAINABLE DEVELOPMENT

In September 2015, more than 190 countries signed up to the 17 Sustainable Development Goals of the 2030 Agenda for Sustainable Development. The Agenda sets out a vision to end poverty, protect the planet and ensure prosperity for all. Delivering the Global Goals by 2030 will require collaboration across the private sector, governments and civil society, massive investment and innovation, as well as fresh thinking about economic paradigms and business models. Achieving the Global Goals is not only a moral imperative, **but the task also presents a significant commercial opportunity, estimated at \$12 trillion a year in revenue and cost savings and 380 million new jobs by 2030.**⁷⁴




Google: AI For Social Good

GOOGLE RUNS A PROGRAM FOR SOCIAL GOOD THAT FOCUSES THEIR AI EXPERTISE IN EFFORTS TO HELP SOLVE HUMANITARIAN AND ENVIRONMENTAL CHALLENGES.⁷⁵

By applying core Google research and engineering resources in partnership with experts, Google is using its AI technology to benefit the world.

Following a global open call, twenty organisations are receiving customized support to implement their ideas for social good. For example, Google was a founding partner of Global Fishing Watch. This service brings big satellite data, the latest machine learning, and cloud computing technology to bear on human interactions with our oceans' natural resources. The resulting near real-time map shows where vessels are fishing, even at night. This allows better monitoring of our global fish reserves.



Screenshot from the near real-time Global Fishing Watch Map⁷⁶

Microsoft: AI For Earth

THROUGH A PROGRAMME CALLED 'AI FOR EARTH', MICROSOFT IS MAKING ITS CLOUD AND AI TOOLS AVAILABLE TO THOSE WORKING TO PROTECT THE ENVIRONMENT.⁷⁷

Microsoft helps to build a sustainable future by providing grants to projects that use artificial intelligence to address the critical areas of:

Agriculture

Biodiversity

Climate change

Water

Organisations can apply for grants to help them label datasets, which are then made public and used to train AI models with sustainability goals in mind. Alternatively grants are awarded to those whose data is ready to analyse. In the latter case, grants take the form of Microsoft Azure compute credits.

For example, Microsoft has partnered with Chesapeake Conservancy to generate land cover data with one square meter resolution to enable precision conservation. Also, Microsoft support helps machine learning models to scour millions of images and thousands of hours of field recordings to detect species of interest to conservationists.



SOLVING CLIMATE CHANGE WITH AI

The website www.climatechange.ai was launched in June 2019 following the publication of the paper *Tackling Climate Change with Machine Learning*⁷⁸ co-authored by some of the world's leading international academics and tech industry researchers. The paper shines a light on how machine learning techniques can be used to reduce greenhouse gas emissions and help society adapt to a changing climate. The figure below illustrates the many potential problem / solution mappings.

TABLE 3: Climate change solution domains, along with areas of machine learning that are relevant to each. This table should not be seen as comprehensive.

	COMPUTER VISION	NLP	TIME-SERIES ANALYSIS	UNSUPERVISED LEARNING	RL & CONTROL	CAUSAL INFERENCE	UNCERTAINTY QUANTIFICATION	TRANSFER LEARNING	INTERPRETABLE ML	OTHER
Electricity Systems	•	•	•	•	•		•	•	•	•
Transportation	•		•	•	•	•	•	•	•	
Buildings and Cities	•	•	•	•	•	•	•	•		
Industry	•		•	•	•	•		•	•	
Farms and Forests	•				•			•		
CO ₂ Removal			•				•	•		•
Climate Prediction	•		•				•		•	
Societal Impacts	•	•	•		•	•	•	•	•	
Solar Geoengineering			•		•		•			•
Tools for Individuals	•	•	•	•	•	•			•	•
Tools for Society		•	•	•	•	•	•	•	•	•
Education		•			•					
Finance		•	•				•			

SOURCE: climatechange.ai

MOVEMENT TOWARDS REGULATION OF AI

In light of the rapid advances in AI, and the exceptional profitability of global scale AI driven online retail, advertising and social media businesses, there have been international calls for regulation of both the technology and the business models built on top of it.⁷⁹ The 2018 Cambridge Analytica scandal where the political firm gained access to private information from more than 50 million Facebook users has stoked demand for more stringent regulation.

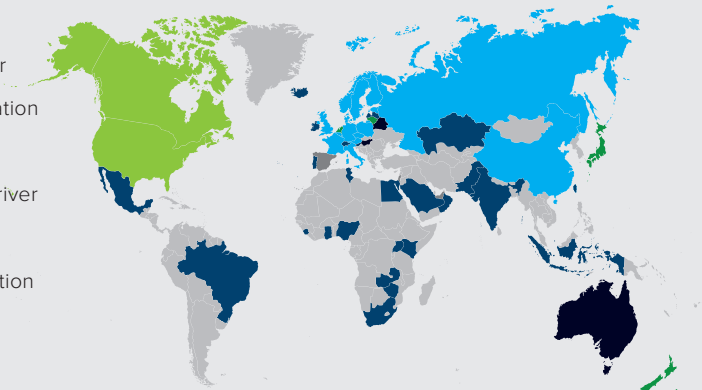
However, there is a broad public perception that laws cannot keep up with advances in technology⁸⁰ and sometimes regulatory gaps persist for fear of impeding economic growth and innovation.⁸¹ It has been argued that AI is still in its preliminary stages and strict regulation is neither necessary nor desirable. The speed of technological change and historically low efficiency of regulations and perceived rent seeking by governments are all used as arguments against regulation.⁸²

Furthermore, regulation of government and public-sector uses of AI raises different issues from regulation of private sector uses of AI, even though the AI technologies used in these two areas are essentially the same. There is increased public accountability for public sector uses, which create different standards, especially in relation to transparency.



FIGURE 8: Regulation of Autonomous Vehicles Worldwide

- Allowed without human driver
- Allowed only with human driver
- Allowed but not specific legislation
- Countries not in study
- Allowed with/without human driver
- Varies by state/province
- Not allowed/no specific legislation



SOURCE: Law Library of Congress Analysis.

Currently there is negligible AI specific regulation worldwide.

The US Law Library of Congress has published a review titled, 'Regulation of Artificial Intelligence in Selected Jurisdictions,' which concludes that the regulation of AI is still in its infancy with the most advanced regulations being in the area of autonomous vehicles.⁸³ Many countries have established commissions to look into regulation but only the EU has published a framework for designing trustworthy AI. The Netherlands and Lithuania have passed legislation enabling the testing of driverless vehicles. With respect to lethal autonomous weapons, the majority of countries agree that meaningful human control is necessary. In May 2019, the City of San Francisco, USA voted to ban facial recognition technology by Police and other agencies,⁸⁴ since joined by the cities of Somerville MA and Oakland CA.⁸⁵

The chairman of Alphabet, Google's parent company, John Hennessy, advocates for technology firms to self-regulate, claiming the industry moves too quickly for governments to be effective.

“If you don't self-police, if you don't preserve the trust of your user community, then you will get regulated.”

JOHN HENNESSY, Chairman, Alphabet⁸⁶

Calls to Regulate AI

Who is responsible if AI systems harm us? In its 2018 report, the USA think tank The AI Now Institute identified five issues arising from the fundamental question of accountability. These include:

1. The growing accountability gap in AI, which favours those who create and deploy these technologies at the expense of those most affected.
2. The use of AI to maximize and amplify surveillance, especially in conjunction with facial and affect recognition, increasing the potential for centralized control and oppression.

3. Increasing government use of automated decision systems that directly impact individuals and communities without established accountability structures.
4. Unregulated and unmonitored forms of AI experimentation on human populations.
5. The limits of technological solutions to problems of fairness, bias, and discrimination.

The report recommends sector-specific approaches that do not regulate AI technology as a whole but focus on its application within a given domain:

AI is fundamentally a set of technologies and like all technology, it can have both helpful and harmful applications. Regulating the technology itself rather than its applications would potentially stifle positive advances and innovation enabled by AI.

“Governments need to regulate AI by expanding the powers of sector-specific agencies to oversee, audit, and monitor these technologies by domain... Domains like health, education, criminal justice, and welfare all have their own histories, regulatory frameworks, and hazards. However, a national AI safety body or general AI standards and certification model will struggle to meet the sectoral expertise requirement needed for nuanced regulation.”

AI NOW INSTITUTE⁸⁷

CASE STUDY CUSTOMER SERVICE

FACEME: Digital Humans

FaceMe is a digital human platform that uses AI to create a natural human-like interface for machine to human communication. While a chatbot enables conversation, a digital human enables a natural, individual, visual connection. Digital humans can be either customer facing or assist employees.

FaceMe's origins were with real-time video into contact centres and using approved customer data to train an EQ system that today drives real-time interactions with their Digital Humans. It has since realised an opportunity to create personalised experiences and a brand-aligned presence within digital contact channels. Its investment in a global network for real time media enables the AI to process and respond to hundreds of thousands of interactions at the same time. FaceMe's solution can operate in conjunction with any chatbot solution and can inhabit a kiosk, or be an online presence. The digital humans translate chatbot output into realistic conversation, with emotion and animation. The inputs, such as what it sees, hears and knows are used to determine how to appropriately respond. For example, if an end user appears confused by a response the digital human may offer to clarify its answer. FaceMe's customers and digital humans include, Vodafone (Kiri), UBank (Mia), Ministry of Primary Industries (Vai) and ASB Bank (Josie), UBS and Southern Cross.

It is expected that digital humans will increasingly become a key part of everyday customer experiences. The emotional connection is becoming more important as is the power of embodied brand in a digital world. Bringing 'humanness' into digital experiences can result in increased sales conversions and higher customer advocacy.

However, the digital humans market is in the early adopter phase. Market education and gaining traction with customers can be a challenge for service providers. FaceMe currently operates in Australasia and recently expanded to the USA and Europe. It is investing in research and development in New Zealand and expects the development team to double in size over the next 18 months.

Organisations who want to explore the option of digital humans should have a clear digital



experience strategy and roadmap. FaceMe CEO Danny Tomsett says that although customers risk 'biting off more than [they] can chew', a bigger risk is 'doing nothing'. FaceMe's guidance for deploying a digital human includes:

- Understand what topics end users are comfortable speaking to an AI about.
- Use an iterative process so organisations can develop a deep understanding of their customer's problems.
- Companies should consider how they can capture, analyse and act upon the insights gained from deploying digital humans.

FaceMe.

Learn more at www.faceme.com

CASE STUDY BANKING AND FINANCE

ANZ: Assisting with General Banking Enquiries



In July 2018, ANZ New Zealand launched their new digital assistant pilot, called Jamie.

The AI technology was developed with New Zealand company Soul Machines. Jamie was initially programmed to answer questions based on the 30 most frequently searched for online topics. The pilot has since been extended and is increasing Jamie's workload as it learns to answer a broader range of customer enquiries and increase its Te Reo Māori vocabulary.

In its first 100 days Jamie had more than 12,000 conversations with people visiting the site. The most common question was how to open a bank account, which Jamie was asked nearly 1200 times. Liz Maguire, ANZ's Head of Digital and Transformation, reports that Jamie was able to answer approximately 60 percent of customer queries.⁸⁸

To avoid robotic responses to standard banking questions, considerable effort was made to optimise Jamie's friendly persona. This included

spending time observing the ANZ contact centre team. In addition, Jamie's creators developed a backstory and personality.

Currently, Jamie is only able to assist with general banking queries. However, it is expected that Jamie will soon be able to carry out personal banking tasks.

"When you're driving in the car you might go: 'Hey, Jamie, I really need to pay the babysitter \$50.' And Jamie would do that for you," says Maguire.⁸⁹



Learn more at www.anz.co.nz

CRESCENDO OF AI ETHICS

The field of AI ethics has grown rapidly in recent years. Perhaps more than any other emergent technology in history, voices from academia, civil society and the technology sector itself have engaged with the concept of artificial intelligence and its implications for the ethical values which underpin much of our civil society. AI creates the potential for profound social justice implications if its use results in divergent access, systemic impacts, or increases in discrimination and inequities.

A number of the ethical issues around AI are related to the concepts of data rights, data sovereignty and privacy issues, which we discuss in detail in Chapter 5 below. However, other key ethical questions posed by AI include issues around: who is accountable for the actions of autonomous systems? How do we achieve transparency and explanation when many of these technologies are 'black boxes' and opaque even to their creators? How can we ensure fairness, given the potential for bias and injustices that are common in datasets? How can we ensure diversity and inclusion? And how can human dignity, autonomy, and rights be protected in practice?

These issues are by their nature global – and as a result the AI Ethics debate has highlighted the variability in ethical values worldwide. The ethical correctness of deploying AI-enabled surveillance technologies varies from country to country – see for example the differences between China's approach to facial recognition and the tracking of its citizens⁹⁰ and that of European nations where the GDPR has come into effect. San Francisco has banned the use of facial recognition software by police,⁹¹ whereas Chinese AI firm Cloudwalk is cooperating with the Zimbabwean Government on a mass facial recognition project.⁹²

Recent work from the Berkman Klein Center for Internet and Society at Harvard University aims to map consensus and divergence in ethical and rights-based approaches to AI principles. The research catalogued more than 32 sets of AI Principles. Berkman Klein's resulting visualisation further highlights the increasing level of AI Ethics activity.⁹³

Dave Heiner, Strategic Policy Advisor at Microsoft said in his talk at AI DAY 2019:⁹⁴

“Where do we start? What are the highest level values that we are trying to uphold? Ethics gets challenging right away, because values differ around the world... but there are some values that are universal. Recently (December 2018) the world celebrated the 70th anniversary of the UN Declaration of Human Rights. These are the closest thing to universal values we have come up with which have stood the test of time.”



Table 4 displays the declarations that embrace the following categories of AI Principles to varying degrees: **Human Rights, Promotion of Human Values, Professional Responsibility, Human Control of Technology, Fairness and Non-discrimination, Transparency and Explainability, Safety and Security, Accountability, Privacy.**

TABLE 4: Declarations of AI Principles	
CIVIL SOCIETY	MULTISTAKEHOLDER
<ul style="list-style-type: none"> • Top 10 Principles for Ethical AI (UNI Global Union) • Toronto Declaration (Amnesty International, Access Now) • Future of Work and Education for the Digital Age (T20: Think 20) • Universal Guidelines for AI (The Public Voice Coalition) • Human Rights in the Age of AI (Access Now) 	<ul style="list-style-type: none"> • Tenets (Partnership on AI) • Asilomar AI Principles (Future of Life Institute) • The GNI Principles (Global Network Initiative) • Montreal Declaration (University of Montreal) • Ethically Aligned Design (IEEE) • Seeking Ground Rules for AI (New York Times)
GOVERNMENT	INTER-GOVERNMENTAL ORGANISATION
<ul style="list-style-type: none"> • Preparing for the Future of AI (US National Science and Technology Council) • Draft AI R&D Guidelines (Japan) • White Paper on AI Standardization (Standards Administration of China) • Statement on AI, Robotics and ‘Autonomous’ Systems (European Group on Ethics in Science and New Technologies) • For a Meaningful Artificial Intelligence (Mission assigned by the French Prime Minister) • AI at the Service of Citizens (Agency for Digital Italy) • AI for Europe (European Commission) • AI in the UK (UK House of Lords) • AI in Mexico (British Embassy in Mexico City) • Artificial Intelligence Strategy (German Federal Ministries of Education, Economic Affairs and Labour and Social Affairs) • Draft Ethics Guidelines for Trustworthy AI (European High-level Expert Group on AI) • AI Principles and Ethics (Smart Dubai) • Principles to promote FEAT AI in the Financial Sector (Monetary Authority of Singapore) 	<ul style="list-style-type: none"> • European Ethical Charter on the Use of AI in Judicial Systems (Council of Europe CEPEJ)
	PRIVATE SECTOR
	<ul style="list-style-type: none"> • AI Policy Principles (ITI) • The Ethics of Code (SAGE) • Microsoft AI Principles (Microsoft) • AI at Google: Our Principles (Google) • AI Principles at Telefonica (Telefonica) • Guiding Principles on Trusted AI Ethics (Telia Company) • Declaration of the Ethical Principles for AI (IA Latam)

The Berkman Klein visualisation available online exposes a perceived Western 'bubble' within which these AI ethics developments have been taking place. Work on AI ethics happening in China often refers to collective behaviours and values, as well as notions of social harmony, influenced by the diverse philosophical and cultural traditions in China, including blends of Confucianism, Daoism and Buddhism. Most recently the Chinese Government issued a series of principles to regulate the research and application of AI, in an attempt to ensure the 'safe, controllable and responsible use' of the technology amid rising privacy concerns in the country.⁹⁵

Ethical Oversight

ETHICS BOARDS/COMMITTEES

Ethics boards and committees are important in many industries for monitoring and encouraging best practice. Some businesses have attempted to set up ethics boards to oversee AI developments and monitor the appropriateness of innovations. However, appointments to Google's Advanced Technology External Advisory Council (ATEAC) resulted in a petition signed by thousands of Google employees and it was cancelled in its first week.⁹⁶

A number of companies and organisations have trialled an internal approach to ethical oversight without success. These entities can suffer from having the wrong personnel, not being impartial, supporting non-ethical issues or possessing the wrong skill sets. It may be that ethical oversight of AI development needs to come from outside the industry, accompanied by an understanding of government, research and business, plus an understanding of ethical theory and the ability to investigate, be impartial and able to act autonomously with authority. This is not an easy task but external ethics committees have a precedent in industries such as health research and clinical care. In New Zealand, the Health and Disability Ethics Committees reviews a large number of proposals, helping researchers and clinicians to ensure their work is developed ethically.

Clearly, self-regulation does not always work and this is why joint declarations of principles such as those in Table 6 have been developed.^{96a} The United Nations (UN) Special Rapporteur's report on AI and human rights also emphasises that it's all very well for companies to have ethics policies, but they must respect human rights too.⁹⁷

As an example of ethical assessment in action, the Canadian Government has issued a Directive on Automated Decision-Making. This manual includes an Algorithmic Impact Assessment (AIA) tool to determine exactly what kind of human intervention, peer review, monitoring, and contingency planning a given AI tool will require. Companies building AI solutions for government can access the AIA online and quantify the risks to determine exactly what kind of intervention they may need.⁹⁸

There is at least one international example of a commercial audit service which examines algorithms for accuracy, bias and fairness, the "ORCAA Seal".⁹⁹



The Australian Human Rights Commission published *Artificial Intelligence: Governance and Leadership* in January 2019, highlighting a number of ethical concerns relating to AI. Following this, the Australian Department of Industry, Innovation and Science released the report *Artificial Intelligence: Australia's Ethics Framework* which examined key ethical issues, case studies and a framework of core principles for AI development. Also in Australia, the Gradient Institute is a partnership between

IAG, CSIRO's Data61 and University of Sydney launched in December 2018. The Gradient Institute has the utopian goal of creating a "world where all systems behave ethically" through research, policy advocacy, public awareness, and training people in the ethical development and use of AI.

Looking Outward: Partnership on AI



The AI Forum of New Zealand is a member of the international Partnership on AI (PAI) which was established with support from a number of major industry players including Amazon, Facebook, Google, DeepMind, Microsoft and IBM. PAI now consists of over 80 partners worldwide. The key aims of PAI are to shape best practices, research, and public dialogue about AI's benefits for people and society.


This is achieved through a work programme based on six thematic pillars:

1. Safety critical AI
2. Fair, transparent and accountable AI
3. AI, labour and the economy
4. Collaborations between people and AI systems
5. Social and societal influences of AI
6. AI for social good.¹⁰⁰



Partnership on AI All Partners Meeting, San Fransisco, November 2018.

FIGURE 9: The Partnership on AI illustrates the effectiveness of a diverse partnership with commercial and non-commercial members.



European Union (EU) Guidelines for Trustworthy AI and OECD Recommendation on AI

The European Union (EU) and the OECD are moving towards encoding international ethical guidelines and standards of AI, which may in future form the basis of regulation.

In April 2019, *Ethics Guidelines for Trustworthy AI*¹⁰¹ was published by the European Commission. It highlights the need for AI to be lawful, ethical and robust. The guidelines offer technical and non-technical methods for trustworthy AI development and deployment, and an assessment list for AI practitioners to adopt. These will be tested as part of a pilot process to gather practical feedback on how the checklist, which operationalises the key requirements, can be improved. This may then lead towards formal regulation.

In May 2019, the 36 member Organisation for Economic Cooperation and Development (OECD) formally adopted the Recommendation on

Artificial Intelligence¹⁰² which includes the first international standards agreed by governments for the responsible stewardship of trustworthy AI.

The Recommendation sets out five principles for responsible stewardship of trustworthy AI:

1. sustainable development and wellbeing
2. human-centred values and fairness
3. transparency and explainability
4. robustness, security and safety and
5. accountability.

The principles explicitly reference the Sustainable Development Goals set out in the 2030 Agenda for Sustainable Development adopted by the United Nations as well as the 1948 Universal Declaration of Human Rights as their basis.

Most recently, on 9 June at the G20 summit in Japan, the G20 Ministers (including New Zealand) agreed on AI principles based on those adopted by the OECD.¹⁰³

New Zealand AI Developments

In the year since the AI Forum of New Zealand published its first research report, *Shaping a Future New Zealand*, there has been plenty of local activity.

While there are key areas of activity in Government and industry, it is important to note there are also significant areas of inactivity.

POLICY DEVELOPMENTS

Recent policy developments include a report on the automation of work in New Zealand by the Prime Minister's Business Advisory Council advocating increased training. The Productivity Commission has published an issues paper, *Technology and the Future of Work*¹⁰⁴, outlining four plausible scenarios for the future of work and called for submissions on 33 relevant questions. The new *Government Data Strategy and Roadmap* was published by the Government Chief Data Steward (GCDS) and calls for expressions of interest for members on a GCDS Advisory Group on Trusted Data Use.¹⁰⁵ There is also a new policy research centre, the Centre for AI and Public Policy at the University of Otago.

A Data Ethics Advisory Group has been formed by the GCDS and Chief Executive of Statistics New Zealand. The group is independent and tasked with helping government agencies use data more effectively, while ensuring that New Zealanders can have trust and confidence in the way their data is collected and used. The group will provide expert impartial advice on trends, issues, areas of concern, and areas for innovation.

The Government is also currently working on a Digital Government Strategy which will replace the 2015 Government ICT strategy.¹⁰⁶ At present development of the strategy is at stakeholder engagement stage. Work is being undertaken for a Cabinet Paper on the Strategy during the second half of 2019. Stakeholders can expect more information to be published on the Strategy by the end of 2019 or early 2020.

One key development is the 2018 publication of the *Algorithm Assessment Report*, an audit of New Zealand Government algorithm use by the GCDS and GCDO. This report conducted an assessment against the 'Principles for the safe and effective use

of data and analytics' and identified computer algorithms used for decision making in 14 Government departments. It concluded that humans, rather than computers, are ultimately making the significant decisions in Government agencies. However, as new algorithms are developed it will be important to maintain fairness, transparency and human oversight. This will further ensure there is appropriate consideration of the people who receive or participate in Government services. It will be essential for agencies to collaborate and share good practice.¹⁰⁷ None of the report's recommendations related to the need for new legal protections. The government is currently considering a response to the recommendations made by the report.

The Social Investment Agency (SIA), recently conducted 83 hui themed, 'your voice, your data, your say'. The SIA aims to analyse data, to better understand investment opportunities for the social wellbeing of New Zealanders. The SIA heard that New Zealanders want to know how information is used and what insights are gained. Service providers would also like to be able to access data for research, to inform service planning. Meanwhile, service users want to know how their personal information is used and who it is shared with.¹⁰⁸

A Māori Data Futures Hui was held in March 2018, which reached consensus that a tirohanga Māori (Māori perspective) is the most appropriate lens to view Te Ao – the World, so that any information collected from Te Ao Māori can truly reflect and represent Māori.¹⁰⁹

The Government has set up a Digital Economy and Digital Inclusion Ministerial Advisory Group with the purpose of providing government with advice and diverse perspectives on the digital economy and digital inclusion.¹¹⁰ As this report was going to press, the outgoing Minister of



Economic Development launched the Government's new approach to Industry Policy aimed at growing more innovative industries in New Zealand and lifting the productivity of key sectors, in particular: Agritech, Digital Tech, Forestry and Food.¹¹¹ AI is likely to be a key enabler for this policy.

Given the need for a skilled future workforce, the Digital Skills Forum held the inaugural New Zealand Digital Skills Hui in March 2019. The hui provided an opportunity for 300 delegates from industry, Government and NGOs to come together for a powerful, action focused day to shape actionable suggestions for policy, work programmes, education and employers.¹¹²

Also in 2019, MBIE has called for research programme proposals for \$49 million of funding from the Strategic Science Investment Fund for data science research to build New Zealand's existing data science capability.¹¹³


Important policy research includes the report *Perception Inception – Preparing for deepfakes and the synthetic media of tomorrow*¹¹⁴ by Brainbox Ltd. Also, the University of Otago's Artificial Intelligence and Law in New Zealand Project (AILNZP) has published *Government Use of AI In New Zealand*.¹¹⁴ Both these research projects were supported by The New Zealand Law Foundation.

In May, the political leaders of New Zealand, France, Canada, the UK and other countries from around the world joined together with leaders from international technology companies and civil society organisations to adopt the Christchurch Call. This was exactly two months following the terrorist attack on Christchurch mosques on 15 March in which 51 people lost their lives.



SYNTHETIC MEDIA

Perception Inception is a research project funded by the New Zealand Law Foundation's Information Law and Policy Project, and affiliated with the University of Otago.



The project examines ethical, legal and social questions arising from state-of-the-art emerging technologies. These technologies are capable of producing audiovisual information in ways we have not seen before, and of kinds we are unfamiliar with.

The project's report¹¹⁵ finds that new kinds of audiovisual technologies can produce increasingly realistic images, sounds and videos by creating and manipulating digital data using computers. These 'synthetic media' are problematic because they can make it look and sound like something happened that did not. There is potential to misrepresent and manipulate what people have said and done.

The report surveys New Zealand legislation and legal precedent around privacy, data protection, freedom of expression, interpersonal harms and copyright in the context of a world with synthetic media.

The report concludes that, 'the law in New Zealand is broadly well-equipped to deal with the impact of technologies which make it look like something happened when it didn't happen,' while noting some specific gaps in the law.

The report recommends caution in developing any substantial new law, but also recommends that we bring creation of synthetic media that represents identifiable individuals under the umbrella of privacy, and the 'reasonable expectation' against creation of such synthetic media. There is also a recommendation for indicators (labels or context) that show that content is synthetic.

The Christchurch Call is a commitment by governments and tech companies to eliminate terrorist and violent extremist content online. While it is not the central concern of the Christchurch Call, there are many opportunities for the application of new AI technologies to help curb the dissemination of extremist content online. For example, AI is being used to train systems to automatically flag concerning content for people to review.

ECOSYSTEM ADVANCES

Beyond Government, New Zealand’s AI business ecosystem has continued to expand in the last 12 months with the AI Forum’s New Zealand ecosystem map now numbering over 194 organisations.¹¹⁶

In 2018, New Zealand’s top 200 tech companies brought in over \$11 billion dollars in revenue, including \$8 billion dollars in exports, with highest ever profitability and growth, and employing more than 47,000 people in tech sector jobs. Some analysts believe that the tech export sector has the potential to become our country’s leading source of offshore income¹¹⁷ and a growing proportion of this income will be attributed directly to investments in AI.

In industry, we have seen some major AI events such as the Microsoft Future Now Conference in Auckland (25 October 2018), which focused on how to harness AI as a catalyst for growth.

Five hundred people attended the AI Forum’s annual AI-DAY 2019¹¹⁸ two day March conference in Auckland to discuss how AI affects society and business. There were 33 local and international speakers, with one highlight being a closing panel discussion on using AI to make safer societies in the wake of the Christchurch mosque shootings.

Throughout the year, the AI Forum has partnered with other organisations such as FinTechNZ and NZTech to host a range of other AI related events.

The AI Forum adopted the themes of last year’s research report, and launched six working groups to foster collaboration across the ecosystem and to address some key challenges including:

- Forging a coordinated AI strategy for New Zealand
- Creating awareness and discussion of AI
- Assisting AI adoption
- Increasing trusted data accessibility
- Growing the AI talent pool
- Adapting to AI effects on law, ethics and society.

As well as inaugurating the working groups, the AI Forum of New Zealand has signed the international Lethal Autonomous Weapons Pledge (LAWS) to ban autonomous killing machines.¹¹⁹ The pledge calls for the decision to take a human life to never be delegated to a machine. It has been signed by 247 organisations and 3253 individuals worldwide.



AI DAY 2019 closing



AI Forum Working Groups

TABLE 5: AI Forum working groups agenda 2018-2019.**THEME 1*****Forging a Coordinated AI Strategy for New Zealand***

Recommendations:

- 1.1 Develop a coordinated national AI strategy as part of New Zealand's wider Digital Strategy.
- 1.2 Ensure AI features strongly in the national cybersecurity strategy.

THEME 2***Creating Awareness and Discussion of AI***

Recommendations:

- 2.1 Advance AI awareness and understanding.

THEME 3***Assisting AI Adoption***

Recommendations:

- 3.1 Develop 'how to' best practice resources for industry and Government.
- 3.2 Accelerate Enterprise AI deployment.
- 3.3 Support SME adoption of AI.
- 3.4 Support AI startups and exporters.

THEME 4***Increasing Trusted Data Accessibility***

Recommendations:

- 4.1 Increase data availability and accessibility.

THEME 5***Growing the AI Talent Pool***

Recommendations:

- 5.1 Increase the supply of AI talent.
- 5.2 Encourage AI student diversity.
- 5.3 Teach AI in schools.

THEME 6***Adapting to AI effects on Law, Ethics and Society***

Recommendations:

- 6.1 Establish an AI ethics and society working group.
- 6.2 Review employment practices, law and obligations.
- 6.3 Review high priority legal implications of AI.

INVESTMENT

Early Stage Investment

Historically, early stage funding has been difficult to obtain for New Zealand tech startups.¹²⁰ There is currently some investment support for early stage AI businesses in New Zealand, although more investment could easily be absorbed.

An April 2019 report published by PwC and the Angel Association of New Zealand¹²¹ showed early stage investment steadily increased from over \$30 million in 2008 to over \$100 million in 2018. There is more startup investment than 10 years ago (up from \$19 million to \$55 million) and also at the early expansion stage (up from \$5 million to \$21 million).



Venture capital investors also help supply funding to expansion stage business. In New Zealand key venture capital sources now include KiwiSaver funds as well as the New Zealand Venture Investment Fund (NZVIF).¹²² In April 2019, Simplicity KiwiSaver announced that it was committed to investing \$100 million over the next five years into Icehouse Ventures. Icehouse Ventures is a managed fund designed to accelerate the growth and development of Kiwi companies with global aspirations.¹²³

Furthermore, in the May 2019 Budget, the Government announced a new \$300 million investment into the New Zealand Venture Capital sector, directing the New Zealand Superannuation Fund to invest \$240m to 'fill a gap' in funding for startups that are seeking to expand. The NZVIF is committing a further \$60m of its existing capital.

This is an unprecedented increase in the pool of investment available to New Zealand expansion stage businesses. Many of the businesses will be incorporating AI into their products and services, and it promises to further unblock the traditionally narrow channels of capital into the local tech sector.

Corporate Return On Investment

Corporate investment in AI is beginning to pay off for some large Australian and New Zealand businesses who are early adopters. For example, ANZ Bank's parent announced its half year results and told shareholders that a highlight was improvements it made in automation in its institutional business,

TABLE 6: Recent Venture Capital Investment in New Zealand AI

COMPANY	DATE	AMOUNT RAISED (NZ\$)	VALUATION
Soul Machines ¹²⁴	Oct 2018	\$30m	Not disclosed
FaceMe ¹²⁵	Nov 2018	\$15m	\$74m (reported)
PredictHQ ¹²⁶	Nov 2018	\$15m	Not disclosed
RoboticsPlus ¹²⁷	Nov 2018	\$15m	Not disclosed
Ambit ¹²⁸	Mar 2019	\$1.75m	\$9.75m

through robotics and machine learning. ANZ said it has reduced turnaround times by up to 40 percent in trade, credit and customer service.¹²⁹ Likewise, insurance group IAG noted in its FY18 financial results the contribution of AI to its operations by “deploying pricing capability using machine learning and real-time pricing models, across core personal line portfolios and embedding cognitive capabilities across the organisation, such as chatbots and computer vision”.¹³⁰

Similarly, Spark New Zealand recently credited its investments in digitisation and automation initiatives as underpinning its strong earnings performance in 2019.¹³¹ Automation initiatives have resulted in cost savings and customer experience benefits. An example is their award winning bot which processes fibre connection orders, plan changes and change of address requests automating processes that used to be manual and highly repetitive. (see case study below).


INACTIVITY IN THE LAST 12 MONTHS

During the last 12 months there have been notable areas of inactivity. Primarily, New Zealand still has no national AI strategy and no targeted activity in Government to create one. This is despite indications a year ago that the New Zealand Government would move quickly on developing an AI action plan.¹³²

Development of a national strategy was a core recommendation in the AI Forum’s *Shaping a Future New Zealand* report. In lieu of such a strategy, the AI Forum’s Working Group on *Forging a National AI Strategy for New Zealand* has taken some steps towards developing a draft strategy for discussion,¹³³ but this is a largely voluntary effort and a significant amount of work remains.

Secondly, although the Government did attempt to appoint a Chief Technology Officer (CTO), the ensuing controversy¹³⁴ means that this position remains unfilled over a year after it was advertised. It is unlikely that a CTO is the position that the Government intends to fill, and some have instead called for a Chief Technology Advisor, similar to the Government Chief Science Advisor, or alternatively a Technology Advisory Council.

As yet, there is no research funding targeted specifically for AI projects or PhDs. Although there is \$49 million earmarked over seven years in the Government Strategic Science Investment Fund for big data research, this was not awarded in 2018, and was only in May 2019 open for proposals.¹³⁵ The low levels of AI research funding is a notable gap compared to some of our usual international comparators such as the UK where the NZ\$1.9 billion AI Sector Deal alone includes data science and AI research funding of over NZ\$800 million.¹³⁶



3. An AI Roadmap for New Zealand

“Areas of need include New Zealand’s poor mental health outcomes, significant numbers of children living in poverty, the country’s high levels of greenhouse gas emissions, unequal growth and low productivity, and significant disparities across indicators of wellbeing between Māori and Pacific peoples and other ethnic groups.”

FINANCE MINISTER HON GRANT ROBERTSON, Wellbeing Budget Speech 30 May 2019

THERE IS NO SHORTAGE OF MAJOR CHALLENGES TO WHICH AI TECHNOLOGIES CAN BE APPLIED. NEW ZEALAND'S OPPORTUNITY IS TO ACTION AN AI ROADMAP WHICH HELPS SOLVE THEM.

Placing AI in a New Zealand Context

Some of New Zealand’s high priority challenges in 2019 include:

Environmental:

- Climate change
- Water shortages
- Pollution
- Unsustainable agricultural practices
- Introduced predators
- Deforestation
- Species extinction¹³⁷

Social:

- Child poverty
- High suicide rates
- High incidence of obesity and cancer
- High rates of incarceration
- Improving Māori wellbeing
- Privacy infringement
- Housing shortages, homelessness
- Traffic congestion
- High road fatalities
- Transmission of hate speech and objectionable material

Economic:

- Concentration on primary sector
- Flat productivity growth
- Just transition to a low-emissions economy
- Infrastructure shortfall
- Tourism capacity shortfall

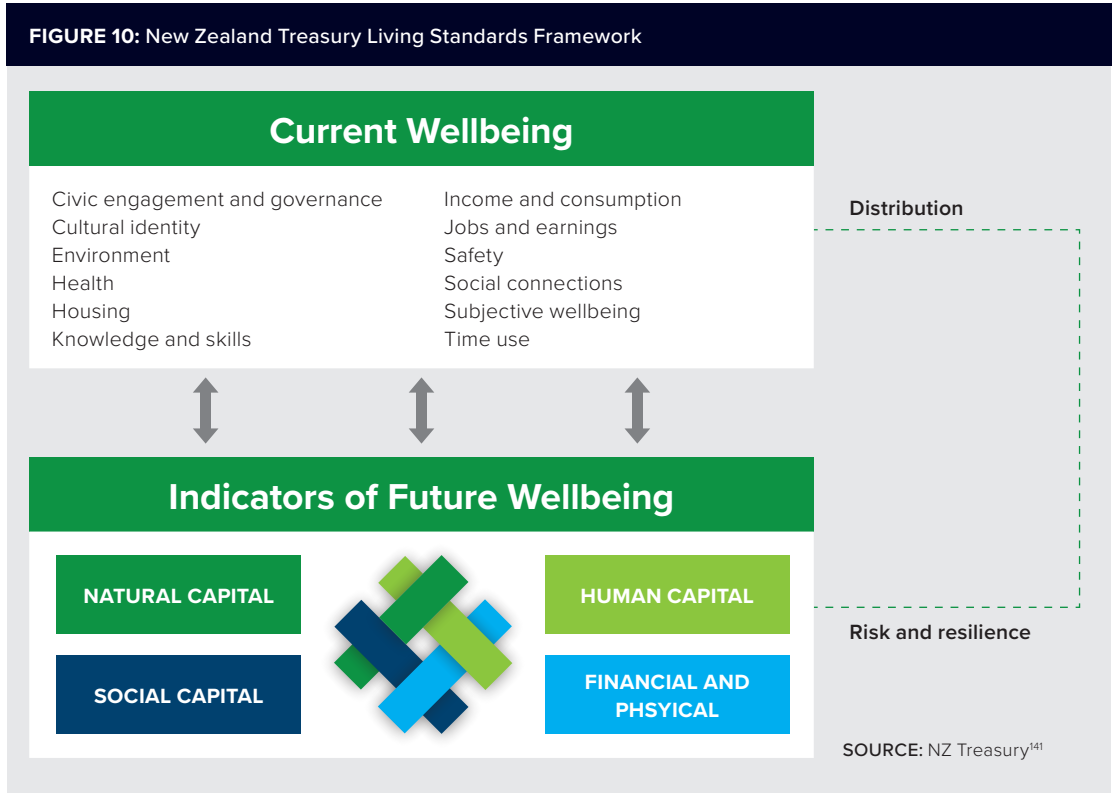
WELLBEING

The New Zealand Government’s Budget Policy Statement for 2019, The Wellbeing Budget, aimed to broaden the Budget’s focus beyond economic and fiscal policy outcomes by using the Treasury’s Living Standards Framework¹³⁸ to inform the Government’s investment priorities and funding decisions.¹³⁹

From 2019 onwards, New Zealand’s Government will measure and report against a broader set of indicators than just Gross Domestic Product (GDP)¹⁴⁰ to show a more holistic measure of success, supported by processes that facilitate evidence based decisions and deliver on Government objectives in a cost effective way.



AI for Wellbeing
AI for Sustainability
AI for the Economy



WELLBEING INDICATORS

Social capability

Culture, relationships, leisure

Empowerment and enablement

Governance, equality, autonomy

Human resource potential

Health, safety, education

Economic self-determination

Work, income and wealth, housing

Environmental sustainability

Land, water, air, waste

Te ao Maori

Partnering with Maori to reflect values, beliefs, and customs

Statistics New Zealand is developing a set of customised wellbeing indicators for Aotearoa New Zealand, based on international best practice. The indicators aim to measure New Zealand’s current wellbeing, future wellbeing (our legacy for future generations), and the impact New Zealand is having on the rest of the world. Examples of potential indicators are air quality, culture, ecosystems and health, human capital, export of waste and more.¹⁴² Statistics New Zealand now provides a suite of resources on its website *Ngā Tūtohu Aotearoa – Indicators Aotearoa New Zealand*¹⁴³ explaining the indicators, their selection and development and how they will be used.

CASE STUDY ENVIRONMENT

ORBICA: Extracting Intelligence from Geospatial Data and Imagery

Christchurch based firm Orbica trains AI systems to identify physical features such as land, water and buildings.

A Geospatial Information System (GIS) stores, analyses and presents geographic and spatial data. This allows scientists to study and analyse information about the physical world and how people use the physical space around them. Often this involves mapping information and also imagery such as satellite views of the earth. Weather and pollution predictions can also be modelled using sensor data sets rather than imagery.

Orbica has developed algorithms that can detect and classify features with a high degree of accuracy and speed, revolutionising traditional methods. This extends beyond buildings, trees, roads and

What is GeoAI?

The combination of Geospatial Information System (GIS) imagery and AI is known as GeoAI. It enables analysis of data in ways that would be impossible for humans to achieve. Extracting intelligence from these datasets helps organisations to manage assets, manage their environmental resources and manage people living in an area.

FIGURE 11: Illustrates how GeoAI can identify and classify buildings, such as houses and garages from imagery. Once overlaid with GIS data about a property this dataset enables use cases in the private and public sector. For example, local councils can automate the process of ensuring new buildings have appropriate building consents. **Supplied by IDC.**

IMAGERY



Satellite or aerial imagery of physical features, in this case, a residential area somewhere in New Zealand.

AI FEATURE EXTRACTION



AI feature extraction detects and classifies buildings. It identifies, for example, houses, garages, sleepouts, swimming pools

GIS DATA



Overlay with Geomaps property information to provide address and therefore property information such as legal descriptions, property area, council valuation, resource and building consents, and land information memorandum.

Concrete Plant, Saudi Arabia



water. The algorithms can be trained to detect and classify any feature on the earth's surface.

At the heart of the 2017 startup was a technology-agnostic approach: the desire to start with the problem at hand, devise the best solution and then identify the best tools. Orbica is pioneering geospatial AI (GeoAI), developing unique data visualisation solutions and enabling organisations to gain near real time visibility of their data.

Orbica aims to create a universal model that it can apply globally. This ongoing development requires extensive machine learning against sophisticated and complex datasets.

On the basis of this work, Orbica won a pitch to solve a challenge for engineering firm Thyssenkrupp. The challenge was how to manage and forecast build progress against schedule on large construction sites. One example, a €900 million concrete plant build in Saudi Arabia is pictured above.

Orbica's solution combines its 3D geospatial and mesh models with AI. Drones capture 3D imagery of the construction site and Orbica grafts the images together into a coherent picture. AI models then

determine the objects in the imagery. The model can identify building construction and ignore irrelevant features such as scaffolding. ThyssenKrupp can get answers to analytics questions such as, how much progress has occurred since the last model runcheck? Or, what are the differences between forecasted progress and actual progress?

Orbica intends to focus future efforts on environmental applications, training its models to extract features such as irrigated lands, roads, forests and vegetation. Orbica can then classify those features and develop insight on how those features change over time. For example, farmers or organisations can better understand how a change in a river impacts local environment and farmlands.



Learn more at www.orbica.world

Māori Wellbeing

ACCOMPANYING THE LIVING STANDARDS FRAMEWORK, TE PUNI KŌKIRI AND THE TREASURY RELEASED THE DISCUSSION PAPER, AN INDIGENOUS APPROACH TO THE LIVING STANDARDS FRAMEWORK (2019).

The paper is a starting point for conversations and a first step towards embedding a focus on Māori wellbeing and Māori concepts of wellbeing into the Living Standards Framework.¹⁴⁴

To realise increased intergenerational wellbeing for Māori, “it will be critical to focus less on the failings of Māori in terms of statistical outcomes and instead look to the potential capability within the Māori population that will support improved wellbeing.”¹⁴⁵

This perspective is echoed by Māori data experts like Kirikowhai Mikaere, who says data about Māori is often collected by Government as part of their service delivery, with the result that the data only reflects one perspective on Māori – often a story of deprivation.¹⁴⁶

A focus on Māori capability and wellbeing may include amendments to the types of data Government collects and more involvement of Māori during the design phase.

The discussion paper proposes that a wellbeing framework that “truly reflects Aotearoa New Zealand” must include an indigeneity lens to address wellbeing for Māori. The proposed three elements for the indigeneity lens include the Treaty of Waitangi/Te Tiriti o Waitangi, te ao Māori and a whānau centred approach. The paper also outlines seven interdependent and interconnected wellbeing domains that, together, describe overall wellbeing for Māori:

- cohesive, resilient and nurturing
- confident participants in society
- confident in language and culture
- living healthy lifestyles
- self-managing
- responsive to the natural and living environment
- economically secure and wealth creating.

Sustainable Development

NEW ZEALAND SUSTAINABILITY GOALS

In September 2015, more than 190 countries including New Zealand signed up to the 17 Sustainable Development Goals of the 2030 Agenda for Sustainable Development. The Agenda sets out a vision to end poverty, protect the planet and ensure prosperity for all.



The 2019 *People's Report on the 2030 Agenda and Sustainability Development Goals*¹⁴⁷ reviews New Zealand's progress towards the SDGs, noting that “the government's new Living Standards Framework and Wellbeing Budget demonstrate a vision [which] aligns in many ways with the SDGs. But currently this vision lacks clear, strong links to the wider, universal vision of the 2030 Agenda and the global framework of the SDGs.”

Providing solid, collaborative data foundations for measuring progress towards the SDGs – and upon which accurate models can be researched and built – will be a key stepping stone towards achieving the outcomes.

Economic Growth

Recently, GDP growth has been slowing in developed countries. There is potential for AI to help reverse this trend, in part through productivity increases, through automation, augmentation, and innovation diffusion.¹⁴⁸ Together these factors mean that AI, economically, is a ‘a capital-labour hybrid’ that provides ‘unprecedented opportunities for value creation’.¹⁴⁹

In line with most other technologies, the adoption of AI is likely to follow an S-shaped uptake curve. However, the exact shape of this curve is presently unknown. For example, whether the lead-in to the upstroke is short or long and how this curve evolves will be critically important to when and how the economic impact appears.

AI investments often fall into the class of assets known as intangibles. These are not physical capital, are difficult to measure, and have often been left out of conventional economic analyses. However, it is becoming increasingly clear that intangibles are a key component of our economic systems. The basic economic properties of intangibles makes an intangible-rich economy behave differently from one predominantly grounded in tangibles. There are at least four ways that intangible capital is different: sunk costs (intangibles may have little resale value), scalability (intangibles may be extended at low marginal cost), synergies (intangibles may be of little use on their own), and spillovers (the idea can be copied).¹⁵⁰



FIGURE 12: S-curves: AI adoption in New Zealand – Rapid, Medium or Slow?¹⁵¹

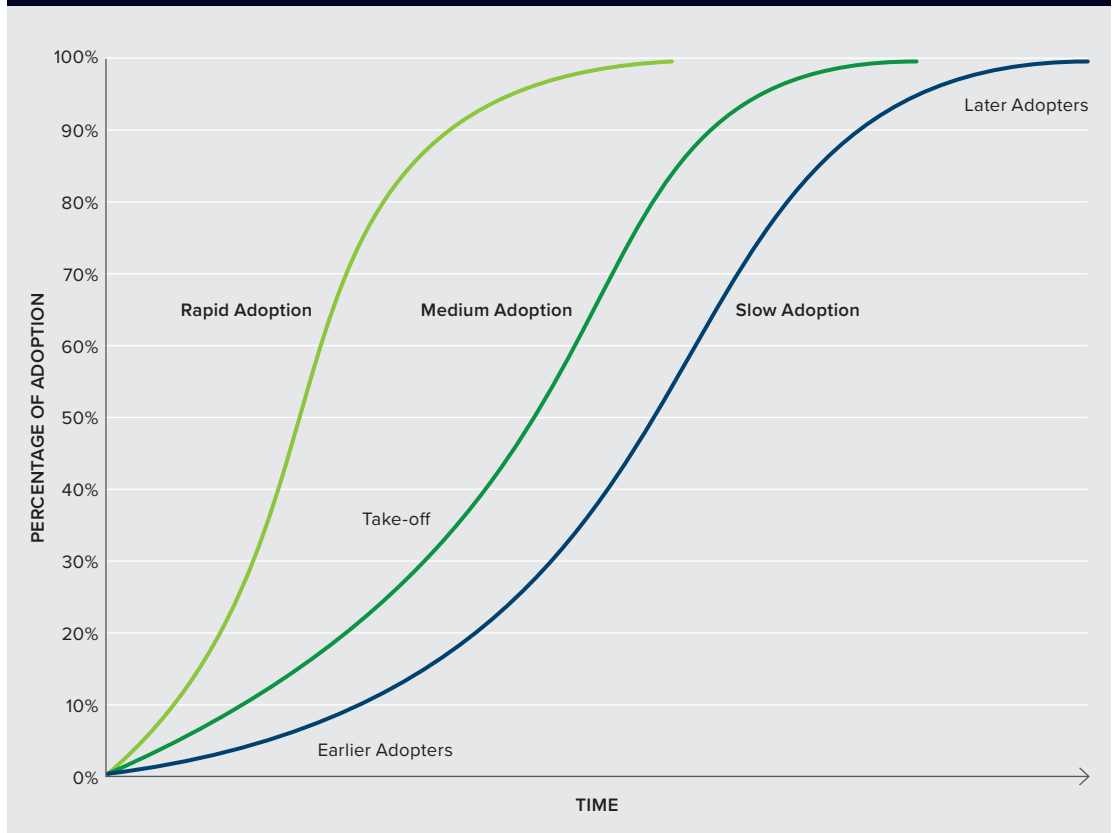



FIGURE 13: Previous technologies have followed a range of 's-shaped' adoption curves, with the internet showing very rapid take-off, and TV a dramatic acceleration phase¹⁵²



SOURCE: Our World In Data, used under CC BY license.

INTERNATIONAL ECONOMIC ANALYSIS

Two major international modelling studies from PwC and McKinsey conclude that the economic impact (as measured by GDP) of AI over the next one to two decades is likely to be significant and positive. Benefits should accrue through productivity gains due to automation and augmentation of work and also through improved product quality. There will be benefits to all sectors and all global regions.

AI will bring productivity gains, but also consumption gains as AI enhances personalisation, utility and saves us time.

Consulting firm PwC (2018) 'quantifies the total economic impact (as measured by

GDP) of AI on the global economy via both productivity gains and consumption side product enhancements over the period 2017-2030'.¹⁵³

PwC's analysis finds that AI will impact all aspects of the value chain and improve products, improve services and reduce frictions across the economy. This will increase value and increase demand for goods and services. The PwC model presents a holistic picture of both production and consumption side responses to AI across regions and sectors.

PwC economic modelling results show there is a clear relationship between AI uptake and productivity. Overall, PwC concludes that global GDP will be 14 percent higher by 2030 with AI compared to

without it. This means a US\$15.7 trillion gain due to an initial phase of automation and augmentation of work, which increases productivity, but flattens after 2025. This is followed by a phase dominated by increased product personalisation and quality, and therefore increased consumer demand. However, this will come mostly after the period of increased productivity. By 2030, 58 percent of all GDP gains will come from consumption side impacts.

Economic Benefits from AI will accrue to the frontrunners

Research by McKinsey (2018) takes a micro to macro approach and comes to a similar conclusion as PwC regarding the direction, magnitude and key drivers of the impact of AI globally.¹⁵⁴ McKinsey considers five key AI technologies: computer vision, natural language processing, virtual assistants, robotic process automation and advanced machine learning.

McKinsey’s modelling predicts 70 percent adoption of at least one of the five AI technologies by 2030. Three production side channels will be particularly important. These are: automation of labour, which will likely add \$9 trillion or 11 percent to global GDP by

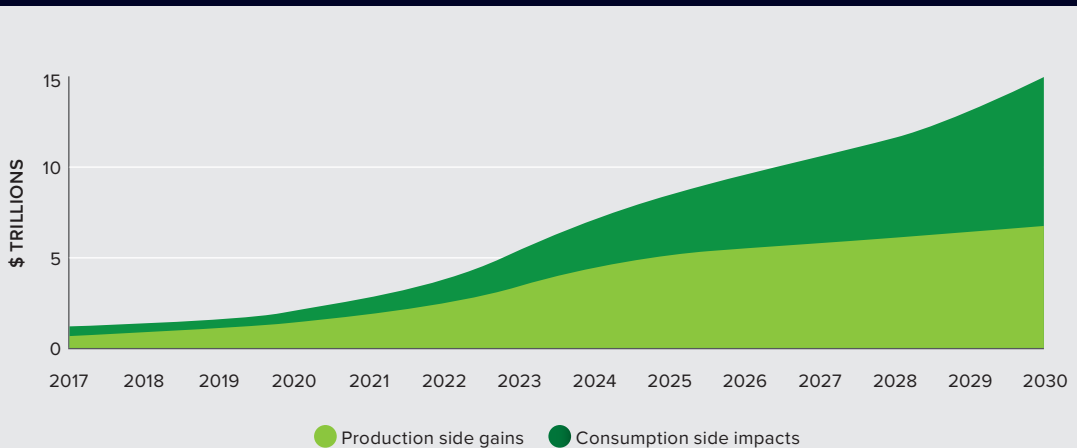
2030; innovation in products and services (\$6 trillion); and negative externalities and transition costs (which will subtract \$7 trillion). These externalities include 14 percent of workers needing to retrain and re-enter the workforce, as well as unemployment and severance costs and the associated reduction in consumption. Overall, AI could raise global GDP 1.2 percent per year and provide \$13 trillion in benefits per annum by 2030. Although 1.2 percent may appear like a small number, the diffusion of IT in the 2000s only drove GDP up 0.6 percent per year. Due to a number of modelling assumptions, McKinsey reports that these results are an upper bound of potential economic impact.

Initial impacts may be small, but as benefits accelerate, there could be a widening gap between countries, companies and workers. Front runners will potentially double returns by 2030, but late adopters could have difficulty generating impact as opportunities pass. It appears that strategic choices, rather than efficiency gains, are likely to have the most economic impact.

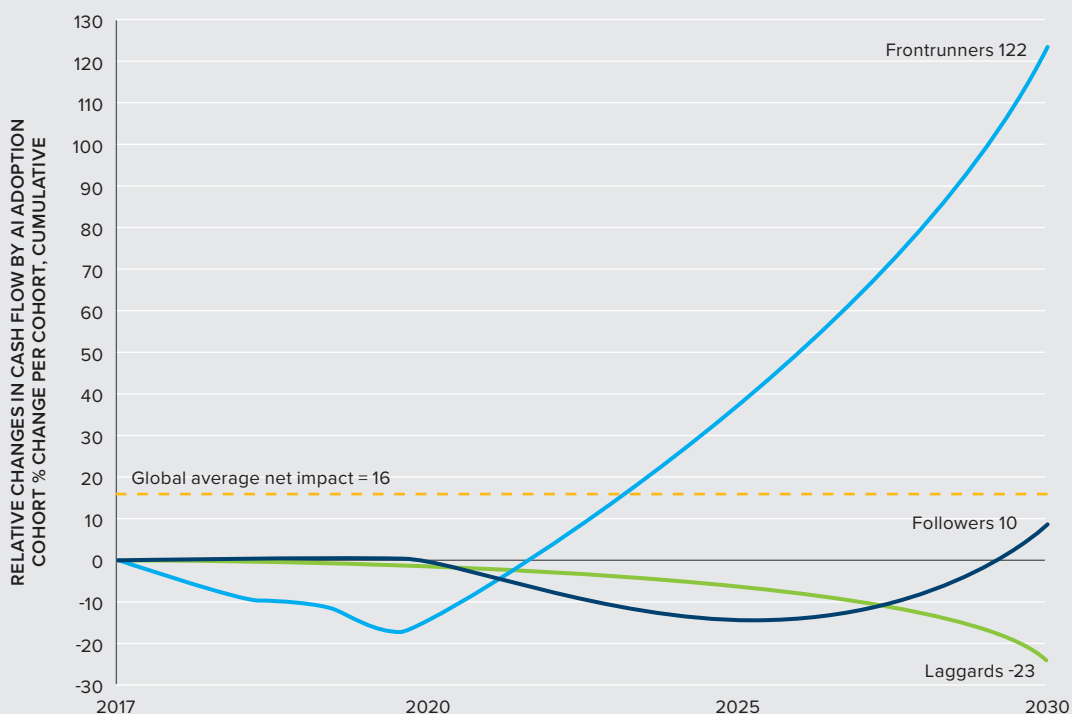
IMPACT BY REGION

According to McKinsey, four groups of countries are classified according to their readiness to

FIGURE 14: Two phases of economic gains from AI: first increased productivity, then increased consumption.



SOURCE: PwC Analysis

FIGURE 15: New Zealand organisations Frontrunners, Followers or Laggards in AI Adoption?

NOTE: Numbers are simulated figures to provide directional perspectives rather than forecasts.

SOURCE: McKinsey Global Institute analysis.

adopt AI. The USA and China constitute Group 1. New Zealand is in Group 2, among 'economies with strong comparative strengths'. This group also includes strong AI performers the UK, South Korea and Canada. An AI divide is likely to emerge between advanced and developing economies.

Likewise, PwC predicts that the greatest gains will be seen in the USA and China, where AI will impact all sectors of the economies. New Zealand is not singled out in PwC's model, but is included in the 'rest of the world' along with Africa and Oceania, a geographical amalgam that is estimated to benefit by only 5.6 percent of GDP; compared to 26.1 percent in China and 10.4 percent in developed Asia.

TABLE 7: Group 1 and Group 2 countries by AI readiness according to McKinsey

GROUP 1	GROUP 2	
China	Australia	Belgium
United States	Canada	Estonia
	Finland	France
	Germany	Iceland
	Israel	Japan
	Netherlands	New Zealand
	Norway	Singapore
	South Korea	Sweden

AI and New Zealand's Economy

As a developed economy with high digital penetration New Zealand stands to gain substantially from AI adoption.

The New Zealand economy is heavily focused on services, which make up about two thirds of GDP.¹⁵⁵ Services account for 74 percent of employment, goods production accounts for 18 percent and primary industry 7 percent.¹⁵⁶

This picture suggests that the New Zealand economy, is likely to benefit from AI in particular ways. Firstly, the PwC model indicates that the services sector, particularly healthcare and education show great potential for personalisation, time saving and improved utility. However, these industries are known as slow adopters of new technology.^{157,158} Although the gains from AI for New Zealand's service dominated economy should come, they may come late.

The modelling indicates that digitally savvy dynamic sectors will benefit disproportionately and already developed economies will benefit to a greater degree than developing economies.


Our Previous Analysis: AI could bring over \$50 billion in benefits to the New Zealand economy by 2035.

The general theory of economic growth driven by AI is premised on the fact that human labour creates value, automation of that labour frees humans from some types of work, and humans can then be reallocated to other productive tasks that AI cannot yet do. Humans can focus on tasks that are 'hard but are necessary to progress'.

Previous analysis by the AI Forum shows that this labour conversion could provide potential economic benefits for New Zealand, across 18 industry classifications, of between \$22.5 billion and \$53.6 billion per annum by 2035.¹⁵⁹

This previous analysis highlights potential labour market impacts for New Zealand that include the creation of jobs, solving some of our demographic workforce problems, and some associated job loss. However, we are unlikely to see mass unemployment.


FIGURE 16: New Zealand employment by sector



SOURCE: Based on Stats NZ data

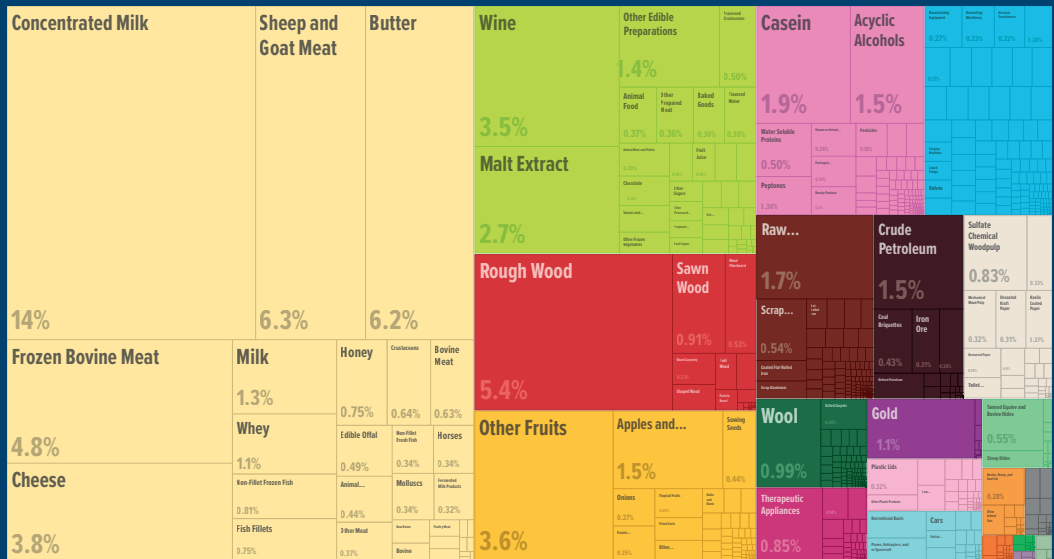
Economic Complexity as an Indicator of Income Inequality

Sometimes analysis of existing data yields new insights for economic and social policy. The MIT Center for Collective Intelligence publishes engaging visualisations of international trade data (see below New Zealand's product exports from 2017). They have published extensive research showing that a country's mix of products predicts its subsequent pattern of diversification and economic growth.¹⁶⁰ In particular, it shows that greater economic complexity is a significant predictor of lower income inequality for a country.



WHAT DOES NEW ZEALAND EXPORT? (2017)

TOTAL **\$37.3B**



SOURCE: OEC, used under Creative Commons Attribution-ShareAlike 3.0 Unported (CC BY-SA 3.0) License.

AI and Public Policy

STRATEGIC RESPONSES AND DECISIONS

There is a growing body of international literature investigating AI related policy options and strategic responses.

As a general purpose technology, AI and its many applications touch on public policy portfolios across all areas of government and society. It is challenging to unpick the horizontal direct AI policy areas from those related to the many application areas.

“...one overarching policy challenge is how best to introduce expertise about AI and robotics into all branches and levels of government so they can make better decisions with greater confidence.”

RYAN CALO¹⁶¹

FIGURE 17: "Heatmap" analysis of 18 different AI strategies across 8 areas of public policy

	RESEARCH	AI TALENT	FUTURE OF WORK	INDUSTRIAL STRATEGY	ETHICS	DATA	AI IN GOVT	INCLUSION
Australia	High	Medium	Low	High	Medium	Low	Low	Low
Canada	High	High	Low	Medium	Medium	Low	Low	Low
China	High	Medium	Low	High	Medium	Medium	Medium	Medium
Denmark	Medium	Medium	High	High	Medium	Low	Low	Low
EU	High	Medium	Medium	Medium	High	Medium	Medium	Medium
Finland	High	Medium	Medium	High	Low	Medium	Medium	Low
France	High	High	Low	Medium	Medium	Low	Medium	Medium
Germany	High	Medium	Medium	High	Medium	Low	Medium	Low
India	High	Medium	Medium	Medium	Medium	Low	Low	High
Italy	Low	Medium	Low	Low	Medium	High	High	Low
Japan	High	Medium	Low	High	Low	Medium	Low	Low
Mexico	High	Medium	High	Medium	Medium	Medium	Medium	Low
Singapore	High	Medium	Low	Medium	Medium	Low	Low	Low
South Korea	High	High	Low	Medium	Low	Low	Low	Low
Sweden	High	Medium	Medium	Medium	High	Medium	Low	Low
Taiwan	Medium	High	Low	High	Low	Medium	Low	Low
UAE	Low	Medium	Low	High	Medium	Low	High	Low
UK	Medium	High	Medium	High	Medium	Medium	Medium	Medium

SOURCE: Analysis, Tim Dutton, Building an AI World, Report on National and Regional AI Strategies CIFAR.

Over twenty countries have now released some form of national AI strategy.¹⁶² The first was Canada, in March 2017. Another ten governments, including New Zealand, have indicated an intention to present something in the coming year. However, New Zealand is yet to release a national AI strategy and at the time of publication there is no clearly defined work happening in Government towards this.

In Chapter 2 we compared the magnitudes of per capita public funding announced for nine national (and EU) strategies.

These strategies vary substantially from country to country in their focus. Analyst Tim Dutton, previously of the Canadian Institute for Advanced Research (CIFAR) identifies eight main areas of public policy across 18 national AI strategies: Research, AI Talent, Future of Work, Industrial Strategy, Ethics, Data, AI in Government and Inclusion.

Further analysis by Dutton suggests:

- Encouraging private sector AI adoption is a priority for eight of the national strategies
- Scientific research and building AI talent are priorities for seven countries.
- Work on ethical AI standards, broader data and digital infrastructure developments or to expand the use of AI in government are priorities for a few countries.
- Making sure AI is inclusive, and skills development for the future of work are the least and second-least prioritised policy areas, respectively.
- Only five strategies touch on all eight policy areas, while seven are quite focused and touch on four or fewer areas.

Specific examples include Canada and South Korea, which both have research oriented strategies that include the establishment of new AI research institutes and provision of funding to attract, train and retain AI talent in the country.

The UK's strategy touches on all eight policy areas, lifting funding for AI research and education, encouraging private sector investment, creating a centre to explore issues in data ethics and a new AI Council to advise the government, and establishing new arrangements to make trusted data sharing easier.

“Continued American leadership in Artificial Intelligence is of paramount importance to maintaining the economic and national security of the United States.”¹⁶³



PRESIDENT DONALD J. TRUMP

The American AI Initiative was announced by President Trump in February, 2019. The American AI Initiative,¹⁶⁴ is a concerted effort to promote and protect the USA's AI technology and innovation. It aims to implement a Federal whole-of-government strategy in collaboration with the private sector, academia, the public and like-minded international partners. This announcement came after a sustained period of inactivity following the end of the Obama administration in 2016.

“We need to speed up building China into a strong country with advanced manufacturing, pushing for deep integration between the real economy and advanced technologies including internet, big data, and artificial intelligence.”¹⁶⁵



CHINESE PREMIER XI XINPING

China's Next Generation Artificial Intelligence Development Plan is a three phase strategy to make China the world leader in AI theories, technologies, and applications by 2030. It contains efforts across all eight policy areas, including breakthroughs in fundamental AI research, making China the world's primary AI innovation centre, and leading the world in standards setting and AI ethics.



The EU, Singapore, Australia and New Zealand co-funded European-Pacific Partnership for ICT collaboration in research, development and innovation (“EPIC”) aims to strengthen the ties between these regions through academic, industrial and political partnerships.¹⁶⁶ The project identifies 11 fields of action across national AI strategies:

1. Research
2. Innovation
3. Startup and risk capital
4. People
5. Infrastructure
6. Data
7. International
8. Regulation
9. Government
10. Social, Labour and Ethics
11. Strategy

The project recently published a policy brief providing guidance for EU and New Zealand on how these regions can collaborate to successfully deploy AI technologies in fields such as Agritech.¹⁶⁷

OECD RECOMMENDATION ON AI

In May 2019, the 36 member Organisation for Economic Cooperation and Development (OECD) formally adopted the Recommendation on Artificial Intelligence¹⁶⁸ which includes the first international standards agreed by governments for the responsible stewardship of trustworthy AI.

The five principles for responsible stewardship of trustworthy AI are:

1. sustainable development and wellbeing
2. human-centred values and fairness
3. transparency and explainability
4. robustness, security and safety and
5. accountability.

The principles explicitly reference the Sustainable Development Goals from the 2030 Agenda for Sustainable Development adopted by the United Nations as well as the 1948 Universal Declaration of Human Rights.

Consistent with these value based principles, the OECD also provided five recommendations to governments:

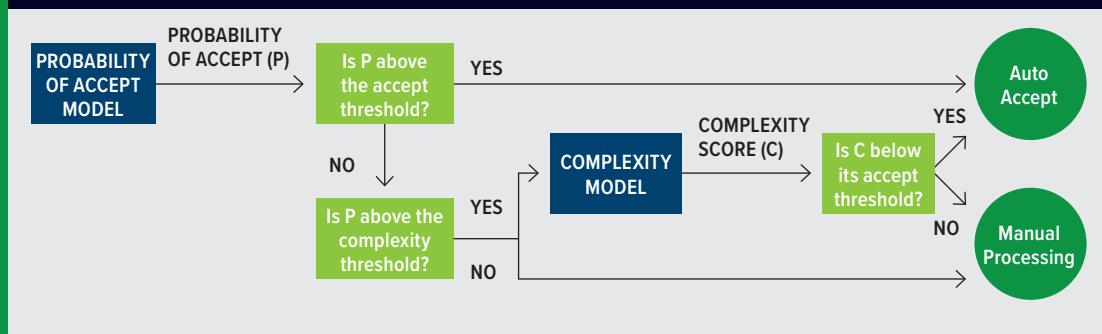
1. **Investment in Innovation** Facilitate public and private investment in research and development to spur innovation in trustworthy AI.
2. **Foster AI Ecosystems** Foster accessible AI ecosystems with digital infrastructure and technologies and mechanisms to share data and knowledge.
3. **Policy Environment for Trustworthy AI** Ensure a policy environment that will open the way to deployment of trustworthy AI systems.
4. **People and Skills** Empower people with the skills for AI and support workers for a fair transition.
5. **Cooperate Across Borders and Sectors** Cooperate across borders and sectors to progress on responsible stewardship of trustworthy AI.

No matter what path is chosen for New Zealand, it is always helpful to keep an eye on what other countries are doing. The OECD now usefully provides a clearing house for AI related developments around the world. Following the Recommendation, it also plans to launch in 2019 an online AI Policy Observatory to bring together evidence, analysis and options on how to ensure the beneficial use of AI.¹⁶⁹

As with other general purpose technologies, the widespread use of AI will raise broader policy questions, for instance in education systems, labour market matching, the regulation of data use, income support, competition policy, and industry structure, amongst others.

CASE STUDY PUBLIC SECTOR

ACC: Claims Approval Algorithm

FIGURE 18: ACC's Cover Decision Algorithm¹⁷⁰

The Accident Compensation Corporation (ACC) reviews and processes over two million claims per year, with a total claim acceptance rate of 96 percent. This manual process previously took the efforts of 60 staff and the result was that approximately 90 percent of claims are straightforward acceptances.

Given that most claims are approved, the ACC decided to investigate big data and AI solutions to automate the claims approval process. The goal was to find a way of automatically approving straightforward cases, leaving unclear or complex cases for humans to decide. Process automation would need to reliably identify which claims clearly fit within the criteria of the Accident Compensation Act 2001.

ACC partnered with Nicholson Consulting and determined that a combination of business rules and a set of statistical models could be used to automate the process, enhancing consistency and saving labour.

The predictive models were built based on twelve million anonymised claims lodged between 2010 and 2016. The model determines the probability that a given claim would be accepted. This automated solution combines big data, predictive analytics and process automation.

One digital product owner at ACC said that quality assurance evaluation shows the automated system makes as few as two incorrect approvals in 18,000. Humans now only need to review approximately 14 percent of claims. Highly experienced service

coordinators deal with the more complex claims or sensitive issues, including sexual abuse or assault. Most importantly, any decision to deny cover will always be reviewed and dealt with by a team member.

ACC's automation ensures faster, more accurate processing, while ensuring the algorithm cannot deny claims.¹⁷¹ Experience developing and deploying statistical models like ACC's can help build a data science foundation to facilitate subsequent use of AI solutions across the organisation and across government.

"Thanks to the services of predictive models, a doctor can now hit submit on the practice management system in a GP surgery and the client could receive a cover notification via text within seconds,"¹⁷² says Nicholson Consulting senior data scientist Ernestynne Walsh.



Te Kaporeihana Āwhina
Hunga Whara
prevention. care. recovery.

Learn more at www.acc.co.nz

CASE STUDY SOCIAL ENTERPRISE

SUSTAINABLE COASTLINES:

Cleaning up New Zealand's Coastline



Sustainable Coastlines support large scale coastal cleanup events and public awareness campaigns. The charitable trust raises awareness of marine litter and inspires communities to take care of beaches and waterways.

Sustainable Coastlines was recently awarded a Microsoft AI for Earth grant for its litter-busting efforts. A national litter database is being established so schools, iwi, community groups and businesses can view the data and trends in their areas and work to change local behaviour and assess which measures are the most effective in reducing rubbish long term. The platform uses United Nations Environment Programme methodology combined with Microsoft Power BI and Azure technology to chart litter along the coast.

Sustainable Coastlines' new collaboration with Enlighten Designs will provide an open and free global litter database. Data will be collected by trained citizen scientists and analysis of the data will inform decision making to reduce litter. The citizen scientists will use an approved consistent method to ensure rigor of the data.¹⁷³ The intelligent platform also evaluates international weather, tide and ocean current data to help predict future litter accumulation.



Learn more at www.sustainablecoastlines.org

FIRST STEPS FOR AI POLICY IN NEW ZEALAND

“In a world of changing and interlinked policy measures, data science and AI can provide policy makers with unprecedented insight: from identifying policy priorities by modelling complex systems and scenarios, to evaluating hard-to-measure policy outcomes.”

THE ALAN TURING INSTITUTE (UK)¹⁷⁴

There is huge potential for using AI to benefit New Zealand and help solve the most pressing challenges which are facing New Zealanders now and in the years ahead. Specifically:

- Improving wellbeing and solving complex social and environmental challenges
- Increasing the productivity and international competitiveness of New Zealand businesses and the resilience of the New Zealand economy
- Enabling better Government service design and delivery, higher public sector productivity and value for money.

Furthermore, AI supports continuous policy improvement. Over time, AI promises more dynamic predictive models based upon wellbeing indicators combined with many other data sources, enabling higher precision and policy insights. In theory this leads to a virtuous cycle of improvement in evidence supporting policy development.

Achieving this potential will require significant investment and policy interventions.

A POLICY MAP FOR AI IN AOTEAROA

All policy makers in New Zealand need to have a perspective on AI and how it impacts decisions in their domain.

Below we introduce an AI policy map, tailored for the New Zealand context, provided for further discussion and development. For consistency with existing international work the structure follows the recent policy recommendations around Trustworthy AI from the OECD.

The key insight arising from compiling this map is that AI's potential effects cut horizontally across nearly all Government portfolios. The New Zealand Government's response to AI will be distributed widely across the machinery of the public sector.

Furthermore, many of these policy issues transcend national borders in scope and scale and can only be effectively addressed through international cooperation. For these issues, New Zealand needs to apply its resources to actively participate, contribute and benefit from international efforts to create effective AI governance and regulation which spans the globe.

We look forward to exploring these diverse policy issues and options in more detail in future AI Forum events and publications.



FIGURE 19: AOTEAROA AI POLICY FRAMEWORK

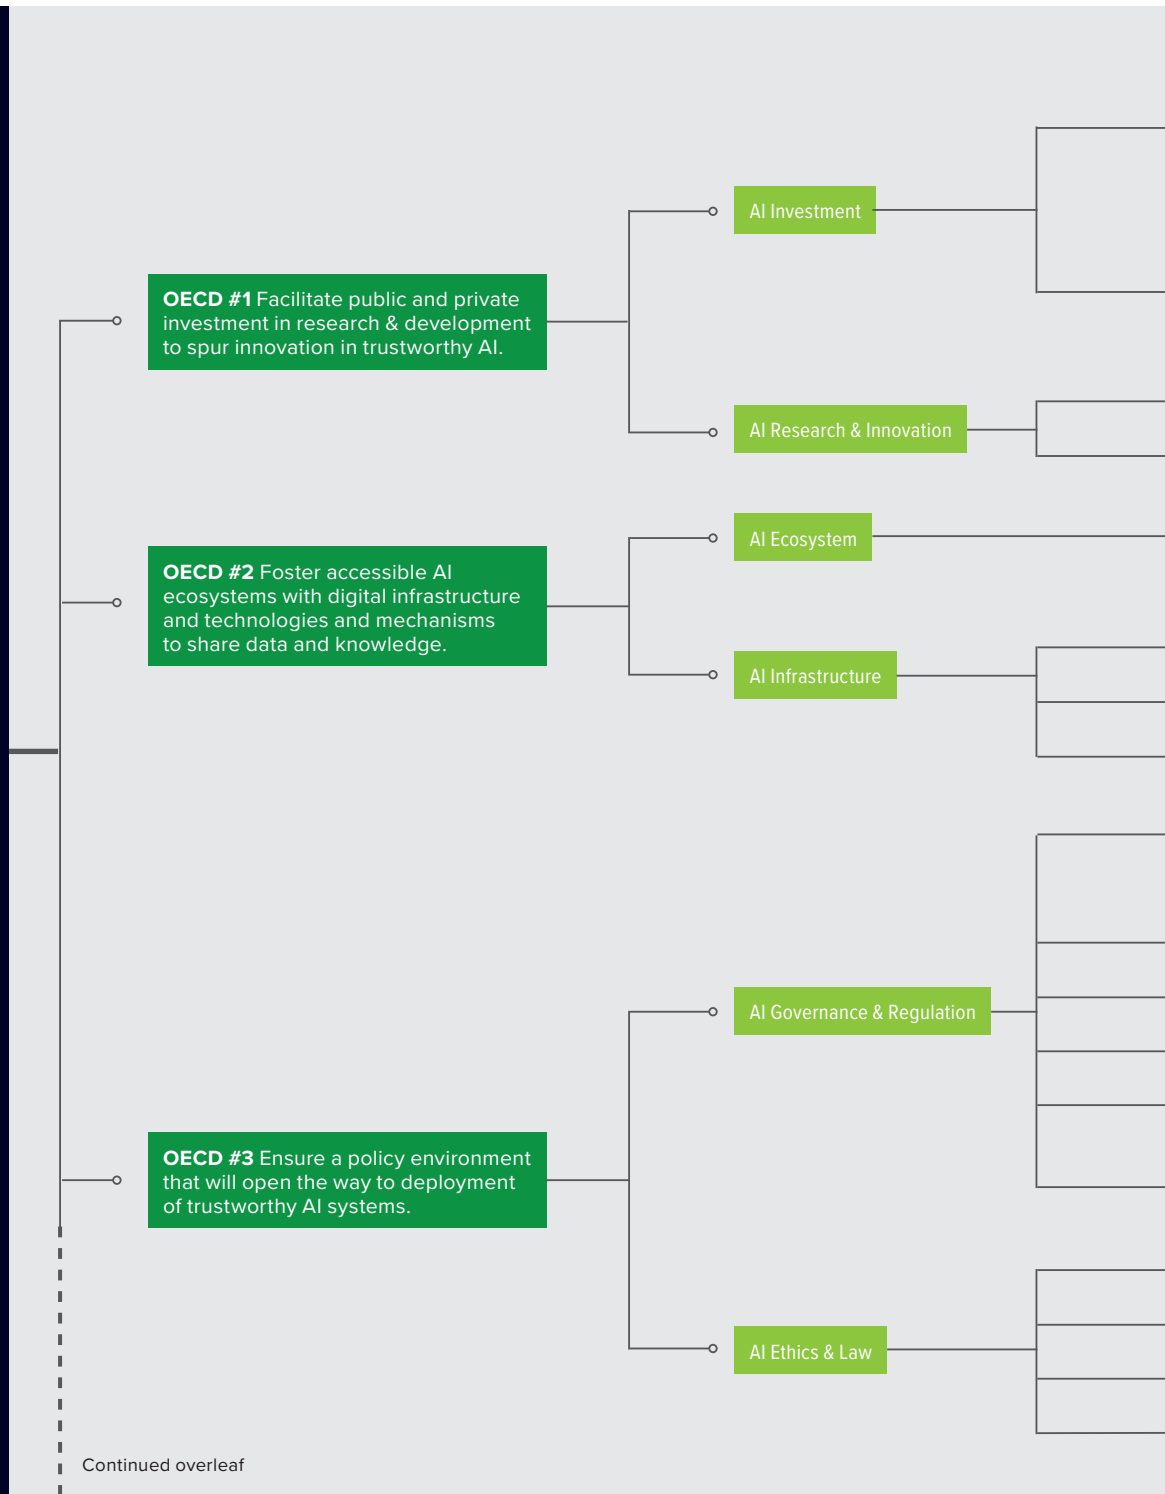




FIGURE 19: AOTEAROA AI POLICY FRAMEWORK





WHAT DOES A FUTURE WITHOUT AI POLICY OR REGULATION LOOK LIKE?

One way to decide where we need policy and what it may look like is to consider a future without AI policy. A non-exhaustive list of potential risks includes:

- no AI infrastructure policy**
 Without a big data and AI infrastructure policy, New Zealand is vulnerable to being turned off in conflict (or other) situations if we depend heavily on third party applications. We face a computational security issue.
- no AI ethics policy**
 Without systematic consideration of AI ethics, New Zealand could be vulnerable to a large proliferation of biased and unjust algorithms, or unethical, exploitative business models.
- no AI safety policy**
 Without regulation of safety standards for critical AI applications, New Zealand could be vulnerable to breakdown of critical systems or harm to individuals.
- no AI training and education policy**
 Without a plan for upskilling our workforce, New Zealand may forego the option of widespread AI adoption and industry may become increasingly susceptible to international disruption.
- no data protection policy**
 In the absence of appropriate data protections, New Zealand could be vulnerable to the emergence of big data organisations that behave in manipulative and totalitarian ways.
- no workforce policy**
 An uncoordinated workforce policy could see a few hungry sectors snatch all the AI talent, or see academia drained of talent to fill industry positions, universities could languish, training throughput could decrease and New Zealand slips behind economically.
- no AI defence strategy**
 Without serious consideration of cyber defence, New Zealand could be vulnerable to critical AI driven cyber attack.
- no AI strategic policy**
 Without a coordinated national (not just Government) AI strategy, New Zealand risks missing an opportunity to use intelligence to solve its grand challenges, and risks falling behind economically.

New Zealand Shouldn't Wait for an AI Strategy

Earlier we noted the lack of momentum to advance towards a national AI Strategy for New Zealand.

It is imperative that New Zealand institutions and businesses act with agency rather than adopting a passive “wait and see” approach for AI to roll over us – AI technologies are ready to be applied to achieving our existing goals.

The introduction of the Government's national wellbeing objectives has provided a ready made framework towards which AI policy and investment can be targeted. The UN's 17 Sustainable Development Goals provide a global set of goals towards which AI investment can be targeted directly, whether through private, public or other partnership funding

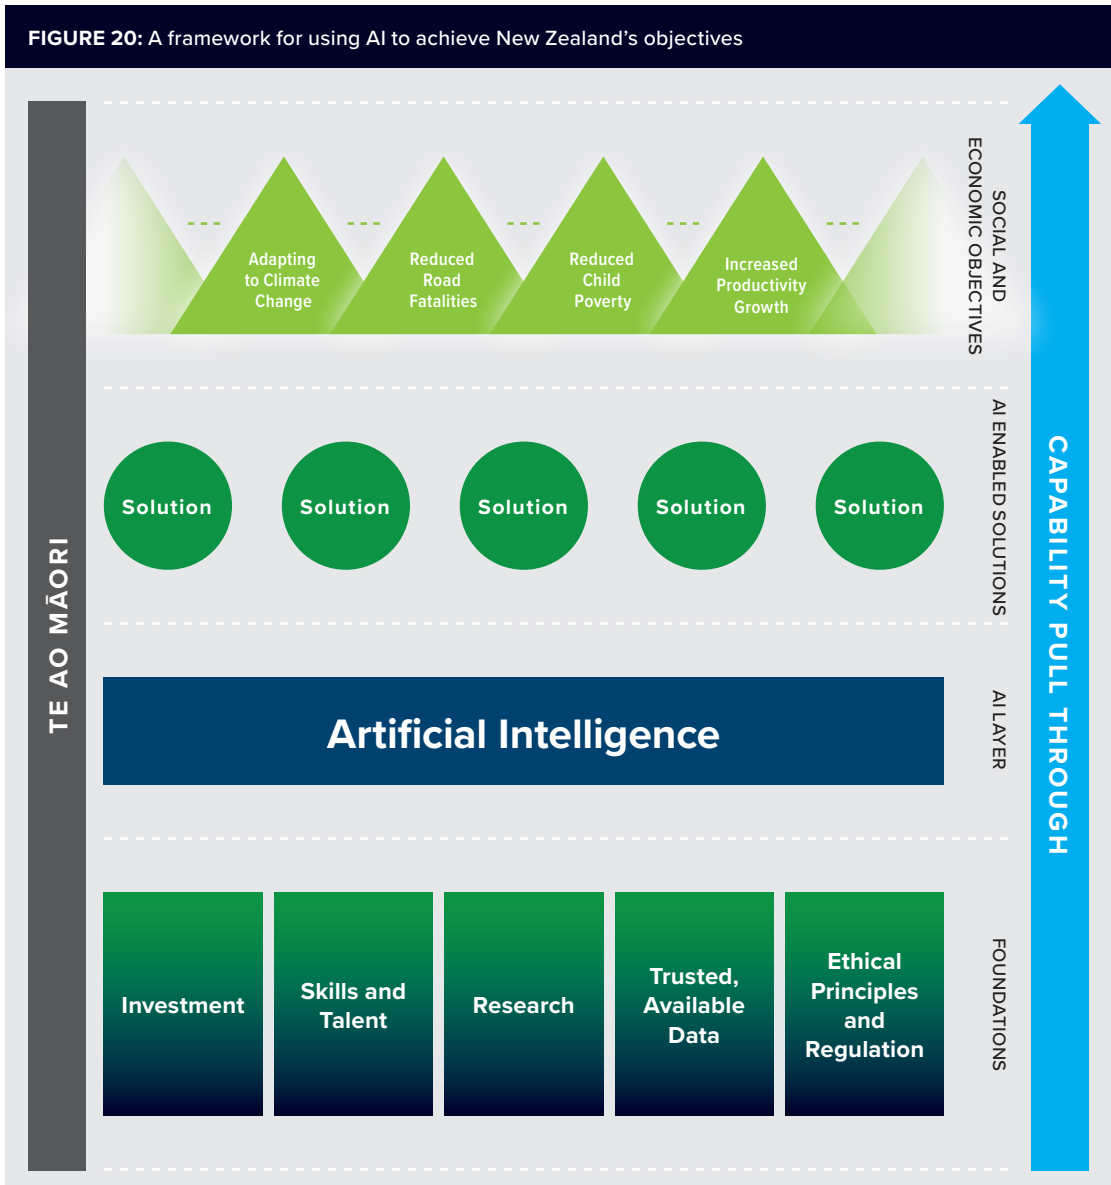
models. Furthermore, AI as the newest general purpose technology should feature in every one of the Government's forthcoming Industry Development Plans for economic growth.

AI, when applied with appropriate checks and balances – promises powerful tools to support New Zealand's Wellbeing, Sustainability and Economic objectives.

In addition, New Zealand developed AI solutions in areas of global importance, such as those identified in the SDGs, offer the opportunity of a worldwide export market for New Zealand's fast growing technology sector.

ALIGNING AI TO NEW ZEALAND'S OBJECTIVES: A FRAMEWORK FOR ACTION

In the Executive Summary we introduced a framework for New Zealand to rapidly identify the opportunities and enabling pillars to apply AI to high priority national objectives. In the following sections we use examples to illustrate how the framework could be applied to achieve wellbeing outcomes in Health, Transport and Conservation.



Health

REDUCING DEATHS FROM BREAST CANCER

Over 600 women die from breast cancer each year in New Zealand.¹⁷⁵ With an AI augmented screening programme many of these deaths might be preventable. Such a programme could pay for itself, given that the total cost of breast cancer in New Zealand is over \$126 million per year, with \$45,000 spent per diagnosed case.¹⁷⁶ By 2021 the cost of cancer care is predicted to rise by 20 percent.¹⁷⁷ Once proven such an approach to screening could then be generalized to other cancers.

AI can be used to predict breast cancer, potentially five years in advance,¹⁷⁸ diagnose it from medical images with comparable ability to radiologists,¹⁷⁹ and decrease the error rate of pathologists by 85%.¹⁸⁰

Early detection of breast cancer dramatically improves the outcome for the patient. In New Zealand, improved screening for breast cancer over the past 20 years has reduced deaths from the disease by 27%.¹⁸¹ **If AI can further augment and improve breast**

cancer screening by even 10% this would result in saving 60 women's lives each year and reduce the cost to the taxpayer by \$12.6 million a year.

The need for early diagnosis underpins breast cancer screening. But rather than a one size fits all approach, AI driven predictive analytics could mean that screening in New Zealand is offered according to risk, thereby targeting more resources to those most in need, and sparing those at very low risk from false positive results and unnecessary interventions. Twelve percent of women will develop breast cancer in their lifetime, but this does not mean that each woman has a 12 percent risk. Risk is much higher with certain genes, such as the BRCA gene and with a family history of breast cancer.¹⁸² There are potentially other patterns of risk that we don't yet fully understand.

New Zealand has an opportunity to link and leverage our comprehensive health datasets and use this big data to train intelligent predictive models for breast cancer.

Combining local and national datasets such as primary health records, electronic health records, the cancer registry, mortality collections, and adding

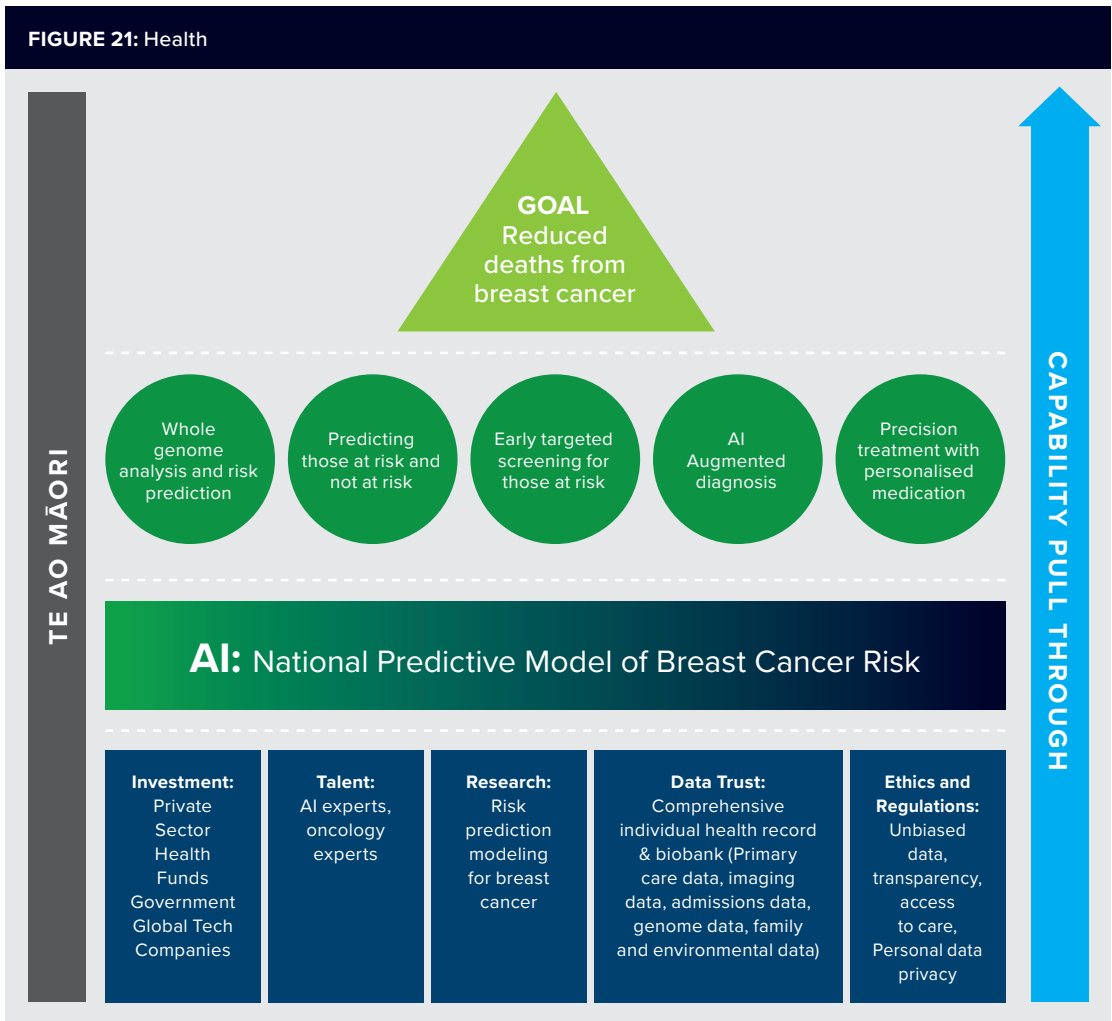


new datasets over time such as a biobank, could yield a national-level predictive model for breast cancer risk, able to provide far more granular and relevant predictions for high disease risk. This model could be used by GPs and DHB oncology services to enable intensive monitoring and prevention or pre-emptive treatment of women at high risk for breast cancer, as well as reducing the screening frequency of those with vanishingly low probability of

cancer, thereby minimising discomfort, anxiety around false positive results, and freeing up resources.

Such a solution relies on AI foundations: skilled AI talent with domain knowledge and expertise building predictive models in health and medicine, research on applying AI to disease predictions and ensuring that data, algorithms and models employed are thoroughly evaluated (initially in parallel with usual care) and reviewed against ethical principles.

FIGURE 21: Health




Transport

REDUCING DEATHS FROM ROAD ACCIDENTS

The use of AI has been found to improve safety on Nevada highways by up to 17%.¹⁸³ By combining traffic data, weather data and in-vehicle information the system develops prediction models that reduce congestion and identify emerging accident areas. The system then alerts traffic agencies where to take preventative action. During the pilot study in Nevada, 91 percent of drivers reduced speeds where AI enabled preventative measures were deployed.

The New Zealand government has calculated that the total social cost for New Zealand of motor vehicle crashes in 2017 was \$4.8 billion dollars. This includes the costs of all injuries recorded by the Police, hospitals and ACC. Additionally, 378 people died on the roads that year and a further 42,398 were injured.

Driving too fast for the conditions accounted for 32% of the road accidents in 2017.¹⁸⁴ If New Zealand developed a similar system to Nevada, **AI has the potential to save 64 lives a year, prevent 7,200 injuries and over \$800 million in social costs.**¹⁸⁵



Combining national datasets on road accidents, geospatial road data, vehicles, drivers, emergency services response, weather could yield a national-level predictive model for road accidents, able to provide far more granular and real-time predictions for high accident risk. This model could be exposed directly to drivers – for example through an app or in-car device, through automated road signage, variable speed limits and potentially enable emergency services to optimise where their vehicles are positioned to optimise response times

Such a solution relies on AI foundations: skilled AI talent with domain knowledge of transport,

research on applying AI to road traffic predictions and ensuring that data, algorithms and models employed are reviewed against ethical principles.

Conservation

IMPROVING PRODUCTIVITY FOR WILDING CONIFER MANAGEMENT

Wilding pines are invasive conifer trees in the high country of New Zealand. They are present conservation land as well as privately owned land and other public land such as roadsides. Wilding



pinus are considered to be a threat to biodiversity, farm productivity and landscape values.

20 percent of New Zealand will be covered in unwanted wilding conifers within 20 years if their spread isn't stopped. They already cover more than 1.8 million hectares of New Zealand and in 2017 they were spreading at about 5 per cent a year.

If wilding conifer management productivity was significantly improved through better precision and automation, the government would save a big proportion of the \$27 million per year it is currently spending on wilding control.¹⁸⁶

A real time “digital twin” national Wilding Pine AI model could be created which takes in satellite and aerial photography, maps demarcating conservation land, private land and commercial forests. This model could be trained to detect young pinus with high precision. Emerging applied research in robotic forestry and invasive tree control could be combined with autonomous drone technology to deliver an automated management solution which manages to control and ultimately eradicate wilding pinus from the New Zealand landscape.

Such a solution would need regulatory changes allowing free-flying drones and ensuring safety of autonomous robotics.

Call to Action – AI Roadmap Next Steps


These three examples illustrate how the AI Policy Framework above can be leveraged to identify and develop areas of AI importance, and help direct AI investment and activity to achieve NZ's social, environmental and economic objectives.

- The role of AI should be consciously considered for all of New Zealand's priority policy objectives.
- AI should feature in every one of the Industry Development Plans for economic growth.
- New Zealand should collectively and quickly identify a manageable number of high-value flagship opportunities to apply AI – for example in health, transport and the environment.
- Identify mechanisms to leverage public-private-iwi-academic partnerships to invest in these flagship projects.
- Use these AI investments to accelerate the pull-through of core AI capabilities needed for New Zealand – this will inform the development of a pragmatic strategic approach – rather than waiting for a formal national AI strategy to emerge fully-formed.

PHOTO: Wilding pinus looking west towards the Southern Alps.



Credit: Alan Liefing



PART B:

Enabling AI Adoption in New Zealand

4. Research

“Creating value from [the] flood of data requires novel capability and breakthrough research, and connections to global centres of data innovation.”¹⁸⁷

NEW ZEALAND GOVERNMENT STRATEGIC SCIENCE INVESTMENT FUND STRATEGY 2017–2024

THIS SECTION OUTLINES CURRENT RESEARCH ENGAGEMENT ACROSS NEW ZEALAND’S RESEARCH INSTITUTIONS. IT ALSO DETAILS LOCAL RESEARCH FUNDED BY GOVERNMENT, INSTITUTIONS AND THE PRIVATE SECTOR.

Research is a key driving force for successful AI deployment. It provides support for the raw resources of a vibrant national AI ecosystem including data, data standards, social license, skills, talent, organisational capability, tools and infrastructure. Both applied and fundamental research are critical to the continued advancement and innovation of core/underlying AI technologies (such as algorithms or computation), as well as deploying AI solutions within and across industries. As such, research is frequently highlighted as a key strategic focus in other national AI strategies.¹⁸⁸

The New Zealand Government Strategic Science Investment Fund Strategy 2017–2024 states that, “Creating value from [the] flood of data requires novel capability and breakthrough research, and connections to global centres of data innovation.”¹⁸⁹

For an indication of the scale of research happening worldwide, the institutions of the Chinese Academy of Sciences have published over 20,000 research papers on AI, and have 2500 AI patent families (where a patent is claimed across additional countries).¹⁹⁰



CASE STUDY HEALTHCARE

PRECISION DRIVEN HEALTH:

Triaging Cardiologist referrals from General Practitioners

Managing patient referrals has traditionally been a paper-based process. The time consuming manual process requires prioritising referrals according to urgency levels and protocol.

Founded in 2016, Precision Driven Health (PDH) is a research partnership between Orion Health, the University of Auckland and Waitemata District Health Board. PDH is investing \$38 million over seven years in collaborations to personalise healthcare by running machine learning on data. Studies and early technology adoption are key focuses of this partnership.

Waitemata District Health Board (WDHB) identified cardiologists were spending too high a proportion of their time triaging patient referrals and reviewing paperwork. Unsurprisingly, cardiologists would prefer to spend more time treating patients and less time processing referrals!

PDH aims to use transfer learning to reduce the time taken for hospital cardiologists to prioritise patient referrals from general practitioners. Reducing the triage time allows cardiologists to spend more time with patients, reduce patient waiting time and improving delivery of the right treatment and care.

Dr Edmond Zhang, a data scientist at PDH, is using different machine learning models on different data types, including images, scans, text-based data, and structured data such as laboratory test results. These outputs are combined into a unified model to triage patients.

PDH will use 14,000 existing WDHB cardiology triage records to create the base machine learning model for this project. To complete the transfer learning aspect of the project, this base model will then be fine tuned for other contexts. As at June 2019, PDH is currently mid way through base model production and a functional triaging model for cardiologists should be ready later this year.

Ultimately, PDH will develop the research outputs into a commercial platform. Orion Health's Smart Data Platform will enable clinicians to use these tools to support their decision making.

Challenges that PDH are facing on their journey include:

- **Data Access:** PDH has access to 14,000 cardiology triage records from WDHB, however gaining access to a wider set of historical triage records has proven difficult. PDH needed to gain data ethics approval as well as anonymising free text data.
- **Moving the models into clinical practice:** The time delay between a proof of concept and deploying a model can mean situations change, inconsistencies develop, or datasets may look different
- **Gaining access to a high performance computer:** Due to privacy and for security, the computer remains on the premises during the project.

The benefits of using the Smart Data Platform include:

- time and cost savings
- combining multiple data sets means better insights
- transferring a base model to different specialisations and hospitals

The plan for the Smart Data Platform is to include a base model that will apply to specific industries and offer Data Science as a Service, either on premise or cloud based.



Learn more at www.precisiondrivenhealth.com

CASE STUDY HEALTHCARE

VOLPARA SOLUTIONS:

Tapping AI in the Fight Against Breast Cancer



Wellington-based Volpara Solutions' mission is to reduce the mortality and cost of breast cancer by providing clinically validated software that underpins personalised, high quality breast cancer screening.

It can be difficult to isolate cancers in breast tissue that is particularly 'dense', as both cancer and dense breast tissue appear white on X-rays. Volpara's software, built on Microsoft Azure, assesses more than 100 variables to help radiographers position the patient, capture higher quality images and automatically and objectively assess breast density.

These assessments mean that radiologists can quickly and objectively identify which women may benefit from supplemental imaging such as ultrasound. Not only does this enhance quality control, it also helps decrease costs by reducing the number of women recalled for repeat screening. This is important in the context of waiting lists and a global radiologist shortage.

The company recently reported annual recurring revenue growth of 86 percent to NZ\$6.63 million and full year cash receipts growth of 83 percent to NZ\$5.6 million. Volpara's technology is now used by 7.1 percent of the USA breast screening market.¹⁹¹ Additionally, through Microsoft's Co-Sell programme Volpara has arranged a trial of their technology with healthcare providers in Singapore.

volpara[®]
solutions[™]

Learn more at www.volparasolutions.com

Examples of AI Research in New Zealand

FUNDAMENTAL AND APPLIED AI RESEARCH

The AI Forum's previous research report, *Shaping a Future New Zealand* surveyed the AI research occurring in New Zealand. It identified that research is often fundamental, rather than applied. Furthermore, pipelines to commercialization are rarely in place. There is a risk novel AI technology developed in New Zealand may be late to be deployed commercially. Potentially local AI products may enter the market after similar technology is already cheaply available through cloud services elsewhere. However, our snapshot of AI research in New Zealand in 2019 suggests more applied research is now taking place and with greater visibility across the ecosystem.

Overall, New Zealand is still likely to benefit from better coordination and increased funding for applied AI research.¹⁹² However, we must also preserve deep institutional research capability as many of the most fundamental advances in AI have emerged from universities where there is less pressure to turn an immediate profit. If too many academic researchers migrate to industry, then this could undermine longer term progress through key major transitional breakthroughs.¹⁹³

Following is an overview of AI research throughout New Zealand.

GOVERNMENT INITIATIVES

The **Marsden Fund** supports a large number of projects in blue-sky research in all fields, including a number of funded projects in AI. These include a number of projects in machine learning, data mining/big data, and planning and scheduling. Projects currently include:¹⁹⁴

- An adaptive predictive system for life-long learning on data streams (University of Auckland)
- Automatic design of heuristics for dynamic arc routing problem with genetic programming (Victoria University of Wellington)
- Deep learning without the headache: computationally efficient extraction of features from data with many correlated variables (University of Waikato)

- Dimension reduction for mixed type multivariate data (Victoria University of Wellington)
- Distributed data-intensive service composition (Victoria University of Wellington)
- Genetic programming for dynamic flexible job shop scheduling (Victoria University of Wellington)
- Large-scale evolutionary feature selection for classification (Victoria University of Wellington)
- Novel decomposition techniques for multiobjective optimisation (University of Auckland)
- Pattern discovery from big medical data (Massey University).

MBIE has a range of contestable investment funds that include support for research.¹⁹⁵

These include the Endeavour Fund, Envirolink Scheme, PreSeed Accelerator Fund, National Science Challenges, Regional Research Institutes Initiative, Strategic Science Investment Fund, and the Te Pūnaha Hihiko Fund.¹⁹⁶

Funding from MBIE supports active AI research through the Endeavour and Partnership Funds totalling approximately \$37 million. Projects currently include:

- Machine learning for convective weather analysis and forecasting (Met Service)
- Machine learning to automate image based identification of species (Manaaki Whenua Landcare Research)
- Map and Zap: automated and high throughput mapping and environment friendly laser treatment control of weeds (AgResearch)
- Autonomous forest pruning and data collection of tree metrics (University of Canterbury)
- Data informed decision making and automation in orchards and vineyards (University of Auckland)
- Reconstructing real world lights and reflectance models for augmented and mixed reality (Victoria University of Wellington)
- Situated visualisation to enrich sports experience for on-site spectators (University of Otago)
- Growing NZ optometry device industry, through developing novel technology for screening blindness (University of Auckland)
- Precision Driven Health (Precision Driven Health).

Government National Science Challenges

The Government supports a number of National Science Challenges including the Science for Technological Innovation Programme (SfTI). This programme aims to develop world leading science and technology relevant to New Zealand with a focus on building enduring partnerships between researchers, business, and Māori organisations.¹⁹⁷ One key SfTI project investigates adaptive learning robots to complement the human workforce. The project has been running for two years with a vision to be at the forefront of research for highly flexible and easily adaptable robots within the next 15 to 20 years.¹⁹⁸ The interdisciplinary research programme involves robotics experts from The Robison Robotics Institute at Victoria University of Wellington, Lincoln Agritech and Scion, as well as researchers and PhD students from Auckland, Canterbury, Massey, Otago, and Waikato Universities.¹⁹⁹

Another key SfTI project involves the use of machine learning, computer vision and IoT techniques and is investigating precision farming to improve the seafood industry particularly mussel and salmon farms. The project has been running for two years and SfTI has awarded another three years to further develop intelligent systems for the seafood primary industry. The interdisciplinary programme involves AI and machine learning experts from the University of Canterbury, aquaculture experts from Cawthron Institute, and IoT experts from the University of Auckland, as well as postdoc researchers and PhD and Master's students from the University of Canterbury, Victoria University of Wellington and the University of Auckland.²⁰⁰

Regional Research Institutes

In 2015, the New Zealand Government funded four Regional Research Institute initiatives.²⁰¹ Some of this work focuses on the development of AI solutions. For example:



PlantTech based in Tauranga has identified that application of AI technologies to enhance premium crop growing industries is a research gap, particularly with respect to small and medium

business applications. Its research focuses on two key areas. Firstly, the application of deep learning for feature extraction relevant to crop growing. Secondly, developing autonomous, sensing machines that can operate in growing environments. The work of PlantTech incorporates multimodal learning, transfer learning, application parameters and technology integration. Three demonstration projects are testing the technology include kiwifruit yield management, plant wound identification and treatment to prevent botrytis, and autonomous crop management, including obstacle recognition and crop analysis. The goal is a new applied AI capability that will be uniquely capable of supporting the development of applications for premium growers.



Xerra (previously the New Zealand Centre for Space Science Technology) is using Earth observation technology, which involves satellites collecting imagery and geospatial data, along with sensors. This data is then paired with machine learning technologies to enable an understanding of human behaviour and the natural world, in ways and at a scale, not previously available. Xerra's approach can help decision making within New Zealand industry, business and Government.²⁰²

Crown Research Institutes

Several Government funded Crown Research Institutes are exploring applications of AI, including the following projects:



The Institute of Environmental Science Research (ESR) has an AI research project developing machine learning methods to streamline and automate the detection of illicit drugs when processing samples.²⁰³ From 2018 to 2023, the ESR plans to target the rapid adoption of analytic projects that can leverage their data via AI and machine learning. Funding from the Government Strategic Science Investment Fund (SSIF), including the Pioneer Fund, will support this new programme.²⁰³



Plant & Food Research are using machine learning and AI to extract information from photographs of fish. This approach allows users to identify individual fish and identify multiple phenotypic traits. With sufficient computing power every fish in a population could be recognised. This has important applications in aquaculture, breeding, fish farming, conservation, and is a less invasive and less labour intensive method than traditional track and trace approaches.



AgResearch have been researching strategies for providing tailored weed control advice. Their 2018 Annual Report identifies a next stage product incorporating AI to both deliver advice and also collect more data to improve performance.²⁰⁵ The organisation has been exploring the use of AI to analyse vast quantities of scientific research, identify causes and effects described in this literature and deduce new hypotheses. This work is profiled in the case study above.



GNS Science (previously the Institute of Geological and Nuclear Science) has been experimenting with the use of cloud computing services to speed up calculations on big data. They have used the Google Earth Engine (and associated multi-petabytes of data archived in the cloud) applying machine learning algorithms to analyse New Zealand's wetland, flood plains, forestry and agriculture. In one project, these techniques were used to map wetlands and floods, such as Otago flooding (2017) and Edgecumbe flooding (2017) to a resolution of 10m x 10m.²⁰⁶



Scion is investigating the use of robots for automating work in dangerous or harsh environments such as forestry work (See Lincoln University below).



Manaaki Whenua Landcare Research is investigating machine learning in a number of areas within their science. They are collaborating with Canterbury University (see below) to apply transfer learning approaches to automate image-based identification of species, developing machine learning based algorithms for cloud removal from aerial and satellite imagery where temporal changes matter, and have collaborated with the University of Hawaii to help with invasive plant detection from UAV footage. Future work includes using machine learning to count pollen by species from microscope slides, detection of feral goats from audio recordings, and counting Manuka flowers in images of native vegetation. Manaaki Whenua is part of the New Zealand e-Science Infrastructure (NeSI) collaboration and uses this capability to scale data processing and analytics.

UNIVERSITIES



Auckland University of Technology (AUT) has particular strength in neural computing models (for example, Professor Nikola Kasabov's group, KEDRI) and now has a Centre for AI Research, which lists over a dozen research projects on its website.²⁰⁷ These include research projects in computer vision, sports science, forensic intelligence, smart shopping, bio and health informatics. The Centre for Social Data Analytics, for example, is using predictive risk modelling and machine learning to aid in making better decisions in the child welfare sector.²⁰⁸



Lincoln University has a Complex Systems, Big Data and Informatics Initiative, which aims to 'transform the way we perceive and solve complex problems that are in the confluence of biological, agricultural, environmental, ecological, or social domains.'²⁰⁹

Lincoln University-owned Lincoln Agritech is a partner in the Robotics Spearhead project of the New Zealand National Science Challenge. In collaboration with the CRI Scion (New Zealand Forest Research Institute Limited) and other universities, the project aims to enable new systems for flexible, highly adaptable and learning robots to ensure agile and safe behaviour in workspaces with people in unknown, highly dynamic and often harsh environments. This could, for example, improve safety in the forestry sector.



Massey University has some AI researchers, in particular working on machine learning, and runs several research projects in AI.²¹⁰ For example, they are investigating how the use of simple sensors (such as motion sensors, power sensors registering the use of appliances, or contact switches on cupboard doors) can be combined with machine learning to build intelligent houses. Intelligent homes are able to take appropriate action and provide support, reassurance, assistance and comfort to its residents. The aim is to enable older people in particular, to live independently in their own homes, for longer.



The University of Auckland has the country's largest computer science department and has a substantial and strong group working in a range of areas in AI including an 'intelligent systems and informatics group.'²¹¹ The university has a number of academic staff who conduct AI research. Professor Bruce MacDonald is leading a newly funded collaborative robotics in agriculture programme, which benefits from a \$16.8 million MBIE grant. The project has two phases, firstly using sensors to learn how expert human workers do manual tasks. The second phase will deploy augmented reality to train workers to improve their skills, finally robots will be able to supplement the human workforce to overcome chronic shortages of seasonal workers.²¹² Collaborators

include the University of Waikato, Plant and Food Research and Tauranga-based RoboticsPlus.

Also at the University of Auckland, the recently appointed Professor Michael Witbrock is establishing the Broad AI Lab, conducting research at the intersection of machine learning, reasoning and natural language understanding. The lab has 'an additional focus on achieving the best social and civilizational impacts of increasingly powerful AI.'²¹³



University of Canterbury has a number of researchers working in many aspects of AI across its colleges and departments including both fundamental research and applied research in areas as diverse as biosecurity, earthquake damage prediction, climate model evaluation, social media, education and astronomy. The AI Research Group undertakes fundamental research into algorithms that mimic aspects of human intelligence while researchers are also involved in understanding the emerging philosophical and ethical context of AI in human society. Machine learning for prediction and classification is a foundation of much of UC's applied AI research. Dr Varvara Vetrova is leading an MBIE funded project in collaboration with Manaaki Whenua Landcare Research to rapidly identify biosecurity threats from photographs of pests for quick decision-making by biosecurity officers and farmers. Another application of AI is object and pattern recognition with Professor Simon Brown leading an MBIE funded project developing low power neuromorphic computer chips. These devices mimic the behaviour of neurons in the brain through nanoscale switches that act like synapses and provide human-like recognition capabilities. Likewise, Professor Richard Green's neural networks and computer vision work is enabling the automation of tasks with human-like precision in a number of domains such as agriculture, viticulture and silviculture amongst others. Dr Olivia Erdelyi, a member of the OECD Expert Group on AI (AIGO), is active in developing sustainable policies to ensure beneficial development and societal adoption of emerging technologies while proactively tackling economic, ethical, legal, societal and policy issues and implications.



The University of Otago has expertise in various areas of AI within the Departments of Computer Science and Information Science. Strengths include machine learning, computer vision, pattern recognition, autonomous agents and multi-agent systems, and evolutionary computation. Application areas include analysis of medical images and data, text recognition in historical documents, mining norms from social media, and robotics. The University is also a New Zealand focus for research on the impacts of AI on society, and on AI regulation. This research is grouped around the Centre for AI and Public Policy (CAIPP), and the New Zealand Law Foundation Centre for Law and Emerging Technologies. A key current project is the 'AI and Law in New Zealand' project, funded by the NZ Law Foundation. This project recently released a report into uses of AI in the New Zealand government, and is now researching AI and employment issues in New Zealand.



Victoria University of Wellington conducts research in a range of areas within the discipline of AI² Led by Professor Mengjie Zhang, the team has over 50 AI researchers (over 15 academic staff members, about 10 postdoctoral fellows and over 30 PhD students) including two IEEE Fellows and one Fellow of the Royal Society Te Apārangi. The team works on machine learning and data mining, computer vision and image processing, autonomous systems, planning and scheduling and combinatorial optimisation, and robotics. VUW has particular strengths in the area of evolutionary computing. The Evolutionary Computation Research Group are working on practical solutions to problems by combining machine learning with discovery techniques.²¹⁴ For example, an optimal solution to a problem may be found through interaction of the learning agent with an environment. This could be a robot navigating an environment or a data mining program determining the best stock market portfolio. The Royal Society Te Apārangi has highlighted VUW's research in evolutionary machine learning and its applications to NZ primary industry.²¹⁵



University of Waikato has a strong group in machine learning, and are the creators of the well-known Waikato Environment for Knowledge Analysis (WEKA) platform. Researchers at Waikato are exploring ways to make machine learning techniques more readily available. They are doing this by creating a set of critical tools for enabling AI and have incorporated several standard machine learning techniques into their WEKA software workbench.²¹⁶ WEKA is rated comparably to Microsoft Azure Machine Learning Studio and TensorFlow (4.5/5) and offers different functionality, such as deep learning and an ML algorithm library, but its few reviews suggest it is not as widely used as these newer competitors.²¹⁷

One rough proxy measure of research strength is the number of times research work is cited. Searching Google Scholar with a filter for ".ac.nz" and the label "Artificial Intelligence" suggests that the strongest AI research at New Zealand universities is taking place at Auckland, Otago, Victoria, Waikato, and AUT.

INDUSTRY

A number of industry players in New Zealand have active AI research programmes.

Examples of these include:

- **FaceMe** scientists are developing world leading technology with a focus on creating simple interactions without delays or complex navigation, bringing together AI, 3D technology, speech services, and video to deliver personalised and intelligent customer experiences. "FaceMe's goals for AI centre around understanding what a user is saying, understanding a user and what they are doing, understanding the environment the user is in, and as a general tool for improving the quality of the digital humans we produce," says FaceMe's Lead Scientist, Jason Catchpole.
- **New Zealand Merino and neXTgen Agri** are learning about sheep behaviour by monitoring the jaws of sheep with sensors.²¹⁸
- **Orion Health** are major contributors to the Precision Driven Health research partnership, as profiled earlier.
- **Soul Machines** continues to undertake proprietary AI research.

Where to Turn for AI Research Support in New Zealand

CallaghanInnovation

Callaghan Innovation provides a range of research and development (R&D) co-funding grants to New Zealand businesses of all sizes. The grants can help businesses get started, develop an R&D intense project, expand R&D or bring R&D-focused students on board.²¹⁹ As well as co-funding, Callaghan provides services to help businesses innovate, including science and engineering technical support, as well as providing commercial advice and connections. Callaghan Innovation’s services include scientists and technologists with expertise in IoT, data analytics, machine learning and deep learning. It also offers a data discovery service to help those new to big data and analytics explore ways to leverage their own data and a Data DIY service for those who want to build their data science capabilities.²²⁰

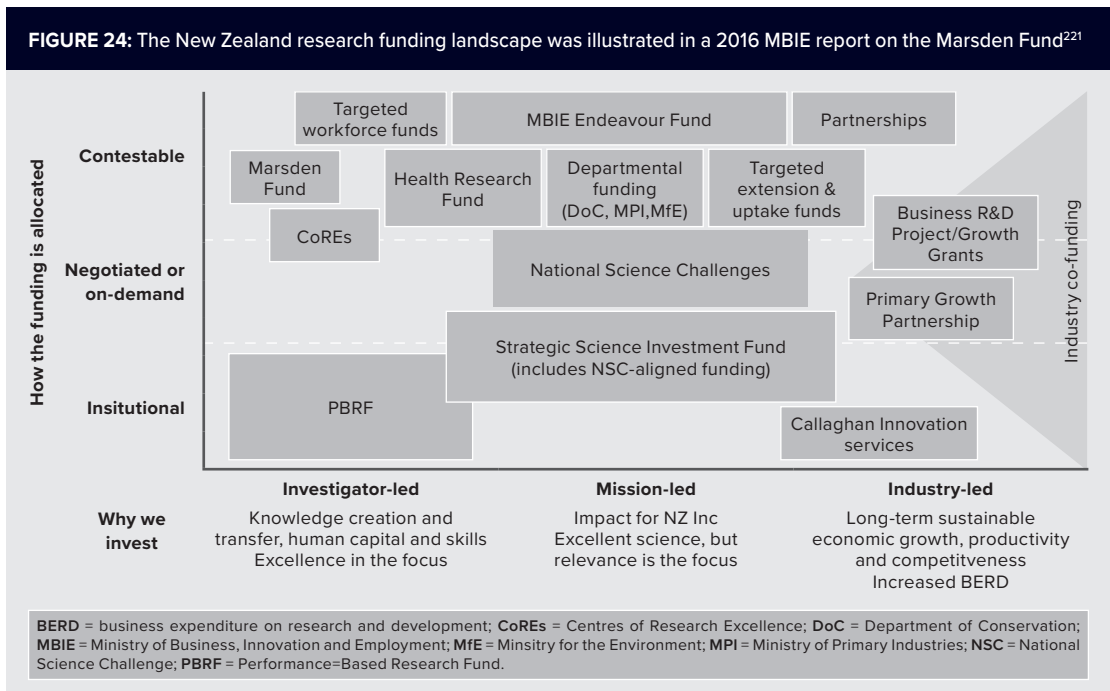


The Marsden Fund is a contestable research fund with annual application rounds. Marsden funds innovative blue skies research (where real world applications are not immediately apparent) across all disciplines. Marsden has granted \$12.2 million in funding to 24 active projects that self identify as working in AI, machine learning, advanced robotics, neural networks, computer vision or natural language processing.



The New Zealand Government National Science Challenges provide targeted research funding for projects that address priority issues for New Zealand. These challenges include the Science for Technological Innovation challenge which in 2019 is looking to fund up to 25 new proposals of up to \$200,000 each.

FIGURE 24: The New Zealand research funding landscape was illustrated in a 2016 MBIE report on the Marsden Fund²²¹





The New Zealand Government Strategic Science Investment Fund Plan 2017–2024 states that the Government will "consider investing in a platform that delivers national research capability in big data and analytics".²²² MBIE has \$49 million available over seven years for research programmes in data science and is calling for proposals for research into big data. MBIE aims to support projects that include Agritech, environment, urban planning, space industry, health sector (in particular precision health) and energy. Proposals will be assessed in July and August 2019.²²³ This funding has taken some time to distribute and the 2019 call for proposals follows an aborted funding round in 2018.

Partnerships with private enterprise are another possibility, for example MinterEllisonRuddWatts has a \$2 million joint venture with Goat Ventures to learn about AI for legal services (2017) and has attracted four PhDs.²²⁴

International not-for-profits also provide competitive funding for AI research globally. Examples include the AI Grant²²⁵ who fund small AI projects and the Open Philanthropy Project, which targets funding to AI safety projects.²²⁶



The New Zealand eScience Infrastructure (NeSI) supports AI research in New Zealand. NeSI is an unincorporated body, with investment from New Zealand universities, CRIs and MBIE. NeSI provides high performance computing, consultancy, training and data services.²²⁷



How is New Zealand Tracking with AI Research?

Despite the range of New Zealand AI research initiatives detailed here, government and business spending on R&D as a share of GDP is low in New Zealand compared to a number of other OECD countries (see Figure 13).

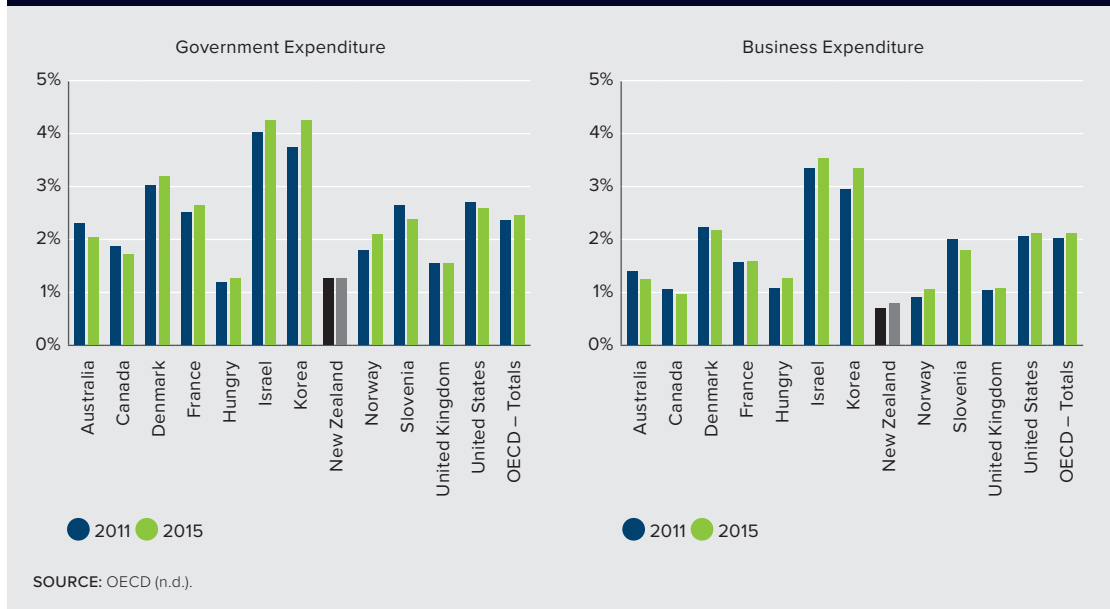
One way to increase research is to increase the number of specialty researchers. According to one New Zealand academic researcher interviewed, this comes down to the simple fact that more funding is needed, however he acknowledged that money is not always easy to source. Essentially, if there was more research funding for AI projects (including in industry), this would attract more AI researchers and postdoctoral fellows. However, international competition for such people is very high. In Canada, for example, Montreal’s ElementAI received CA\$137.5 million dollars in venture capital funding,²²⁸ Quebec has pledged CA\$100 million

dollars in research funding into AI,²²⁹ and the Vector Institute of AI was originally earmarked to receive significant funding from federal (CA\$50 million), provincial (CA\$50 million) and private (\$CA80 million) sources,²³⁰ however, the funding contribution from the Ontario Provincial Government was subsequently cut.²³¹ All of these examples are in a country whose population is only eight times that of New Zealand.

Similarly, the UK’s AI Sector Deal includes £100 million funding for 1,000 new AI PhDs, together with a new prestigious fellowship scheme for top AI researchers.

The New Zealand Government has a goal to raise research and development spending to 2 percent of GDP over 10 years. To help more businesses invest in R&D the government is implementing an R&D tax incentive which is a credit rate of 15 percent on up to \$120 million of eligible expenditure, with a minimum threshold of \$50,000 per year.²³² This Research and Development Tax Credits Bill was passed by Parliament in May 2019.²³³

FIGURE 25: New Zealand spend on R&D as a percentage of GDP compared to selected OECD countries



CASE STUDY AGRICULTURE

ROBOTICS PLUS:

Machine Learning and Robots to Solve Real World Agricultural Problems

Tauranga-based Robotics Plus has a vision of “robotics to feed the world by powering up productivity in food and fibre value chains”.²³⁴ The company has been focused from the outset on solving the growing challenges in horticulture including labour shortages, sustainability and yield security. The innovative business, which is part of the Plus Group of companies, has a number of technology products that use a combination of robotics and AI to solve real world agricultural problems. Its products include Āporo, a robotic apple packer which was launched in 2018. The Āporo, which is already in use in New Zealand and the United States markets, packs up to 120 apples per minute into trays.²³⁵ Another product is a log scaling machine that photographs and measures logs which are set for export. The product was developed with log exporter ISO Limited to increase worker safety and speed up the transport chain. The process of measuring logs was previously done by manual processes. “The machine is essentially a camera bar that takes dozens of photographs of the logs and can measure the outside edge,” Robotics Plus chief executive Matt Glenn told Stuff.co.nz. “Artificial intelligence was developed so the camera can depict the outside edge of a log and make calculations to its volume.”²³⁶ The company has other inventions in the works, including an autonomous vehicle that can drive around an orchard by itself. Functional units, like kiwifruit picking arms or pollination systems can be put on top of the vehicle.²³⁷ Robotics Plus was the anchor partner on the Multipurpose Orchard Robotics project, in collaboration with Plant and Food Research, University of Auckland, University of Waikato and 8 industry partners. The four-year

project began in 2015 with over \$10 million in funding and had 21 engineers and scientists working on the project with the aim of automating the harvesting and pollination of kiwifruit and apples. The company was named Callaghan Innovation Hi-Tech Kamupene Māori o te Tau – Māori Company of the Year.²³⁸ The company has also attracted major investment, with Yamaha Motor injecting US\$10 million into the company.^{238a} What the future holds Matt Glenn reflected that “building the future” isn’t just about investment, it also involves partnering with the right people who can provide the ability to scale. Focus has gone into increasing capability, with the team building from 15 to 50 in just 12 months, identifying manufacturing parties to enable scale as the demand for their products increases and building opportunities for their products. The Āporo robotic apple packer has now been sold into France, UK and the USA with a total of 24 packers in use. Uptake on the log scaling machine is underway with orders placed by ISO limited (an international port logistics company) with a number of units already installed into New Zealand. Expansion plans are in place into international markets. Robotics Plus has continued to focus on the design of unmanned ground vehicles and their applications with a view that they have a large global market segment for medium sized vehicles. When asked on the role of AI, Matt’s response was that AI is vital in all the systems that Robotics Plus build, including navigation and safety systems within autonomous systems, creating exciting opportunities.



RoboticsPlus



RoboticsPlus

Learn more at www.roboticsplus.co.nz

CASE STUDY SCIENCE

AGRESEARCH:

Automated Scientific Discovery

The faster scientists can find solutions to agricultural problems the more advantage the New Zealand agricultural sector gains on the world markets. However, scientists at New Zealand Crown Research Institute AgResearch recognised that there is too much technical literature in the world to keep up with, and sifting through it is a slow process. AI can accelerate the process of extracting and classifying content and then use the data to provide advice or recommendations. For example, recommending a reading list, finding testable hypotheses, or suggesting a product or service for a customer.

With little understanding of AI before embarking on their project, an AgResearch team is creating an AI solution to this problem. The solution is being produced in house and has the advantage of developing expertise without committing to a full scale investment. The findings will help inform the organisation's next research activity.

The system being developed assists scientists to pinpoint related data or facts within large amounts of literature on a particular subject. The initial field of interest is nitrogen leaching. Given the number of expert nitrogen scientists in New Zealand, this subject will be a good test of the system to see if it can find hypotheses scientists have not yet discovered.

The system uses natural language processing (NLP) and works by analysing the text of research papers to find content linkages and infer solutions. For example, in one document a statement may say, 'a affects b'. In another, 'b affects c'. The system can then infer a relationship between 'a' and 'c'.

The system will do three main things:

1. store and manage published data using research publishing houses' APIs
2. examine and extract content from inbound documents into a large textual database find relevant and contextual linkages between statements by applying NLP across the database.

The project's challenges include the lack of an existing skill set in AI, a multidisciplinary team that need to speak the same language, and understanding and accessing the literature – which is often housed



in proprietary databases – and mindset changes. For example, a switch from using the technical literature for hypothesis testing to using it for hypothesis generation requires a major shift in perspective.

According to AgResearch, the two largest benefits of using an AI driven literature discovery are:

- time saved discovering novel solutions
- helping emerging researchers come up to speed in a field.

Overall, AgResearch are exhibiting key best practice approaches to AI adoption. These include:

- learning by doing.
- building a cross-functional project team including developers, domain experts and users.involving end users from the start
- assessing the decision between buy versus build
- setting smaller goals, such as producing a usable subset of the full system.



Learn more at www.agresearch.co.nz

New Zealand AI Research Opportunities

There are a number of immediately apparent opportunities to target research into AI in New Zealand, largely based upon the strength of the underlying domain knowledge and datasets of our traditional industry sectors.

AI AND HEALTH RESEARCH

The quality and scope of big health data in New Zealand is considered very strong and includes linked data in the Statistics New Zealand Integrated Data Infrastructure (IDI), which includes other relevant information on education, tax and revenue, as well as other research databases (PREDICT/VIEW).²³⁹ If this and other data were more easily accessible to researchers, with appropriate privacy and health data constraints, then it could mean novel insights can be generated. We will examine AI and health in depth in our forthcoming *AI in the New Zealand Health Sector* report.

AI AND AGRICULTURE RESEARCH

New Zealand has a strong background of technical knowledge within the agriculture sector. We have access to unique primary sector datasets including from dairy and kiwifruit industries.²⁴⁰ The low added value on primary products, potential future shifts in global food consumption patterns and promising results in the automation of agriculture, may mean that New Zealand AI research can leverage our existing knowledge to become world leaders in agricultural AI as a service.

A Regional Research Institute, the New Zealand Centre for Space Science Technology, now known as Xerra, produces satellite geospatial datasets that could be leveraged by AI to enhance agriculture and forestry. Perhaps more importantly, the techniques have the potential to be sold as a service overseas where the agriculture and forestry industries are far larger than in New Zealand. Hyperspectral imaging using AI is also a growing field, particularly in horticulture. Callaghan Innovation is supporting industry projects to spot crop disease sooner, count yield more effectively and identify the ideal harvest time for crop with medicinal qualities. We will examine AI and agriculture in depth in our forthcoming *AI in the New Zealand Agriculture Sector* report.

AI AND CREATIVE INDUSTRY RESEARCH

New Zealand's well established multimedia industries (including animation, film and CGI) suggest that research into AI, synthetic media and digital forensics might be a fruitful avenue with potential international collaborations.

AI Research as a Career

Two New Zealand based academic AI researchers spoken to, feel that an aspect of AI research that is missing in New Zealand is clear support for a career path for emerging researchers. Funding is often ad hoc and postdoctoral researchers in particular, can struggle to find job security and progression opportunities. These local AI researchers indicated that the UK Alan Turing Institute²⁴¹ is a good example of the sort of connected, collaborative organisation that can support a range of career progression paths for AI researchers. A more holistic approach could be investigated in New Zealand.

There may be a key relationship between research funding provided at universities and the number of international graduate students attracted to New Zealand to undertake research and then enter the workforce. A University of Otago computer science lecturer notes that half his international graduate students remain in New Zealand. However, these students generally only come here in the first place if there are full scholarships available. The Government could investigate the possibility of instigating scholarships for targeted AI research degrees with this long term outcome in mind.

Another structural aspect to reflect on is the competitive nature of research funding and the competition model between universities. At present research groups from various universities submit applications to funding bodies and 'winning' proposals receive funding. This makes for a high-stakes winners-take-all situation in many cases, which leaves many junior researchers unsure of the security of their roles. One approach could be to simply distribute the funding across all AI researchers to foster collaboration and enhance funding security. Given that there are very few AI postdoctoral researchers locally, consideration could be given to supporting the role of Research Software Engineers to work with research teams to develop software. These researchers would increase productivity by publishing code and tools rather than research papers.

5. Trusted, Available Data

“The more users engage with everyday technology platforms, the more data is collected; the more data that’s is collected, the more an AI product improves. And the more competitive the AI product, the more users and investment dollars it attracts.”

STRATFOR WORLDVIEW, 2019

SUCCESSFUL AI ADOPTION IN NEW ZEALAND WILL REQUIRE AVAILABILITY OF LARGE QUANTITIES OF RELIABLE TRUSTED DATA, SUPPORTED BY DATA STANDARDS AND DATA GOVERNANCE SYSTEMS.

The world is currently experiencing a data explosion.

The connections between people around the world continue to grow exponentially. An estimated 56 percent of the world’s population (4.3 billion people) are now connected to the Internet and online services. Humanity now produces 2.5 exabytes (2,500 million gigabytes) of data daily²⁴² and every year the amount of data doubles.

Historically, digital data consisted mainly of documents and structured transactional data. More recently, unstructured media formats including emails, text messages, social media feeds, images, videos, audio, as well as web content and metadata have eclipsed traditional structured formats. Machine sensors are also gathering phenomenal quantities of data²⁴³ and within the next 10 years it is projected there will be 150 billion networked sensors.²⁴⁴ AI has existed theoretically for decades, but only this recent explosion of data has enabled the application of these technologies to advance at speed.²⁴⁵

Industry is also producing data at a phenomenal rate with innovative applications. Data may be streaming from transactions, documents, customer activities, cameras, multimedia, meta-data, fleet sensors, IoT devices and robotic processes. This data proliferation is creating novel business opportunities.

In agriculture, for example, a new array of sensors for measuring weather, moisture, water and nutrient levels, crop maturity and soil characteristics can

generate large quantities of data to help increase efficiency and predict Return on Investment (ROI) for specific crops. The advent of daily satellite photography of the earth’s surface creates a new source of planetary intelligence. Meanwhile, in manufacturing, data capture in production plants is guiding predictive maintenance and reducing down time. Sensors and monitoring software can learn when to detect subtle changes in machine behaviour or raw materials before problems occur.²⁴⁶

Advanced machine learning techniques, depend upon a large quantity of high quality data. With the right data, companies can better understand their current operations, optimise their processes and predict the future (on the assumption that the future will resemble the past). Most industries produce large amounts of potentially valuable data and many companies are digitising in some way.

Consumers are also producing a high volume of data through the use of search tools, social media and mobile technologies. However, much of the data produced is incidental to the task at hand. Incidental consumer data, also known as “data exhaust”, is the trail of data left by users during their online activity and transactions. Data exhaust has proven extremely valuable for some global business models.²⁴⁷

The fact that data has value can lead to issues encompassing rights, privacy, and ownership. These issues mean there must be appropriate governance of data use.

Data Governance

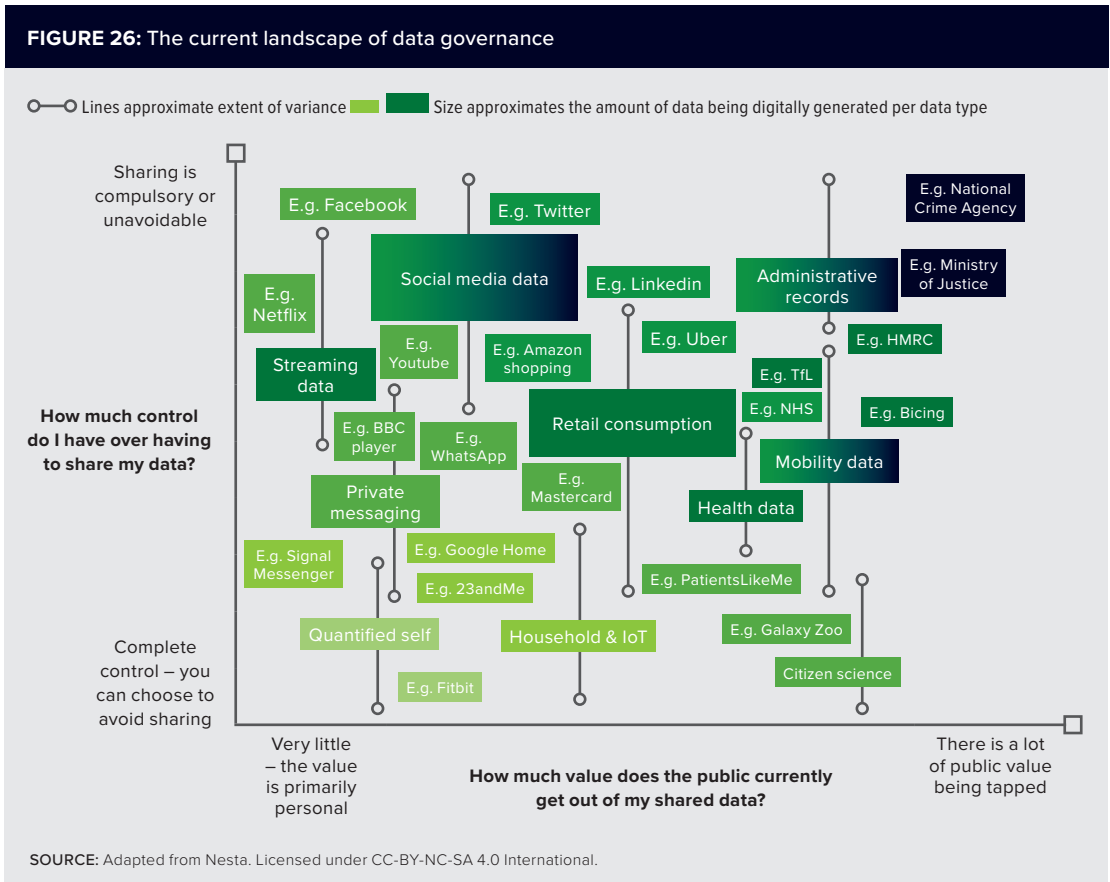
Any well-functioning AI ecosystem will need to work towards agreement on how data is governed. What data is being collected, and why? How is privacy maintained? Who owns the data? How can the data be used? These questions require serious consideration.

"Greater data use needs to be balanced with the protection of privacy rights and ethical use," as outlined in *The New Zealand Data Strategy and Roadmap* (Dec 2018).²⁴⁸ Anyone seeking to collect and use data must familiarise themselves with these issues, which are complex and nuanced.

Awareness of significant variations in ethical norms and regulations relating to data in different countries around the world is also required. For example, as described below, the European General Data Protection Regulation (GDPR) approach is very different to, say, China's rules around data use.

PERSONAL DATA RIGHTS

Nesta, a charitable foundation in the UK, has provided a useful landscape of Data Governance (see figure) which illustrates how individual control over sharing personal data maps against public value for many different types of data.



Use of data must respect the rights of those whose data is being used

Rather than making individuals responsible for every instance of use of data pertaining to them, it may be more effective to articulate and protect data rights, so that data is only shared and used in appropriate ways.

The United Nations (UN) has investigated the issue of AI and human rights. A 2018 report by the UN Special Rapporteur argues that an ethical approach to AI development, particularly in the area of content distribution, is not a replacement for respecting human rights.²⁴⁹

Among the rights to be protected, the report recognises the right to privacy. However, AI systems are currently used to infer private facts about us, which we may otherwise decline to disclose. Information such as sexual orientation, family relationships, religious views, health conditions or political affiliation can be inferred from network activity. Even if not explicitly stated, these inferences can be represented implicitly in neural nets and drive content algorithms. These features of AI systems could violate the obligation of non-discrimination.²⁵⁰

In the *MIT Technology Review*, Martin Tisne argues that a rights-based approach to data use is the preferred mechanism to ensure privacy, safety and ethical treatment of users. In his view, a bill of data rights could be developed and include rights such as:²⁵¹

- The right of the people to be secure against unreasonable surveillance.
- The right to not have one's behavior surreptitiously manipulated.
- The right not to be unfairly discriminated against on the basis of data.

DATA PRIVACY

Maintaining privacy has benefits for society over and above protections for individuals.

Indiscriminate use of data can have negative effects and these include discrimination, exploitation, imbalances of power, invasion of privacy and other undesirable societal effects.

Although it may be considered that individuals can choose to forgo their right to privacy and share their

data with whomever they choose, this does not necessarily mean that as a society we should allow unrestrained use of personal data. A key reason for this is that privacy has benefits for society at the level of society, irrespective of whether individuals wish to permit sharing of their data.²⁵² This is further discussed in Part Two of this report series. There is also an important relationship between 'decision rights' and privacy. Through decision rights individuals can choose whether to keep some piece of data private or not. However, some data gathering operations simply lay claim to these decision rights. Indeed, just by referring to the relevant data as 'data exhaust' minimises the importance of this data and subtly implies that it is permissible to gather what is talked about as if it is 'waste'.²⁵³

BIG TECH AND PRIVACY

Some major tech companies such as Facebook, Google and Apple have called for strengthening of data privacy laws.

Each company has its own view of what such laws might look like. For example, Facebook CEO, Mark Zuckerberg, argues for expanded and internationally standardised data protection rules.²⁵⁴

Microsoft's corporate vice president and deputy general counsel, Julie Brill, has written that people have a right to privacy. She has called on the US Congress to pass privacy legislation giving people control over their data. Julie Brill has also called for more accountability and transparency from tech giants, shifting the burden of data protection from the individual.

Brill says, "The prevailing opt-in/opt-out privacy model in the United States forces consumers to make a decision for every website and online service they visit. This places an unreasonable – and unworkable – burden on individuals. Strong federal privacy should not only empower consumers to control their data, it also should place accountability obligations on the companies that collect and use sensitive personal information."²⁵⁵

Harvard Professor Shoshana Zuboff argues that corporations and tech companies treat human experience as a free raw material, which can be translated into behavioural data and monetised. She calls this process, which considers data and data exhaust to there for the taking, “surveillance capitalism”. However, she asserts, this default shouldn’t be accepted without critical reflection against a society’s values and norms.²⁵⁶

Author Nick Couldry asserts in *The Costs of Connection*: “We are living through a time when the organization of capital and the configurations of power are changing dramatically because of this contemporary form of social relation [data relations]... Data colonialism is global, dominated by powerful forces in East and West, in the USA and China. The result is a world where, wherever we are connected, we are colonized by data”.²⁵⁷

Data and Privacy in New Zealand

New Zealand’s primary piece of privacy legislation is the Privacy Act 1993. Following a Law Commission review (2008-2011) and a report by the Privacy Commissioner in 2016, the Government has introduced a new Bill to reform the Privacy Act. The Bill has now had its second reading. Prior to the second reading the Justice Committee received 162 submissions on the Bill, most welcoming the Bill’s direction.

Fifteen submissions were made on the topic of automated decision making, including submissions from the Privacy Commissioner, InternetNZ, the Human Rights Commission, the Council for Civil Liberties, and the Privacy Foundation. Most of these submissions were calls for increased transparency and human oversight. There were submissions on possible bias, the need for a new offence relating to re-identification of de-identified data, and the right to erasure (or data portability). The Select Committee also proposed a change to the collection principle (IPP 3) in the Bill, which would tighten the need for transparency when collecting information, especially from children and young people.²⁵⁸

The Privacy Commissioner has also called for the Bill to include more significant penalties for breaches, including up to a \$1 million fine for organisations,²⁵⁹

saying that *‘without real and meaningful consequences for non-compliance, rogue agencies will continue to thumb their nose at the regulation, meaning responsible organisations will disproportionately bear the cost of compliance, while cowboys will ignore their obligations. If New Zealand citizens and industry are to reap the benefits of a digital economy, they need to have confidence that their regulatory regime is robust, and that their personal information will be kept safe, and used responsibly.’*²⁶⁰

In the wake of the Christchurch mosque attacks and in the current political climate, some commentators feel there could be some development on this.²⁶¹

Data and Privacy in China

New Zealand’s approach to privacy regulation can be contrasted with China. Chinese Cybersecurity Law provides a distinction between two categories of data, **personal data** and **important data**. The concept of personal data provides autonomy and control over one’s data, but important data is relevant to national security, the economy and people’s livelihoods.²⁶² Chinese Cybersecurity Law ensures sovereignty over Chinese data by mandating that personal data and important data must be stored within mainland China.²⁶³

Consent is required to collect and use personal data. A Chinese Government audit inspected a number of commercial apps resulting in a black list of Chinese apps that excessively collected personal information.²⁶⁴ The Future Today Institute articulates some concern over the Chinese approach to data, primarily the Chinese Social Credit Scoring System (see box) in conjunction with China’s investments in the Belt and Road Initiative. In one scenario, China’s systems could be exported to jurisdictions such as Turkey or Rwanda, where fundamental governance may not be in place.²⁶⁵

There are also infamous cases of data sharing, for example Sesame Credit linked up with the Baihe dating site to provide an opt-in system. Partners could then judge each other on their social credit score. However, such systems are subject to the risk of incorrect data and a few negative scores could, in theory, spark a negative outcome for individuals.

China's Social Credit Score

CHINA'S SOCIAL CREDIT SCORING SYSTEM EXPANDS THE IDEA OF FINANCIAL CREDIT CHECKS TO ALL ASPECTS OF LIFE.

Currently, it is not a single system and differing versions are operated by government and private companies. In some versions, citizens are allocated points while other versions include a credit record. Activities that negatively influence these records could include jaywalking, playing music too loud, not paying fines, stealing electricity or spending too long playing computer games. Positive influences may include being a parent or donating to charity. Data is gathered through shopping habits, facial recognition and a range of other sources. Some citizens can be blacklisted, such as for refusing to pay a fine and banned from buying plane tickets, travelling certain train lines, buying property, or taking out a loan.

In 2018 the face of Dong Mingzhu, the chairwoman of a leading air conditioner manufacturer in China was displayed on a giant Billboard sized screen in Ningbo, a port city in east China's Zhejiang province. The billboard was intended to publicly shame citizens who jaywalk by putting



Billboard displaying Dong Mingzhu in Ningbo

their photos on massive LED screens. But in this case the cameras didn't catch Mingzhu jaywalking – they identified a photo of her in an advertisement displayed on a bus!²⁶⁸

Although there are multiple systems, the data collected by private companies is expected to be aggregated by the government in the future as a permanent record. The impact of these systems is as yet unknown. As with any technology, it will depend on who is using it and why.²⁶⁹ However, a Washington Post survey found that 80 percent of Chinese people support these schemes.²⁷⁰

DATA OWNERSHIP

It may be difficult to recognise who, if anyone, owns a given dataset.

New Zealand law doesn't recognise any enforceable "ownership" rights in data per se, except where there are contractual arrangements between parties to that effect. Two parties to an agreement can agree which one "owns" any particular data. Frith Tweedie and Grace Abbot of EY Law write, however, that these ownership rights will generally only apply between the parties to the contract and not further afield. There is a risk of third parties obtaining or copying data.²⁶⁷

Commercial Data

Commercial data differs from tangible assets because possession is not restricted to one entity at a time.

The proliferation of industrial data poses several problems, including ownership. For example, Agritech (agricultural technology) providers can collect and aggregate farmers' data. However, this raises concerns regarding data ownership, farmers' privacy and commercial sensitivity.²⁷¹ If data generated on a farm is stored by an Agritech provider, do they now own the data or does the farmer still own it? Such centralised digital data represents the same phenomena as farm level data (which might be considered property of the farmer) but once it is stored online, control of the data is no longer exclusive. Although digital service providers may only use the information for the purpose for which they have been given it, additional uses may not always be preventable or even apparent. Commercial data like this differs from tangible assets because possession is not restricted to one entity at a time. Subsequently, ownership may

be better considered with respect to the rights and responsibilities of the entities accessing the data.

When users conduct searches online they generate user profile information or UPI as a byproduct. Just as UPI is the byproduct of the interaction between users and search, there is an agricultural data byproduct of the interaction between farmers and Agritech. This could lead to the potential for market manipulation.²⁷² A data aggregator (such as technology vendor) could benefit disproportionately by generating gains over and above individual contributors of the data who see gains only in their local activities.

Ultimately, building trust will be important, otherwise producers of data will not want to share it with aggregators. Certification of data hosts is one possible response to this problem. For example, the American Farm Bureau Federation has developed the Privacy and Security Principles for Farm Data, and produced an associated list of certified companies.²⁷³ New Zealand has the New Zealand Farm Data Code of Practice that requires organisations to explain how they safeguard farmer data.²⁷⁴ Another approach might be to reclassify industry data as IP or trade secrets.²⁷⁵ Yet another solution is the establishment of data trusts or data coops (see below). There will be similar data issues across many industries including, but not limited to, agriculture, health, insurance, banking and transport.

Personal Data

It is onerous for individuals to monitor who possesses their personal information, so collectors of data should be obliged to use it appropriately.

Personal data, also known as personal information, is any information that relates to a living identifiable human being. The idea of data ownership has been championed by some as a solution to issues of privacy and personal data. The premise is that users or customers who generate the data should own the data and determine how it is used. However, data ownership is abstract, partly because data flows so freely, can be shared across platforms, and because we unconsciously leave a trail of data about our behaviour with almost all use of digital technology.

We create enormous amounts of data daily, for example when we read or shop online, use a fitness

tracking device, or simply travel with our phones in our pockets. The majority of this data is created through what we do, and we are oblivious to this. We do not necessarily consent. The usage of the data we generate can become problematic, even if we 'own' it. Endless examples can be thought up that could misrepresent who you are (or represent you accurately when you don't want that) and have ramifications for your life. Location services might tag you as an alcoholic because you frequent a bar, buying fertiliser after watching a documentary on ISIS flags you as a potential terrorist, and so on.

Recent academic opinion emphasises that individuals are granted little control and oversight over how their personal data is used to draw inferences about them and that protections such as GDPR are insufficient when it comes to sensitive inferences.²⁷⁶

When your data is generated by you simply existing, it should not necessarily be up to you to own, curate and determine how this data may be used by reading an ever expanding list of consent documents and selecting a growing number of personalisation options, of which the setting choices themselves reveal information about you. The onus should be on those who decided to collect this data to use it in a way that respects human rights, including the right to privacy.

A response to those who favour data ownership as the solution is to ask, 'what is the point of owning data that should not exist in the first place?'²⁷⁷ Examples of such data could include big technology companies' inference about your likely sexuality or political views on the basis of what you read (from your web history), and real world locations you frequent (from your location history), even when you have explicitly declined to disclose your sexuality or politics to the company. In one case, retailer Target deduced from 25 data points about purchases made by a teenage girl, that she was pregnant before her family knew about it.²⁷⁸

Because the concept of ownership of personal information is confusing and potentially unhelpful the Government Chief Statistician and Privacy Commissioner have developed principles for safe and effective data use. Adherence to these principles should offer appropriate protections when organisations desire to use personal data (see 'New Zealand Action' below).

NEW ZEALAND ACTION

The New Zealand Data Strategy Roadmap sets out five considerations for safe data use. These include ethical use, privacy and control, safety and security, transparency, and rules and protections.²⁷⁹

The Data Futures Partnership (DFP) was an independent group that worked with New Zealanders to discover their expectations from digital systems.

The Partnership was funded by the New Zealand Government and aimed to ensure the creation of the right 'systems, settings and conditions to allow our data to be put to work making New Zealand a better place'. The DFP published its report, *A Path to Social Licence* in 2017. Contained in the report was a Transparent Data Use Dial, based on the views of New Zealanders. The vision was that organisations could use such a tool to convey details of what they are doing with users' data in a clear and comprehensive way. However, the Partnership ceased in August 2019.

New Zealand has some of the most comprehensive public sector data in the world. Through the Integrated Data Infrastructure (IDI), public data about individuals is linked across a range of Government

departments encompassing health, welfare and revenue, among other domains. Statistics New Zealand hosts this data, and if researchers or Government analysts wish to access this linked data they must make a special application detailing their aims and methodology.²⁸¹ However, the usefulness of the data is limited, because access must occur at a Statistics New Zealand data lab, using the tools provided by the lab.

A SOCIAL CONTRACT FOR DATA USE?

Given the diverse range of views on how data may be used, a key resource for a mature AI ecosystem will be a social contract for data use.

- Societies may need to develop a set of substantive standards for AI systems to preserve data rights and privacy, that are applicable to companies and states.
- Professional standards maybe needed to translate human rights and responsibilities into guidance for technical design.

The Government Chief Data Steward and Privacy Commissioner's Principles for the Safe and Effective use of Data and Analytics.

The following principles support safe and effective data and analytics. They will underpin the development of guidance to support government agencies on best practice for the use of data and analytics for decision-making:

- Deliver clear public benefit
- Ensure data is fit for purpose
- Focus on people
- Maintain transparency
- Understand the limitations
- Maintain human oversight

Full details of the 17 principles falling under these six themes are available from the Office of the Privacy Commissioner's website.²⁸⁰

FIGURE 27: Transparent Data Use Dial



SOURCE: New Zealand Data Futures Partnership

- Codes of ethics (such as those now adopted by most of the major AI companies) may be important but are not a substitute for recognition of human rights. An example of work in this space is the European Commission, which has launched a pilot phase to ensure that ethical guidelines for Artificial Intelligence (AI) development and use can be implemented in practice.²⁸²

The Government Chief Statistician and Privacy Commissioner's 'Principles for the Safe and Effective use of Data and Analytics'²⁸⁹ are a starting point towards a social contract for data use in New Zealand. An emerging workstream, hosted by the Department of Internal Affairs and also the Government supported research at the new Centre for AI and Public Policy at the University of Otago will certainly contribute to this work.²⁹⁰

Māori Data

ISSUES OF DATA, OWNERSHIP AND PRIVACY REQUIRE ADDITIONAL CONSIDERATIONS WHERE THEY INTERSECT WITH TE AO MĀORI. DATA PRODUCED BY MĀORI, ABOUT MĀORI, OR ABOUT THE ENVIRONMENTS MĀORI HAVE RELATIONSHIPS WITH²⁸³ IS CONSIDERED A LIVING TAONGA.

Māori have always been data gatherers and data protectors, but these practices have historically been in the context of the collective, rather than the individual.²⁸⁴ The Western worldview and its conventions around data ownership and governance may not be adequate models for thinking about, or using, Māori data. Scholars like Professor Tahu Kukutai have highlighted the need to consider Māori philosophy and worldview when working with Māori data.²⁸⁵

There is a worldwide Indigenous Data Sovereignty movement, which is concerned with the right of Indigenous peoples to govern the creation, collection, ownership and application of their data.²⁸⁶ The Indigenous Data Sovereignty movement began as a conversation between Indigenous scholars from Canada, Australia, Aotearoa New Zealand and the USA. The movement asserts the need for conversations and legislation about data to account for “the inherent and inalienable rights and interests of indigenous peoples relating to the collection, ownership and application of data about their people, lifeways and territories”.²⁸⁷

In Aotearoa New Zealand, Māori data sovereignty is the right of Māori to access, use, have governance and control over Māori data.²⁸⁸ This raises considerations for the use of data in

Government or commercial contexts, including the need to ensure data for and about Māori can be safeguarded and protected, advocating for Māori involvement in the gathering and governance of data, and requiring the quality and integrity of Māori data and their collection.

The conversation about Māori data and Māori data sovereignty is currently being stewarded in part by Te Mana Raraunga, the Māori Data Sovereignty Network. In addition, a Māori data hui was held in May of 2018, bringing together Māori data thought leaders, iwi and hapū representatives, and other specialists, to discuss the future of Māori data, define priorities, and widen the conversation around Māori data.²⁹²

Māori leaders and Government both stress the importance of honouring and reflecting the Treaty of Waitangi partnership when working with Māori data. Māori should be part of decision making processes, especially when it comes to the work of the Crown. In addition, any Māori data collected, should be used to benefit Māori, not just the individuals or entities collecting the data.

Currently, Statistics New Zealand is working in partnership with iwi and Māori organisations to understand and support their data needs, recognising the value of data for evidence based decision making and strategy. This work includes partnering with a representation of te ao Māori to co-design a data governance approach that appropriately reflects the Treaty of Waitangi relationship between Māori and the Crown. Ideally, the learnings from this process will be transferable to other agencies.

EXAMPLES OF INTERNATIONAL ACTION ON DATA RIGHTS AND OWNERSHIP

The General Data Protection Regulation (GDPR)

is an important piece of European Union (EU) legislation

on the scope and limits of how data may be used and mandates a high threshold for consent and data use. The GDPR is relevant to any organisation collecting, holding or processing data about anyone that is usually resident in the EU. A summary of the implications of the GDPR by Deloitte notes the regulation mandates simple consent and withdrawal of consent processes for permission to use personal data. It also mandates the right of data subjects to know what personal data is being used and for what purpose, and also provides the right to obtain an electronic copy of any personal data. There is also a requirement that systems ensure 'privacy by design'. This means that privacy must be addressed during development and not by retrospective amendment. Individuals will have the right to have any data about them erased and no longer disseminated.²⁹³

Since its inception, a number of cases have been brought, including recently against British Airways which faces a record £183.4 million fine after a website failure compromised the personal details of roughly 500,000 customers. The French Commission nationale de l'informatique et des libertés (National Commission for Information Technology and Civil Liberties, CNIL) fined Google €50 million for not meeting French standards for providing information to consumers about how their data is being used.²⁹⁴

However, there have been major legal challenges by tech companies that have fallen foul of the GDPR. Ireland's Data Protection Commissioner, Helen Dixon has noted that big tech firms have become more combative in this regard.²⁹⁵

However, one may question if much has changed. Since GDPR came into effect, only one percent of Microsoft users have changed their privacy settings. Similarly, less than one percent of visitors to Google's Accounts page viewed their ads settings.²⁹⁶ The phenomenon of having very low opt-out rates is well known. An example is that organ donation rates are generally high



when driver license applications have an opt-out box, and very low when the onus is to opt-in.

Practically in New Zealand, the Office of the Privacy Commissioner offers advice on GDPR for New Zealand businesses.²⁹⁷ There are also law firms providing information about compliance with GDPR. "Breaching the GDPR could result in significant financial and adverse reputational consequences," Simpson Grierson advises.²⁹⁸ Activities that might fall under GDPR include, having a website that allows payment in Euros with content in the language of the relevant EU state or monitoring website activity of individuals who are in the EU while they are in the EU. Clearly, New Zealand firms should undertake an assessment of whether their activities fall under the scope of the GDPR.

Data Portability: In the UK the recent *Unlocking Digital Competition* report made a number of recommendations aimed at ensuring competition, diversity and consumer choice in the digital marketplace. One recommendation is the creation of a Digital Markets Unit, which would ensure a code of competitive conduct and be charged with enabling greater mobility of personal user data across platforms.²⁹⁹

User Ownership of Personal Data: The Estonian e-government implementation is considered a leading example of open and transparent government, providing a foundation for trust between the Citizen and the State, and giving more control to the real 'owner' of personal data – the Citizen.³⁰⁰

Meanwhile, the cities of Amsterdam and Barcelona have been trialling a system where citizens are able to determine which projects they are happy to contribute their data to. A distributed ledger allows each participant to determine how their data is shared and used. The project, named DECODE (DEcentralised Citizen-owned Data Ecosystems) is run by a consortium of 14 organisations across the EU, and gives decision rights to the public.^{300a} Although these are valuable first steps, neither of these approaches puts the obligation on the collector of the data to use it in ways that respect the rights of users.

CASE STUDY TRAVEL AND LOGISTICS

AIR NEW ZEALAND: Optimisation and Predictive Maintenance



The enormous complexity of today's airline operations creates a set of interconnected challenges that involve people, data, predictions, and machinery. While airlines have access to vast amounts of data, using this data effectively to create and implement improvements can be difficult to do at scale. Without collaborative solutions that can extract, analyse, and build insights from complex data, airlines could struggle to adapt to future challenges.

AI can be used to overcome some of these data challenges, build predictive models that analyse airline data provide predictions and recommendations. A few examples of how AI could help improve airline operations include:

- Predictive maintenance to help airline operators avoid costly delays and cancellations by predicting when maintenance and mechanical upkeep will be required on specific components.
- Customer service improvements that use AI chatbots or sentiment analysis to answer customer questions quickly and accurately.
- Crew management that integrates predictive models with an airline operations management system, analysing complex data to help optimise scheduling for crew and ground staff.

The Air New Zealand Data Science team

At Air New Zealand, the airline's data science team works right across the business, looking for opportunities to add value in a range of areas, including customer service, engineering, flight navigation and human resources. The team uses a range of data science and machine learning techniques in its work, matching the approach to the problem or opportunity.

Examples of AI-enabled projects currently underway at the airline include:

Route optimisation: on long-haul flights pilots have a fair amount of flexibility in selecting the route they can take to their destination, with exceptions for high-traffic areas such as around airports. Air New Zealand is working on AI models to produce advice for pilots on optimal flight routes. This is based on a complex set of variables including estimated flight time, fuel efficiency and weather. It's important to note this technology would not take over the decision-making role of pilots, but provide them with more data to support their flight planning processes.

Predictive maintenance: Air New Zealand's fleet of more than 110 aircraft operates to 31 international



destinations as well as 20 destinations throughout New Zealand. All of these aircraft require regular maintenance in line with requirements set out by the Civil Aviation Authority. The airline is working with aircraft manufacturers to gather data on the lifecycle of certain components to support its scheduled maintenance programme. This data helps to inform planning processes ensuring the right number of aircraft are in the hangar at any one time.

Overcoming Challenges: **Air New Zealand's Learnings**

Like many New Zealand businesses, Air New Zealand has faced challenges in its journey to integrate AI into its day-to-day operations, including:

- **Building maturity across the organisation:** Building maturity and understanding of what data science is and how AI works at all levels and areas of the business has been a priority, with the airline's Data Science team working to build awareness and get buy-in.
- **Recruiting talent:** Enterprise AI is still a nascent field, and businesses have to balance between developing internal talent and hiring external experts. The airline has been building relationships with universities in order to foster and develop new data science talent and is working to provide a pathway to employment at Air New Zealand for young graduates.

- **Building trust:** As Air New Zealand develops new AI and data analysis applications, it has to find a balance between providing value for stakeholders and customers, and building and maintaining trust. From balancing privacy concerns with convenience, to ensuring internal teams have faith in data insights, the Data Science team is engaged in ongoing conversations about trust and social license.

AIR NEW ZEALAND 

Learn more at www.airnewzealand.co.nz

DATA TRUSTS


Data trusts can help ensure data is used safely and appropriately

Data trusts are institutions that enable the legal sharing of data for collectively-defined outcomes. The idea is that organisations would become data stewards rather than data owners, with control of data managed by appointed Trustees. Data trusts could help to reduce data abuses, enhance accountability, improve ethical standards and maximise value.³⁰¹

Data held in trusts might overcome the need for data portability. With effective data trusts providing a central repository, data could be accessed on an as needed, and as approved, basis by all parties wanting to process it. When held in trust, data is protected from unlawful or unauthorised dissemination or analysis. The central nature of the repository could also help ensure data integrity, completeness, and that data is current. Rather than having the right to transfer data across platforms, we could have the right to approve entities to access data trusts.

“We also need more research into other data access models such as data cooperatives, data commons and people led data trusts which may sometimes be more appropriate”³⁰²

JENI TENNISON – CEO, Open Data Institute



OPEN TRUSTED DATA ECOSYSTEMS

The international Open Data Institute (ODI) works with companies and governments to build an open, trustworthy data ecosystem, where people can make better decisions using data and manage any harmful impacts.³⁰³

- ODI define a data trust as ‘a legal structure that provides independent stewardship of data’. Trustees provide independent stewardship and decide what data to share and with whom, according to the purpose of the trust.³⁰⁴
- There are existing legal structures that can be applied to data trusts. Data trusts can be a good approach to data sharing when there are conflicts of interest between parties wanting to access the data. Data trusts have a number of advantages, including giving individuals and small organisations greater say about data management.
- ODI presented research on data trusts in March 2019, including three case studies. One example was a data trust that collated commercially sensitive information from electric vehicle car charging point suppliers, car club operators and smart parking sensors about the availability and use of charging points and parking spaces.
- An independent assessment of the economics of data trusts was undertaken by London Economics.³⁰⁵ Their report outlines ‘a number of market failures’ which data trusts should address and notes data trusts need to demonstrate economic benefit before government supports them. Small and medium sized enterprises (SMEs) may benefit from data trusts if barriers are reduced and a directory model is used. The report doesn’t offer strong conclusions regarding the economic benefits of data trusts.
- Regardless of the economics, governments may support data trusts due to their contribution to productivity, development of AI and increased data driven innovation.


A report commissioned and run in collaboration with the Open Data Institute as part of a project funded by the UK Government’s Office for Artificial Intelligence titled *Data trusts: legal and governance considerations* sets out the legal (British and European Union perspectives) and governance issues for Data Trusts and outlines possible structures such as incorporation and contracts or contracts solely.³⁰⁶

Another solution is to treat internet platforms as utilities, on which a range of entities can then conduct their business. The UK’s Furman report recommends that platforms open their data to competitors and such regulation might enable this internet-platform-as-utility model.³⁰⁷

DATA TRUSTS IN NEW ZEALAND

Trusted data can be a national asset and a taonga. Although, to establish this would need data contribution and sharing agreements between entities that are normally competitors. To enable data to benefit the public, there are situations where, rather than being owned, data could be held in trust for transfer across platforms (facilitating competition) or for use by multiple entities for legitimate public good. For example, health, agricultural, tourism or environmental data might be stored in a public data trust. The trust can steward the data, maintain and manage the data, and determine how the data is used, which stakeholders (such as hospitals, primary care

FIGURE 28: A new framework for data governance – Data Trusts, Data Coops and Data Commons



physicians, regulators, farmers, environmental groups, government departments, researchers or private entities) may access it and under what conditions.

Figure 28 is adapted from Nesta's work on Data Governance and Data Trusts. It shows how the concepts of Data Trusts, Data Coops and Data Commons can be applied to practical solutions for personal and commercial data and provide a framework for data governance which enables public value to be delivered while ensuring that data rights are respected.

In a New Zealand context, pilot projects to explore data trusts in sectors including Agriculture, Tourism and Health could be prioritised to provide solid underpinnings for successful AI adoption.

Data as Infrastructure

Increasingly, public data can be seen as key infrastructure of our national future, underpinning a future knowledge economy and society. In moving toward an AI enabled world society needs to invest


in collecting, curating, structuring and linking data. Data and AI models derived from it must be valued like physical infrastructure, because it will support all kinds of business and societal functions. A strategic and holistic view of data can ensure that data is consistent and interoperable, this will ensure scalability for future growth and added economic and social value. Greater integration and sharing of data will generate new insights.

New Zealand has relatively transparent and open government systems and processes, which allow us to build a data environment that is safe and open. To do this, we need to see 'data as infrastructure' and establish rules for standardising data and rules for the benefit and protection of all New Zealanders.

To succeed as critical infrastructure, data will need to be collected appropriately, according to appropriate data standards, be sufficient, fit-for-purpose and free from bias.

The New Zealand Data Strategy and Roadmap commissioned by the Government Chief Data

FIGURE 29: Data as the "Infrastructure of the Future".



CREDIT: Gavin Starks, Shared Under CC--NC-SA 4.0 International.

Steward (GCDS) provides a shared direction and plan for New Zealand’s data system.³⁰⁸ The roadmap identifies growth in global data and explains that the value of data lies in its use.

This work is still in the early days with the current phase of work focused on mapping ongoing and planned data initiatives. These initiatives will establish practical ways the Roadmap can help create value for New Zealand.³⁰⁹

DATA STANDARDS

Why do we need data standards?

The availability of good quality training data is identified as a major bottleneck in AI development.³¹⁰ This, in conjunction with the need for data to be fit-for-purpose, means we should be employing data standards for data collection and management.

The process of collecting data sources, extracting data from those sources, cleaning the data and structuring it can be done in many ways to solve a specific problem. However, if we then want to combine datasets or use data across domains, users, platforms and applications, we run into the problem of non-standardisation. We need data standards to specify requirements for the data and allow interoperable use of that data.

Novel solutions to many of New Zealand’s major social and policy challenges will involve input and analysis of data from multiple government and non-government agencies, across a range of metrics. This is particularly the case for any work taking a societal perspective. Much of the work of the Social Investment Agency falls into this category, for example.

FIGURE 30: Characteristics of data when it is treated as infrastructure.³¹¹



Anyone wanting to use AI to analyse a range of diverse datasets from multiple sources to gain novel insights will find it helpful if the data is organised, stored, labelled and accessible in standard and predictable ways.

When data is collected, it is not always clear whether it might be used by someone else at a later time. Supporting and using data standards makes work easier to reuse later and by others within and across different organisations. It should be noted, however, that technically enabling data integration is a separate issue to the choice to integrate and ethics of doing so. Just because we can integrate data doesn't mean we should or will.

The benefits of data standards include:

- cost-effectiveness in the long term
- facilitation of transparency and understanding
- enabling access
- encouraging reuse
- allowing for comparison
- providing consistency.

STANDARDS FOR DEVELOPMENT OF AI

It may also be important to adhere to standards for development of AI applications.

This is because it is often hard to determine how or why an application worked or did not work because of a lack of documentation about the development process. In other high impact industries such as medicines development, the research and development process must be documented and preserved. The Partnership on AI has a project to define best practices for transparency in machine learning life cycles, which they are calling, "Annotation and Benchmarking on Understanding and Transparency of Machine Learning Lifecycles," or ABOUT ML for short.³¹² Data science is not yet governed by widely accepted standards of quality assurance like other fields of science. However, AI development and experimentation could be documented from lab to market.³¹³

HOW STANDARDS PROLIFERATE: (SEE: A/C CHARGERS, CHARACTER ENCODINGS, INSTANT MESSAGING, ETC.)



SOURCE: XKCD (Creative Commons Attribution-NonCommercial 2.5 License.)

Overall these processes ensure quality. Although upfront investment may be needed for data standardisation, this can lower the costs of integrating with other solutions in the future.³¹⁴

What data standards could we use? There is merit in creating and adhering to a conformed cross-industry data standard.³¹⁵ Ideally the standard used would be a 'reconcile the research on factors that make for quality data'.³¹⁶

DATA STANDARDS: CDDQOS

One example of a reconciled data standard is the Conformed Dimensions of Data Quality Open Standard (CDDQOS), which identifies the following as important features of data:

- Completeness
- Accuracy
- Consistency
- Validity
- Timeliness
- Currency
- Integrity
- Accessibility
- Precision
- Lineage
- Representation.³¹⁷

DATA STANDARDS IN NEW ZEALAND

The New Zealand Government holds a substantial amount of data about New Zealanders and the Government Chief Data Steward (GCDS) and Statistics New Zealand help facilitate a cross Government approach to set standards and ensure effective data capability. Data standards for Government in New Zealand can be found at the Government open data website www.data.govt.nz. Also, some Ministries support domain specific standards such as the Ministry of Health's Health Information Standards Organisation.³¹⁸

Statistics New Zealand runs workshops where agencies can identify gaps in their data standards and experts can then set, maintain and provide guidance and support for standards. As of February 2019 this process is taking place on an ad hoc, rather than systematic basis.³¹⁹

Statistics New Zealand and the Privacy Commissioner have also published *Principles for the safe and effective use of data and analytics*.³²⁰ These principles support and guide Government agencies use of personal data (see page 100).

Some industry sectors support their own standards such as the New Zealand Farm Data Standards, which were developed to support efficient data exchange across the pastoral farming industry. The standards are designed to be used during product development and integration.³²¹ If an organisation wants to be accredited under the Farm Data Code of Practice then it should ensure its systems adhere to these data standards.

SUFFICIENT DATA

Machine learning AI generally requires large training sets of data. If insufficient training data is available, or the data does not represent what they are claimed to, then the predictive ability of supervised learning AI applications will be compromised. This can lead to erroneous and costly mistakes. But more data is not necessarily better. When data is not fit-for-purpose then misleading insights can also result.

QUALITY DATA

While data is growing at an exponential rate, quality data is required to produce quality results. Without sufficiently diverse and fully representative data it is too easy for AI systems to produce ineffective results.



Data must also be 'sufficient, fit for purpose, and free from bias', argue Liz Blythe and Zoe Sims of New Zealand law firm Russell McVeigh.³²² When we use the outputs of AI systems to support decision making, these factors are crucial. If the data is substandard (inaccurate, incomplete, or biased) then our decisions will be ill-informed. These authors note that 'this could carry with it the risk of significant legal liability, including under the Human Rights Act 1993, where relying on the outputs of an algorithm leads to discriminatory decision-making'.³²³

It can also be problematic when data points are missing. For example, The New Zealand Integrated Data Infrastructure (IDI), a research database, contains linked datasets for a host of public sector information. However, not all the datasets are complete. An example of this is the incomplete Laboratory Claims Collection. Statistics New Zealand noted in 2010 that, a 'recent change in the status of Labs, from a key component of the payment system to data warehouse, has meant that there has been a reduction in the comprehensiveness of the recording of tests'.³²⁴ Although the dataset continues to exist, it is not fully maintained and is susceptible to giving rise to erroneous inferences.

Gaps in data can be historic. For example, previously data may not have been specifically collected about people with disabilities, the elderly, Pasifika people and so on. Using data that is weak or incomplete carries with it a high risk in terms of legal liability for discrimination and other illegal practices.

Methods can sometimes be employed to try and minimise data gaps, such as using proxies like post codes, however, these approaches all carry risks and users of the data should exercise caution.

New Zealand ranks fourth for data accessibility and twelfth out of 30 countries for its overall data economy.

WHICH COUNTRIES ARE LEADING THE DATA ECONOMY?
– HARVARD BUSINESS REVIEW 2019



PUBLIC DATA ACCESSIBILITY IN NEW ZEALAND

To generate insights from the data we have, not only does it help if the data is standardised, but it must be accessible. In 2017, Statistics New Zealand was given functional leadership over data and analytics across the New Zealand Government. The agency developed a strategy and roadmap to enable greater data access and use for all New Zealanders (including communities, businesses, government and non-government organisations).³²⁵ In a world enabled by AI, data accessibility could be one of New Zealand's strengths. An article in the Harvard Business Review ranks New Zealand fourth for data accessibility and twelfth out of 30 countries for our overall data economy.³²⁶

Important new data sets are constantly emerging. The Xerra Earth Observatory is a new Regional Research Institute and compiles databases of satellite imagery and geospatial data, which is useful to a range of industries including agriculture and forestry.³²⁷ Other examples of accessible data in New Zealand include the Social Investment Agency's Data Exchange, Statistics New Zealand's Data Ventures project and the Open Government Data Programme.³²⁸



Investing in what works for better lives

SIA DATA EXCHANGE

The Social Investment Agency (SIA) is tasked with analysing data to support the social sector in New Zealand.³²⁹

This work helps to determine where investment is likely to provide social wellbeing return for New Zealanders. SIA operates a Data Exchange, which “provides a safe, secure and controlled cloud-based exchange between social sector organisations, connects social sector organisations in a way that’s easy, consistent and efficient, helps front-line staff improve service delivery by providing better information about the people they work with, and promotes standardised approaches to privacy, data management and data standards”.³³⁰

Operating at a small pilot scale since 2016, it provides transfer of data between government and non-government organisations. Information exchange is two way, privacy centric, secure and automatic in near real time.

This work needs to navigate closely alongside the Privacy Act – a 2017 inquiry found against the Ministry of Social Development for seeking to use data for possible future uses when these were not clear.³³¹

DATA VENTURES

DATA VENTURES

Statistics New Zealand runs a commercial operation called Data Ventures.

Data Ventures brings together sets of New Zealand’s non-government data through negotiation with partners, investment, and sharing of the returned value. Businesses and organisations contribute data that is not usually available. For example, de-identified telecommunications data from multiple companies to give a snapshot of population density or movements of the population, which can help with applications including emergency planning and tourism.³³²

Data Ventures brings the data together, making it confidential, complete and ready to sell to customers to enable fully informed data-driven decisions.³³³

CASE STUDY TELECOMMUNICATIONS

SPARK: Data and AI Supporting Transformation

AI Applications in Telecommunications

The telecommunications industry uses data analysis widely to streamline operations, design and implement business strategies, and build effective marketing campaigns. Telcos hold vast sets of useful data, and the amount of data is growing exponentially. As a result, old data analysis techniques and methods are no longer efficient, which has led many organisations to turn to AI solutions.

There are many uses of AI and analytics applications in the telecommunications industry, including:

- **Customer segmentation:** segmenting the market and targeting content according to each group.
- **Customer churn prevention:** keeping the customer engaged, analysing data to immediately address satisfaction issues.
- **Lifetime value prediction:** measuring, managing, and predicting the customer lifetime value (modeled on customer purchasing behaviour, activity, services used, average customer value).
- **Recommendation engines:** making predictions of customers' future product and service needs.
- **Customer sentiment analysis:** analysing data to assess positive or negative customer reactions to a service or product.

Spark Transformation

Spark has recently transformed its entire organisational structure to use Agile Methodologies. Spark is moving away from a traditional, hierarchical structure enabling the company to iterate, go to market at a faster pace and operate more leanly.

To implement AI and automation at scale, Spark followed a three-pronged strategy:

1. **Robotic process automation:** for automating repetitive use cases.
2. **Cognitive AI:** for pattern detection and unsupervised learning like retention and segmentation models.
3. **Conversational AI** using natural language processing: to communicate and understand context e.g. chatbot and email automation.

Why AI?

In the past, Spark used a lot of data analysis and other manual methods to understand its customers. However, with vast amounts of data available these days there is an opportunity to use machine learning to analyse customer attribute data and find patterns. These patterns can be used to better understand customers and reach out in meaningful ways.

AI Experience

Spark applied a use case-driven approach to develop AI solutions to solve specific business problems.

This use case driven approach needed a combination of business analytics skills, deep data analysis skills and machine learning expertise working together in an agile, collaborative team.

Spark already had in-house capability in business analytics and deep data analytics, and used subsidiary company Qrious and other external partners to provide and develop the machine learning capability in the team.

As Spark develops more AI solutions, it provides an opportunity for internal teams to gain real life experience developing AI-based solutions, and creates a virtuous cycle of developing AI talent pool inside the organisation.

Use case: Customer Intelligence using Data

Spark focused on using AI to deliver business value in three key areas which focused on the customer:

- Gain customer intelligence
- Create frictionless experience
- Implement data-driven decision making

One of the examples of using advanced AI is the churn model built at Spark to address customer retention. An end-to-end machine learning prediction model was built along with a methodology that predicts which customers are likely to churn and automatically tracks which campaigns will be more effective in retaining at risk customers. This AI solution consists of machine

learning models enabled with a live interactive dashboard to understand the context, check the accuracy of the model and track the benefits.

Often growth or retention initiatives are applied across broad customer segments. However, Spark wants to understand customers better by segmenting them into smaller groups using data and machine learning algorithms. Also this enables Spark to understand their likes and dislikes, how to have better conversations with them. This reduces bias and stereotyping of human generated segmentation.

For example, based on the churn likelihood prediction made by the AI solution, Spark released specific campaigns which were 30-40% more effective than earlier.

Spark has found that once it has a template AI solution it can re-use the solution to solve other business problems. For example, Spark tweaked a machine learning model for fibre customer segmentation to segment the entire customer base. The reusability of the AI solutions gets easier as the team get more confident with solving problems.

Buy vs Build

Spark uses plug and play AI models where it is appropriate, typically for generic solutions such as chatbots or digital assistants. However, Spark uses its own capability and resources to create models and solutions for complex problems, specific to Spark's products and services. The telco says it holds the view of a complementary relationship between buy vs build rather than a systemic "one or the other" approach.

Learnings

Spark had many challenges to overcome when starting its foray into using AI to solve business problems:

- **Getting started and building the case for change.** Discovery for finding the right opportunity to apply AI may take time. Instead of starting with a broad variety of use cases the recommendation is to start small for specific problems and scale thereafter
- **Finding people who could understand both data and business.** It may be difficult to find

such candidates in a small market like New Zealand, Spark leveraged data skills of Qrious and also developed people in the role.

- **Human Machine Collaboration.** The interface between humans and learning algorithms need to have clarity of roles. Spark provided names for the algorithms/bots and gave them role descriptions to enable better human machine interaction.
- **Scaling up is hard.** Spark found it could be difficult to scale proof-of-concepts and put them into production throughout the business. Now Spark considers how to solve scale and production issues during the POC stage.

What the Future Holds

The company recognises the value that AI-based solutions bring, particularly understanding of customer behaviour which will transform the way Spark engages with New Zealand consumers.

Some key areas of focus are:

- Understanding customer journeys and predicting failures in journeys
- Linking customer experience and capacity management, especially in networks
- Getting better in the process of delivering AI

Spark believes that there is a huge potential for AI to be used to improve customer experience, support growth and operate more efficiently.

Spark recommends to...

“Start now when it’s more of a choice in the short term, in the long term it will become a necessity.”



Learn more at www.spark.co.nz

6. Talent, Skills and Capabilities

“If your plan is for one year, plant rice. If your plan is for ten years, plant trees. If your plan is for 100 years, educate children.”

(Proverb)

NEW ZEALAND NEEDS TO ENSURE THAT IT HAS THE RIGHT SETTINGS TO NURTURE AND GROW A SKILLED, CAPABLE AND DIVERSE AI TALENT POOL AGAINST A BACKDROP OF INTERNATIONAL COMPETITION.

Organisations require skilled AI talent and organisational capability in order to extract value from large datasets. The skills and capabilities needed include technical AI practitioners (such as machine learning specialists and data scientists), managers who understand the business context and potential of AI, together with a pipeline to ensure supply of future talent. Also needed are suppliers of tools, infrastructure and expertise, as well as appropriate relationships among parties to coordinate these skills and capabilities.

Talent Demand and Supply

GLOBAL DEMAND FOR AI TALENT CURRENTLY OUTSTRIPS SUPPLY

Global demand for AI talent has doubled in the last 24 months, there are now two jobs for every AI professional, and a vast gap between talent demand and supply.³³⁴ Even though machine learning is the most rapidly emerging job in the USA, ‘lack of data science talent’ is cited by AI professionals as the second greatest challenge. The AI Forum’s report *Shaping a Future New Zealand* noted that 44 percent of survey respondents considered education a key barrier to AI adoption.³³⁵ In sum, New Zealand will need to cultivate a broader understanding of AI and train more AI practitioners to successfully increase AI adoption rates.

This talent shortfall applies to digital skills generally. The New Zealand Digital Skills Forum’s report *Digital Skills for a Digital Nation* identified the total number

of IT graduates was 5,090 in 2015, yet the tech sector alone added 14,000 new jobs in 2016. The report forecasts ongoing shortfall as demand for digital skills in New Zealand grows. However, the highest demand growth was for machine learning skills.³³⁶

In New Zealand, in 2017, there were only 1715 bachelor degree completions in IT related fields, 775 graduate certificates or diplomas, 550 honours or postgraduate diplomas, 265 Masters degrees, and 60 PhDs.³³⁷ While these degrees cover a broad range of specialties, a computer science lecturer at the University of Otago said



FIGURE 31: The number of machine learning engineers has increased ten-fold³³⁸

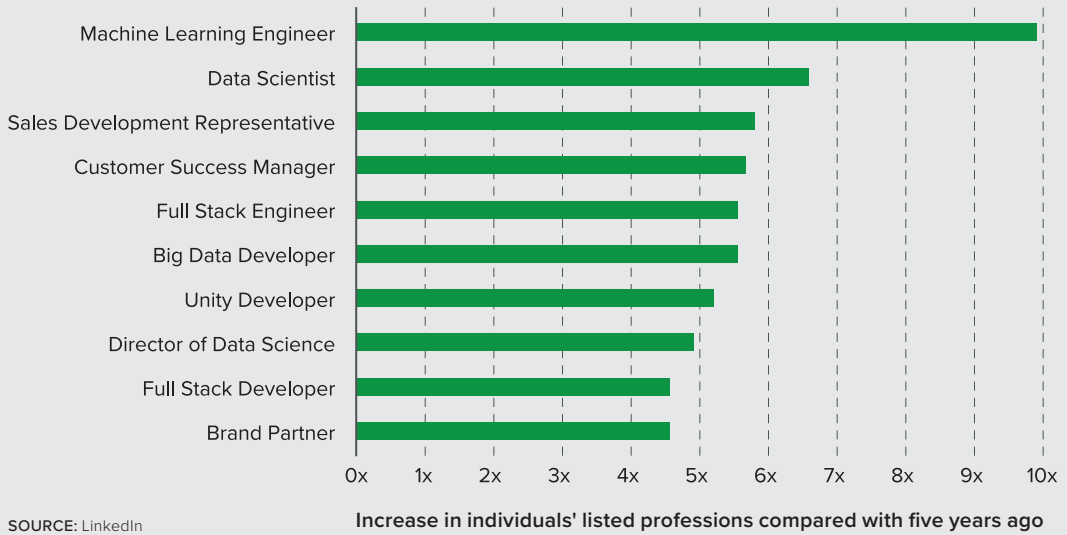



FIGURE 32: Lack of data science talent is 2nd greatest barrier³³⁹



that virtually all Otago computer science graduates will have taken some AI courses in their degree.

INDUSTRY DEMAND FOR AI TALENT COULD HAMPER RESEARCH EFFORTS AT UNIVERSITIES

There is some concern that industry, increasingly hungry for AI talent, has been draining universities of faculty members and researchers. The tech giants Google, Microsoft, Facebook, IBM and Baidu all employ a substantial number of ex-university scientists, attracted by powerful computing reserves and large salaries.³⁴⁰ If sustained, such a trend could ultimately impair the ability of universities to conduct fundamental research, a necessary component of the AI ecosystem. European academics are particularly concerned and have formed the European Laboratory for Learning and Intelligent Systems (ELLIS) in an attempt to keep talent in Europe.³⁴¹

FEW PEOPLE SELF-REPORT AI OR MACHINE LEARNING SKILLS IN NEW ZEALAND

There is not concrete data about practitioners with machine learning experience in NZ. The 2019 *AI Global Talent Report*³⁴² found there are 36,500 people worldwide who self-report AI and machine learning skills on their LinkedIn profile (up 66 percent from 2018). A total of 22,400 unique individuals presented at a set of 21 top AI conferences in 2019 (up 19 percent from 2018). Results show that 124 people in New Zealand self-reported AI or machine learning skills and just five presented at these conferences. However the methodology of this study is coarse and relies on self-reported skills and a selected set of conferences. Also, for almost all countries the number of LinkedIn profiles vastly outnumbered the number of conference presenters, however in China the reverse pattern is seen. This suggests that there is a large unidentified talent pool in China who do not have LinkedIn profiles.

In New Zealand Universities, a tally of faculty only (lecturer or above), revealed approximately 85 faculty members undertaking research into AI and related fields. There are also several companies in New Zealand with PhD level staff doing AI related work (for example, as part of a joint venture between MinterEllisonRuddWatts and Goat Ventures to research AI for legal applications³⁴³), and therefore

the Global Talent Report likely underestimates the number of capable people in New Zealand.

THERE IS A LACK OF DIVERSITY IN THE AI TALENT POOL

This situation just described has set the context for strong competition for AI talent.³⁴⁴ But overall shortages are not the only obstacle in the AI talent market. A lack of diversity in the talent pool is also problematic. Few machine learning professionals and researchers globally are women and many other minorities are underrepresented. Evidence for this comes from a number of reports.

- Research by Nesta (2019) found only 14 percent of AI researchers in the UK are women.³⁴⁵
- The AI Global Talent Report (2018) found only 18 percent of the authors of conference papers were women.³⁴⁶
- The World Economic Forum (2018) found only 22 percent of AI professionals on LinkedIn were women.³⁴⁷
- Element AI (2017) found that only 12 percent of authors who contributed work at three top machine learning conferences were women.³⁴⁸

This lack of diversity may in part explain some of the failure cases in AI discussed above. With Joysy John, director of education at research and innovation foundation Nesta, telling the Telegraph that, “Future technology will not be able to meet the needs of a diverse population if it is being shaped by a small section of society with a singular worldview.”³⁴⁹

INCREASING EDUCATION OPPORTUNITIES AND FACILITATING A TALENT ECOSYSTEM MAY HELP EASE THE AI TALENT BOTTLENECK

As a response to this global talent shortage a number of large international tech companies have been priming the talent pipeline by offering free courses online. For example, Microsoft has created an ‘AI School’.³⁵⁰ These free courses are something that New Zealand talent could utilise.

Deloitte’s 2019 Tech Trends report suggests that with the current system settings only big companies with vast resources will be able to keep industry

leading AI talent on their payroll. Deloitte argues that retaining employees full time will become dated and we will see a diverse gig talent ecosystem.³⁵¹

However, Andrew Ng the co-founder of Google Brain has a different view and lays this out in his 2018 AI Transformation Playbook. Ng shows how large organisations can prepare themselves to adopt and deploy AI. This includes an initial phase where external talent is leveraged to sustain pilot projects, but then an AI team is built internally.³⁵²

ANDREW NG'S 5-PHASES IN HIS AI PLAYBOOK³⁵³

1. Execute pilot projects to gain momentum
2. Build an in-house AI team
3. Provide broad AI training
4. Develop an AI strategy
5. Develop internal and external communications

Finally, despite the potential for AI to help solve the grand challenges that society faces, the technology and financial services sector is absorbing 60 percent of AI talent globally.³⁵⁴ New

Zealand may need to find innovative ways to ensure appropriate distribution to sectors in need.

It is possible that with time the skills bottleneck will ease. This is because there is a trend to greater usability in AI development tools. For example, the open source platform TensorFlow³⁵⁶ and similar tools allow users to write code at their preferred level of abstraction. At the extreme, Lobe³⁵⁷ has produced a simple visual drag and drop interface that allows people with no deep learning experience to train AI applications in machine vision, with no need to download anything. The company plans to add more functionality in the future.³⁵⁸ The fundamentals of AI should soon be available to anyone in just the same way that drag and drop web page builders are commonplace. Being proficient in JavaScript might give a user endless flexibility in creating exact web design and content, but using WordPress allows a beginner to do a fairly decent job. The same is likely to be true for AI applications at some point in the future. However, for now, all available AI talent is being absorbed by the market.

NEW ZEALAND AI TALENT SHORTAGE?

Greg Cross of New Zealand AI company Soul Machines feels that New Zealand is running out of the talent needed to grow the AI industry.³⁵⁵ Cross told PwC that the job market constantly evolves and the skills people need to find work change. Soul Machines currently has about 120 staff in Auckland,

Melbourne, San Francisco, Los Angeles, New York and London and recently Tokyo.

The company has doubled in size in the last 12 months. Soul Machines now has 26 staff in the US and continues to grow its footprint internationally.



The fast growing Soul Machines team.

THE STRUGGLE FOR SOME NEW ZEALAND ORGANISATIONS TO ACQUIRE AI TALENT

Members of the AI Forum's Growing the AI Talent Pool working group have discussed the AI talent shortage with New Zealand employers and recruiters.

The employers' viewpoint: Although the working group has not completed a formal survey of employers, anecdotally New Zealand organisations have found smaller New Zealand cities may be seen as a lifestyle option for overseas technical AI candidates, but New Zealand candidates don't appear interested in roles outside Auckland or Wellington. There are difficulties attracting talent from academia too. Even with salaries of \$110-130,000 offered at scientist level and \$140-150,000 at Principal Researcher level, candidates see a move out of academia, or outside the main urban centres, as a career stagnating step, with no clear pathway for progression. This is particularly difficult for new organisations and some have secured external recruitment services because online advertising is not generating suitable applications. For these reasons, recruitment may well be an ongoing obstacle to development of successful and profitable agricultural (and other regional) AI capability in New Zealand.

Recruitment agencies viewpoint: Recruiting for AI candidates is relatively new. Also, some hiring organisations themselves are new and unproven businesses. Coupled with a desire, in many cases, to recruit senior AI scientists, this has made recruitment

a challenge. There is a large talent pool in academia, however encouraging these candidates to accept roles is challenging because academic AI researchers are used to certain institutional structures, payscales, progression pathways, conferences and autonomy. It may be necessary to recreate some of these familiar structures to entice academics. Possible solutions include engaging a candidate search agency to avoid treading on toes in a small market, considering part time roles to help candidates transition, presenting candidates with a compelling vision for the organisation, and establishing a visible, tangible business before attempting to recruit senior candidates.

Developers' viewpoint: The interview experience is key to landing top talent. According to the 2019 HackerRank Survey of developers, applicants were most likely to be turned off by employers 'who don't provide enough clarity around roles or where they will be placed.' Also, 49 percent of developers said that a lack of alignment between their values and those of the employer was a turn-off. Finally, 14 percent looked for diversity on the interview panel.³⁵⁹

In New Zealand, anecdotal evidence suggests there is a mismatch of expectations. Graduates report being frustrated they can't gain roles due to lack of experience, but they can't gain experience without an entry level opportunity.



PHOTO CREDIT: David Straight

Growing AI Talent in New Zealand

Sources of AI talent in New Zealand include school students who follow through to tertiary education, industry training schemes, in-house training, self-taught developers, and a supply of international talent through immigration.

IMMIGRATION

Overseas talent may be needed in the short term, but immigration is unlikely to be a long term solution to New Zealand's AI talent needs.

Immigration is one solution to the current shortfall of AI talent in New Zealand, however, in-demand AI and data science skills including ML are not explicitly included in the long term skills shortage list.

Overseas workers wishing to work in New Zealand can apply for a temporary work visa or seek permanent residence under the Skilled Migrant Category (SMC).³⁶⁰ The process for receiving a temporary work visa is straightforward for workers with high demand skills. By contrast, permanent residence is more expensive and takes longer; 80 percent of successful applicants are already in New Zealand, having already lived here under temporary visa status.

The Essential Skills pathway³⁶¹ is the most obvious route for an AI practitioner to secure a temporary work visa. Eligibility depends on meeting one of two key criteria:

1. having the specified work experience, qualifications and (if relevant) occupational registration to work in an occupation included in the Long Term Skill Shortage List (LTSL)
2. an offer for work in that occupation, or being supported by an employer who is unable to recruit local workers for a particular vacancy.

The LTSL includes a wide range of degree level ICT, electronics and telecommunication occupations that may currently be broad enough to encompass valuable AI skills. This includes degrees focusing on computational intelligence, information science, information systems, systems design and the like. However, at present, the need for particular AI and data science skills such as machine learning and machine vision is not explicitly signalled.

The New Zealand Government also runs an Immediate Skills Shortage List but AI expertise is currently not included. Some ICT occupations are listed (cableers, cable jointers, telecommunications technicians), but these occupations do not signal the level of skill and expertise required to attract AI talent.

Overall, the current immigration rules may be sufficiently permissive for New Zealand firms to bring individuals with AI skills into New Zealand via temporary work visas and such individuals are likely to be well placed to satisfy the Skilled Migrant Category requirements for permanent residence. However, the need to be an approved employer and the time lag from application to visa approval still hamper some smaller New Zealand companies trying to move quickly.

Consulting with Immigration New Zealand on the content and wording of the Skills Shortage Lists during the current or next update could raise the profile of AI opportunities in New Zealand.



SKILLS DEVELOPMENT

A number of high income economies are investing heavily in AI education

Around the world, countries are investing in AI education. The UK has announced £200 million to fund 1000 PhD students in AI. The research and training will take place at 16 Research and Innovation Centres for AI Doctoral Training.³⁶² South Korea is investing US\$2 billion in AI research and development and plans to train 1370 AI talents by 2022, which will include 350 key researchers, along with 4,500 domestic AI scholarships. Short term talent gaps will be alleviated by funding six month intensive training courses for 600 people by 2021.^{362a} The US aims to provide US\$200 million in grant funding per year to address shortages of STEM teachers and also expand access to high quality computer science and STEM education.³⁶³

It is not just larger economies like the UK, South Korea and the US that are investing heavily in AI talent. Finland is aiming to be a world leader in practical applications of AI technology and has started to train one percent of its population in the fundamentals of AI with an entry level 'elements of AI' course.³⁶⁴ The course is targeted at working professional people, particularly in business and government, but it is free and anyone can take it. To prove the point, Finland is offering the course to inmates in the country's prison system.³⁶⁵

Singapore has also developed training courses: AI for Kids, AI for Students and AI for Everyone.^{365a}

FIGURE 33: Elements of AI course created in Finland and freely available worldwide



In comparison, New Zealand is behind in providing resources specifically targeted to AI skills development and has yet to develop a programme of systematised AI education. Policymakers need to ensure that any AI curriculum is not overly specific because technologies and the needs of industry continue to change rapidly. Any focus on developing AI talent through a pipeline involving schools must keep this in mind. The focus in schools may be on strengthening basic skills in mathematics and coding, with classroom examples of AI, such as programmable robots, or examples of AI use across diverse industries to support core content.

NEW ZEALAND SECONDARY EDUCATION

Curriculum development and teaching should ensure school students develop foundational knowledge so that further AI education remains an option

The *Digital Skills for a Digital Nation* report³⁶⁶ concluded that machine learning skills are desperately needed both in New Zealand and globally. While New Zealand needs to increase the total number of students entering the technology field, a dramatic increase in the number of people with digital skills entering the workforce is also required. The report also identified problem solving and critical thinking skills as the most important soft skills for a future workforce.

If New Zealand wants to grow its own AI talent then we need to increase interest in the fundamentals of mathematics, computer science, critical and computational thinking at school. Despite a range of approaches being tried to boost the number of students entering computing heavy courses, a shortfall remains.

FIGURE 34: The numbers game: broad interest in AI fundamentals at school will produce AI professionals

The Numbers Game



Everyone at school



Those that take STEM at tertiary



Those that graduate with ML/AI Skills



The New Zealand Digital Technologies/Hangarau Matahiko curriculum roll-out is an opportunity

To tackle this issue, the Ministry of Education has rolled out the Digital Technologies/Hangarau Matahiko Curriculum. This curriculum refresh, to be implemented nationwide by 2020, emphasises two key domains of technology thinking: computational thinking for digital technologies, and designing and developing digital outcomes.³⁶⁷

'The digital curriculum is about teaching children how to design their own digital solutions and become creators of, not just users of, digital technologies, to prepare them for the modern workforce',³⁶⁸ says Minister of Education Chris Hipkins.

EDMUND HILLARY FELLOWSHIP

The Edmund Hillary Fellowship and associated Global Impact Visa is an effort to attract and retain a community of important entrepreneurs, investors and changemakers in New Zealand to help solve key global problems.

There are 100 Fellowships per year and following an application process fellows are expected to contribute towards New Zealand entrepreneurship communities, including working on, or investing in, New Zealand initiatives for widespread global impact. The successful Fellows are eligible for a Global Impact Visa, which is a three year open work visa. This enables recipients to work and live in New Zealand. Visas are also available for immediate family members and there is a path to permanent residence.³⁶⁹ Key talent and leaders in AI maybe attracted through these Fellowships.

However, immigration may not be a sustainable solution in the face of high global demand for skilled workers. New Zealand has tried to sell itself as a highly desirable destination to live and work, rather than try to compete with international salaries and giant organisations. However, this is unlikely to be a strategy that will succeed beyond the segment of people seeking lifestyle over income.

The only specific mention of AI in the new curriculum is Progress Outcome 8 for year 13 students. This specifies, 'in authentic contexts and taking account of end-users, students evaluate concepts in digital technologies (e.g., formal languages, network communication protocols, artificial intelligence, graphics and visual computing, big data, social algorithms) in relation to how key mechanisms underpin them and how they are applied in different scenarios when developing real world applications'.³⁷⁰

The Level 2 NCEA Achievement Standard 91898 'Demonstrate understanding of a computer science concept' lists AI as a possible area for focus. However a popular resource used in New Zealand by teachers and students, the *Computer Science Field Guide*³⁷¹ has only one short section on chatbots.

Recognising the opportunity to help support the new curriculum, The AI Forum's working group on growing AI talent is aiming to expand the Field Guide's content in conjunction with the University of Canterbury. However, a lot more can be done.

Barriers to Teaching AI in Schools

Teaching AI goes beyond the skills needed to use digital technology because AI is heavily dependent on conceptual understanding and foundational mathematics and statistics. This along with some other key issues means that there are a number of barriers to teaching AI in schools.

According to the Principals' Federation, only 4000 of 100,000 teachers in New Zealand have the skills to teach the technology curriculum.³⁷² This suggests even fewer could introduce AI concepts to the classroom. IBM developer Dale Lane says the most critical aspect of AI education may be helping teachers to improve their own skills and educate children more effectively.³⁷³

Teaching AI in schools will require knowledgeable educators

The New Zealand Government has committed \$40 million to support teachers to understand the new curriculum and teach its content by 2020.³⁷⁴ However, the Teachers' Union fears this falls far short of what is needed.³⁷⁵ A scarcity of teachers with a grasp of technology in general, and AI in particular, is becoming

“There is consensus that teachers are not replaceable. Instead they will be empowered to improve the outcomes for their learners and will also see their manual and menial work reduce through the use of AI.”

AI IN EDUCATION WHITE PAPER
– X4 Consulting³⁷⁶



a serious problem and there may be need for far closer collaboration between tech firms and educators. For example, TechWomen recently hosted the ShadowTech Teachers event in Auckland. The pilot provided 35 digital technology high school teachers with corporate experience. They were able to make industry connects and take their learning back to the classroom.

As well as Government supported programmes to help transition teachers to teach the new curriculum, there are also private solutions in play. Companies such as New Era Technology provide digital solutions and support to teachers and have a specific focus on supporting the curriculum 2020 implementation goals.³⁷⁷

School teachers are unlikely to need to be AI specialists, but a conceptual understanding of what AI is, and that it is likely to be applicable to virtually all future careers is desirable, so that learning can occur in a context that appropriately emphasises these technologies.

SHADOW TECH

SHADOWTECH BY TECHWOMEN

The ShadowTech programme is an initiative from TechWomen that connects girls studying at high school with women working in tech.

Students get to spend a day with a mentor experiencing what it might be like to have a career as an IT professional. The intent is that more girls will then choose to pursue technology at secondary and tertiary level. This is both an example of a way to increase the number of school students interested in studying technology and also an attempt to increase diversity in the STEM workforce.³⁷⁸ In 2018, ShadowTech had over 600 students, 300 mentors and 100 organisations participate in six cities. In 2019, the Ministry of Youth stopped funding this initiative. A crowd sourcing campaign is open and it will be important to support these important diversity initiatives moving forward.



AI FOR EDUCATION

As well as teaching about AI, schools can use AI to support students.

When it comes to learning maths, software company Jaipuna and the Open Polytechnic's School Strategy division have collaborated to produce Amy, an AI tutor to help students advance their math skills. Amy identifies when students make an error and tailors feedback to keep them on track. Schools that sign up for Amy can now provide a maths tutor to students on demand 24/7.³⁷⁹

Teaching AI in schools does not require new technology

It is not necessarily new technology that is needed in classrooms, but an understanding of foundational concepts. There is a current trend towards more frequent career changes and the need for more generalist skill sets. With such rich availability of free online training materials and tutorials, many workers are now picking up skills in ad hoc fashion at the time and point of need. Moving forward, it is likely that generalist skills such as critical thinking, problem solving and communication will become more and more important.

Also, given parental concerns over screen time, AI concepts when introduced early in the curriculum need to be introduced in age-appropriate ways. A range of approaches exist to support this including a children's book explaining AI,³⁸⁰ various 'unplugged' activities³⁸¹ and lessons that involve basic coding.³⁸²

Students don't need to leave high school with advanced skills in deep learning or neural networks. University of Auckland robotics professor Bruce MacDonald says that the University doesn't even teach deep learning at undergraduate level, but it is widely used by postgrads.³⁸³ This further suggests that in schools, there will need to be a focus on foundational mathematical and conceptual knowledge to generate an interest in AI.

Generating interest in AI as a field of study while students are in high school will be crucial because AI will become a key feature of many careers, for example, law, medicine, engineering, finance and economics. Many of the careers that will feature AI in the future may not traditionally be seen as algorithm-based.

“New Zealand can join other countries by rapidly creating free online course materials and communitites for teachers, students and the general public.”

A NEW ZEALAND PRINCIPAL'S VIEWPOINT

"With respect to AI, what schools need is advice and guidance from industry experts. We don't have enough expert knowledge in this area, or do not know how to access it, or access the right experts. We need industry perspectives on unpacking the sorts of things we should/could be doing in schools (right from Primary) to educate our students and open their minds to the world of AI. We also need to work with industry to map out a strategic intent and then some grass roots work and actual resourcing. The Digital Technologies Curriculum was landed on us without sufficient meat in the sandwich, including resourcing, content availability, and there is a risk that could be already out-of-date because it is not compulsory until 2020 and doesn't encompass AI. Our school has chosen to undertake professional development with a professional learning facilitator, but these facilitators not necessarily connected to industry experts. We need funding and resources so the necessary professional development is not merely ad hoc."

– ANDREW KING, Oropi School Principal.

These thoughts resonate with what is happening internationally where in China there is already strategic alignment, with AI texts targeting preschool³⁸⁴ through high school students.³⁸⁵ Meanwhile, the USA's the Association for the Advancement of Artificial Intelligence (AAAI) announced a national guideline initiative for the K-12 school years to define what students in each grade should know about AI, machine learning, and robotics.³⁸⁶

NEW ZEALAND TERTIARY EDUCATION

Efforts should focus on attracting diverse students to the tertiary AI courses that already exist

New Zealand universities currently offer a wide range of courses specialising in AI. Every university in New Zealand has a number of undergraduate and postgraduate technical courses on AI and related subjects. This includes approximately 23 undergraduate courses on AI, where most commonly AI courses are taught at third year level due to the large amount of requisite knowledge (e.g. programming and mathematics). In addition, there are approximately 40 postgraduate courses on AI at New Zealand universities.

It appears that there are sufficient courses on AI available in New Zealand, but attracting students is the issue. As noted earlier, ensuring diversity of students studying AI will also be important. For example women make up less than 20 percent of enrolments in computer science.³⁸⁷ Few students and a lack of diversity underscores the need to provide exposure to career relevant AI concepts and use cases with wide appeal while students are at secondary school.

Encouraging more students to enrol in AI related courses at New Zealand tertiary institutions is one way to increase the local talent pool. However, students may legitimately ask, 'why should I study at a local tertiary institution at all?' This is because there is a veritable explosion of freely available, industry-led and foreign prestigious institution-led courses on AI available online. However, a tertiary course is much more than providing resources, it's also about support, interaction, and assessment. Free online courses don't offer these, and they are therefore more suitable for those who do not require any support, or can source support elsewhere. The many online resources can be of value for continuing education, and a proportion of students can teach themselves AI based on these resources alone.

DIVERSITY

The AI Forum's working group on Growing the AI Talent Pool recently carried out an employer survey in New Zealand (publication pending). One key finding was that while there appears to be relatively good ethnic diversity generally, there are very few women, Māori and Pasifika currently in the AI talent pool in NZ. Although not specific to AI, NZTech recognises the importance of diversity in tech and is allocating ten percent of its membership fees to address the issue NZTech has created a sub-Board to oversee and govern diversity and inclusion initiatives.³⁸⁸ Other initiatives include the emerging tech group, Pasifika in IT, who aim to encourage more Pacific youth into tech. Initiatives such as 'Tahi, rua, toru, tech',³⁸⁹ T3W, Te Tech Tribe and Hack Tairawhiti should also help increase interest in tech as a viable future for our tamariki. To

THE IMPORTANCE OF DIVERSITY

Liz Blythe and Zoe Sims of Russell McVeigh note the popular theme of AI's 'white guy' problem.

This is where a lack of developer diversity leads to algorithms being coded to reflect the dominant group's particular unconscious biases. This can potentially lead to discrimination when applications are deployed. Diversity in development and due diligence of algorithm performance should both contribute to reliable results, and importantly in a transparent way and without bias.

It may be that organisations need to engage external advisors to help assess algorithms and perform algorithmic audits. Companies with a New Zealand presence such as Deloitte and Accenture provide services in this area.³⁹⁰

“While there appears to be relatively good ethnic diversity generally, there are very few women, Māori and Pasifika currently in the AI talent pool in NZ.”

AI FORUM – Growing the Talent Pool Working Group

realise the benefits of diversity, tech companies could also consider recruiting from non-traditional sources, or using blind CV assessment in recruitment processes

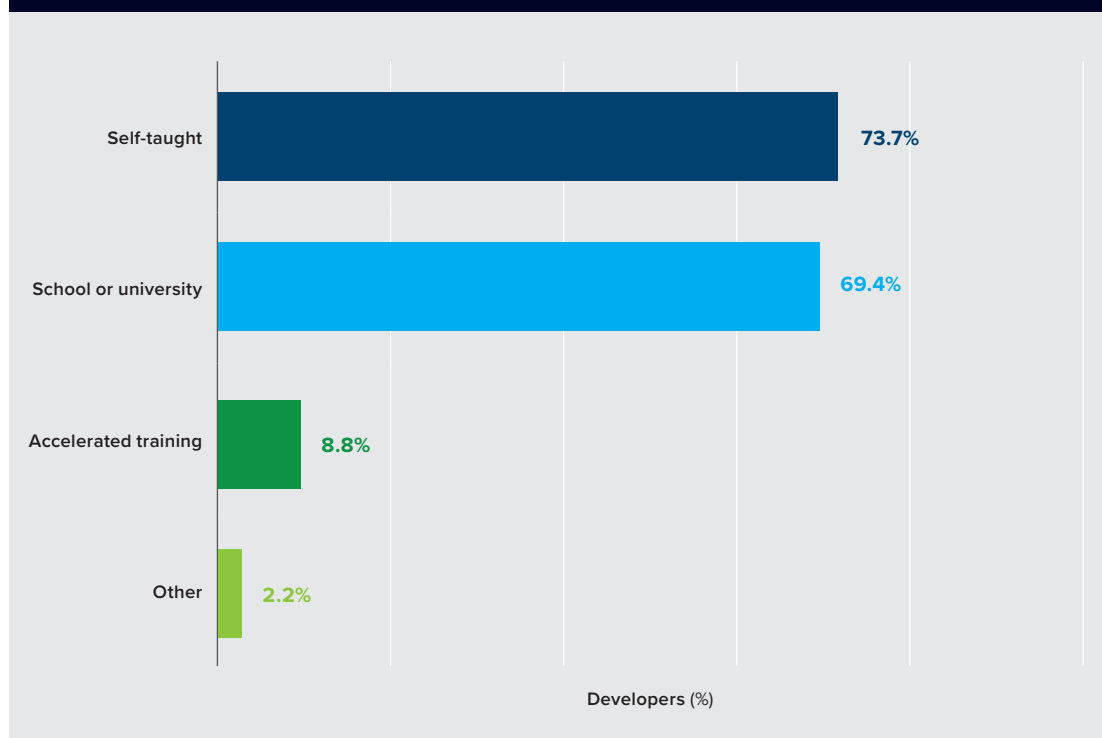
SELF-TAUGHT DEVELOPERS

According to the HackerRank 2018 and 2019 developer skills reports,³⁹¹ although 67 percent of developers have computer science degrees, approximately 74 percent said they were at least partially self-taught. The survey was not specific to AI/machine learning, however, and it has been suggested by the founder and CTO of Get Cloud Ready Consulting, Janakiram MSV, that barriers such as the intensity of mathematics required, the need to understand data analysis, as well as the explosion in number of learning resources, all contribute to the difficulties in self-teaching AI and machine learning.³⁹²

“Nearly 50% of the surveyed developers believe that the humans creating AI should be responsible for considering the ramifications of the technology. Not the bosses. Not the middle managers. The coders”.

IBM'S EVERYDAY ETHICS FOR AI³⁹⁴

FIGURE 35: How did you learn to code? Hacker Rank Skills Report 2019³⁹³



NON-TECHNICAL SKILLS

Soft skills for developers are important, but there is also a need for AI policy specialists and ethicists

For safe, acceptable and effective deployment of AI, we need a pool of talent that can do more than build technical solutions. Non-technical, or 'soft' skills including critical thinking, communication, ethics, professionalism, strategy, teamwork, policy, and the social impact of AI are needed too.

With any technology that affects society at scale there is an obligation to develop it responsibly. Previous examples such as the industrial revolution which resulted in exploitation of factory labour, or fossil fuel energy use which has caused potentially catastrophic climate change, show that major new technologies can cause damaging and unforeseen effects. Currently, we have the chance at the dawn of an AI explosion and Fourth Industrial Revolution³⁹⁵ to avoid such pitfalls.

Often generalist technical programmes of study do not have an embedded programme of soft skills. Nor are there professional certification requirements for a lot of these skilled AI workers.



This is in stark contrast to other professional disciplines such as accounting, law, nursing, teaching, medicine and others where ethics and professional development are core training components.

That said, according to HackerRank's 2019 Developer Survey,³⁹⁶ 58 percent of developers took action in response to concerns about their company in 2018. This included concerns about privacy misuse and censorship. Nearly 40 percent of developers took their concerns to management and 23 percent quit their job or began looking for another position due to their concerns. This suggests that alignment of values between organisations and staff is likely to be critical to the attraction and retention of AI talent.³⁹⁷

Although many computer and information science degrees may include material on ethics, privacy, and so on by embedding it in other courses, New Zealand as yet has no dedicated AI soft skills course. At least one New Zealand University is in the early stages of investigating with industry and potential students the value of non-STEM postgraduate AI qualifications and industry micro-credentials. The purpose of any such courses would be to develop AI leaders, managers, consultants and directors rather than black box technical developers.

Soft skill courses like these, including courses on AI for a non-technical audience are important because directors, CEOs, managers, politicians and other decision makers need to understand what problems AI can be applied to and how they can steer their organisations in a direction that leverages these technologies for greater benefit rather than exploitation. This will involve top-down vision as well as bottom-up technical solutions.

There is no need to train all engineers to be ethicists, but engineers and data scientists should be exposed to enough soft skills concepts that they can recognise when expert ethical and policy analysis is required and what that looks like. This would empower technical experts to responsibly consult with or hire the appropriate talent when needed. These diverse experts can then work with each other, rather than replace each other.

7. AI, Ethics, Regulation and Society

“Effective regulation and governance of AI technologies will require involvement of, and work by, all thought-leaders and decision makers and will need to include the participation of the public, communities and stakeholders directly impacted by the changes.”

The effective and ethical development of artificial intelligence: An opportunity to improve our wellbeing, ACOLA (2019)³⁹⁸

AI WILL CERTAINLY INTRODUCE NEW CHALLENGES IN ETHICS, LAW AND REGULATION. EXISTING LEGISLATION WILL NEED TO BE REVIEWED TO ENSURE IT IS FIT FOR PURPOSE IN THE NEW CONTEXT. SIMILARLY, SOCIETY WILL NEED TO REVISIT DEBATES AROUND FREE SPEECH, PRIVACY, INCLUSION AND PROTECTING DEMOCRACY IN LIGHT OF POTENTIALLY HARMFUL APPLICATIONS OF AI TECHNOLOGY. IN ITS PROGRESSIVE TRADITION, NEW ZEALAND SHOULD BEGIN THE CONVERSATION ABOUT THE ROLE OF AI AND DATA WITHIN OUR SOCIETY NOW.

AI and Ethics

Ongoing debate and consensus on ethical aspects should guide adoption of AI.

The field of AI Ethics has grown rapidly in recent times as detailed in Chapter 2.

Key ethical questions posed by AI include:

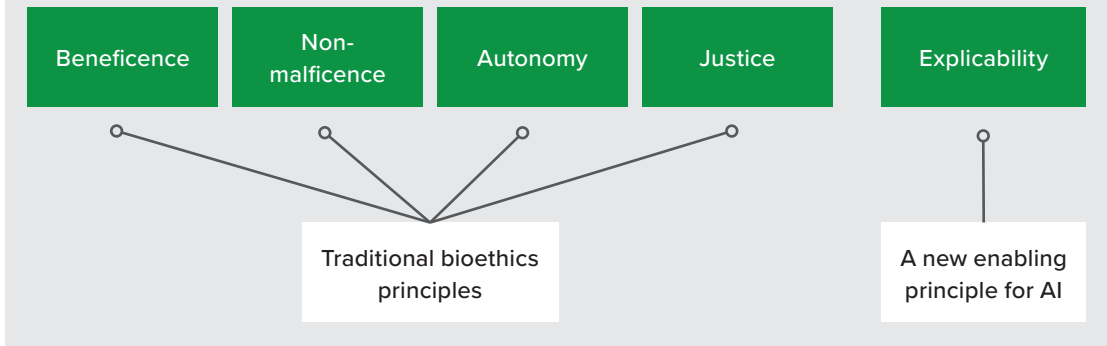
- Who is accountable for the actions of autonomous systems?
- How do we achieve transparency and explanation when many of these technologies are ‘black boxes’ and opaque even to their creators?
- How can we ensure fairness, given the potential for bias and injustices that are endemic in datasets?
- How can we ensure diversity and inclusion?
- How do we ensure AI does not cause harm?
- How can human dignity, autonomy, and rights be protected in practice?
- How do we define and protect data rights, data sovereignty and data privacy? (These have been previously discussed in Chapter 5.)

These issues are by their nature global – and as a result the AI ethics debate has highlighted the variability in ethical values worldwide. As noted in Chapter 2, recent work from the Berkman Klein Center for Internet and Society at Harvard University catalogued more than 32 sets of AI Principles.

It is notable that, for example, the IEEE *Ethically Aligned Design* report³⁹⁹ has only relatively short sections on non-Western values, for example Buddhism or Confucianism. Moving forward it will be useful and important to elicit other non-Western ethical approaches to AI and seek common ground for collaborative progress.

Here in New Zealand, it will be particularly essential to include the Te Ao Māori (Māori world view) perspective into any ethical AI guidelines proposed.

FIGURE 36: Floridi et al: Five AI Ethical Principles



Floridi et al. assessed six important statements on AI and ethics (the Asilomar AI principles, The Montreal Declaration for Responsible AI, the IEEE’s Ethically Aligned Design, European Commission’s European Group on Ethics in Science and New Technologies, the UK House of Lords Artificial Intelligence Committee’s report, and The Tenets of the Partnership on AI).⁴⁰⁰ They then developed a high level ethical approach to AI based on established biomedical foundation.

- **Beneficence:** Aim to do good (for example, promoting wellbeing, preserving dignity and sustaining the planet)
- **Nonmaleficence:** Do no harm, AI (and its developers) should respect privacy, security, and we should be cautious of its power
- **Autonomy:** Individuals should retain the power to make choices and also whether to decide to cede decision making powers to a machine
- **Justice:** Promote prosperity and preserve solidarity (for example, by eliminating discrimination and promoting equity)
- **Explicability:** The other principles should be enabled through intelligibility and accountability of AI systems

The authors articulate this as a “European Approach” but note that it is complementary to other approaches.

Implementing Ethical AI

Practically, how can we be sure that ethical principles are inherent in every new AI development, Government or business strategy?

FAIRNESS AND ACCOUNTABILITY

Along with initiatives such as the OECD Principles on AI and the Statistics New Zealand and the Privacy Commission’s Principles for safe and effective use of data and analytics, policy to promote fairness and accountability in AI adoption should clearly articulate which concepts of fairness the Government supports.

Policy could consider recommending ethics committees; mandating product reports and ensuring developers of critical infrastructure and applications adhere to regulatory codes of conduct.

AI Ethics Committees: Although many large technology firms currently have AI ethics policies, this does not necessarily lead to a smooth process of ethical assessment. Google’s nascent AI ethics board collapsed in its first week, following a staff backlash. There is a strong case for external ethical oversight. One solution is to adopt ethics committees as has been the norm in the healthcare sector for decades. There is precedent in New Zealand with the Health and Disability Ethics Committees, and various institutional review boards. Plans to deploy AI products might benefit from committee assessment with the principles of AI ethics as a benchmark.

Ethical Product Reports: Microsoft plans to add an AI ethics checklist to every product release.⁴⁰¹ One possible policy intervention could be to mandate an ethical product report to accompany each new AI product. There are various policy options for processes of certification and supervision of AI products.

Ethical AI Practitioners: Ethical behaviour begins with ethical people. Many professions have professional codes of conduct and disciplinary bodies. Industry or the Government could support the development of such self-regulatory codes of conduct for data and AI related professions and ensure that ethical duties are included. If people understand the merits of ethical AI they may demand licenced conduct from providers.

Different ethical traditions: Aotearoa sits in a geographically and culturally unique position with indigenous Maori, Western and, to a lesser degree, Eastern cultural ethical and political traditions coexisting within our society. The field of AI Ethics is at the cutting edge and in many ways exposes the new fault lines of global cultural tectonics. This is one area where New Zealand can play an active part in facilitating dialogue between diverse cultural and ethical traditions locally and internationally.

“New Zealand can play an active part in facilitating dialogue between diverse cultural and ethical traditions locally and internationally.”



TRANSPARENCY AND EXPLAINABILITY

The Government's Algorithm Review was an important first step in transparency. However, policy should ensure that individuals affected by decisions based on AI analysis have the right to request an explanation. Government may consider establishing an AI Ombudsperson, or formalise AI request for explanation procedures.

Policymakers may develop a framework, procedures and audit processes to scrutinize algorithmic decisions when necessary. This will be easier where data standards and development standards have been used, including documentation of the development process. Lessons can be imported from the aerospace and healthcare industries where the response processes to harms are often formalised.

SAFETY

AI technologies should not be rushed to market without thorough safety assessment

AI technologies pose a number of risks to safety that will need to be managed. Near term risks include the risk of harm posed by autonomous robotics and unexpected algorithm failures. Long term risks include the possibility that powerful artificial general intelligences may end up with goals that are not aligned with those of humans.

One obvious place that safety risks manifest is in the transport sector. Two high profile examples illustrate the risk. Firstly, a self-driving Uber killed a pedestrian in March 2018 in Arizona.⁴⁰² Other autonomous vehicle fatalities have also occurred. Secondly, all Boeing 737 MAX aircraft were grounded in March 2019 following the airline's second deadly crash to be precipitated by a malfunctioning AI flight system. Independent international affairs and aerospace industry analyst Alessandro Bruno said, 'I think this accident shows the limitations of artificial intelligence and it should convince Elon Musk and Tesla to slow down efforts to push down self-driving cars.'⁴⁰³

AI in healthcare could also pose safety risks. There are precedents where healthcare regulators have approved medical devices that subsequently cause harm. This is in part due to the low threshold for

evidence required by some regulators. For example, surgical mesh, some pacemakers and gastric weight loss balloons have all caused serious harms.⁴⁰⁴ AI has the potential to impact patients at scale through the use of robotics, diagnostics, and a range of treatments.

All industries will need to make safety assessments before deploying AI solutions. This will be especially critical where the technology will affect populations rather than individuals and essential infrastructure rather than leisure activities, for example.

AI, Law and Regulation

Many existing New Zealand laws apply to AI, however some new law may be needed

In earlier chapters we identified how AI can potentially affect privacy, accountability, surveillance and automation of government services which may require new laws or regulations. However, New Zealand has many existing laws and regulations that will apply to uses of AI, and further regulation may struggle to capture the relevant nuances of AI use in particular sectors. This is highlighted in the case of deepfake technology, where the recent Perception Inception report stopped short of calling for new laws without first understanding the complex interaction of existing legal regimes.

Areas of law in NZ that may similarly need analysis and close monitoring in the context of AI include:

Bias and Discrimination: New Zealand already has a legal regime that addresses bias and discrimination. The Human Rights Act 1993 provides protection against unfair discrimination, giving effect to New Zealand's binding obligations under international human rights law. The Government is subject to even more extensive human rights obligations under the New Zealand Bill of Rights Act 1990 (NZBORA) and human rights treaties in international law.

Together, the Human Rights Act and NZBORA address many AI ethical concerns including in relation to bias and discrimination. However, there may be value in clarifying the practical application of the principles in those laws in an AI context. This could be done by drawing upon existing understanding and reporting mechanisms found under international

human rights law and implementation of a human-rights based national ethics framework or code of ethics for the use of AI in New Zealand.

Public Sector use of AI and Algorithms: It is particularly important that issues of accuracy, human control, transparency, bias and privacy are addressed in use of predictive algorithms in the public sector. A University of Otago report entitled "Government Use of Artificial Intelligence in New Zealand" (May 2019) concluded that oversight measures are needed for government use of algorithms, while also noting that many existing public law principles already address transparency and fairness concerns. It may be more an issue of drawing out exactly what those principles mean in the context of AI and getting consensus on this.

AI and Employment Law: AI and automation will almost certainly impact the job market. New Zealand needs to consider how resulting job losses fit within existing categories of redundancy and unfair dismissal. Contractor rights may also need to be beefed up as technology such as AI spurs on the "gig economy". It may be that changes are needed to employment law. Also, if intelligent machines are employed by companies, we might need to define their obligations and rights, or even regard them as "legal persons" for the purposes of the law of obligations. The suggestion of full-blown legal personhood has proved controversial around the world, particularly as it may allow people operating and developing AI systems to hide from responsibility.

AI and Copyright Law: Many of the digital services available today didn't even exist when the last major review of New Zealand's Copyright Act 1994 occurred. To address this, MBIE has produced an issues paper with a view to amending the Act to help New Zealand's copyright regime remain robust and flexible. One issue is that although New Zealand's Copyright Act recognises that works may be "computer-generated" it's not clear who in practice would actually own the copyright of AI-generated work. For example, are the programmers the authors? What if customers have provided data that helps train the AI? What happens if there is "joint" creation of new work by multiple AI systems working together? Legal certainty is needed to assure investors of who owns what rights.

Robo-advice: There are legal and regulatory issues associated with robo-advisors that are currently under scrutiny. For example, in the financial services sector, amendments to financial services legislation (enacted in April but not yet in force) removed the requirement that only a natural person can give financial advice, which will make robo-advice permissible subject to prescribed requirements. It's unclear when this will come into force, but in the meantime, the Financial Markets Authority has announced an exemption is available to those looking to provide robo-advice, subject to certain requirements. Robo-advice may well present significant opportunities to provide greater access to justice to New Zealanders who struggle to get legal advice today.

Consumer Protection: New Zealand's current consumer protection laws, such as the Consumer Guarantee Act 1993 and the Fair Trading Act 1986, have not yet been tested in court as to whether they offer adequate protection to consumers in an age of algorithmic recommendations and automated buying.

Collusion: Some commentators argue that, compared to tacit collusion by human agents, algorithmic tacit collusion is more likely to occur. Tacit collusion can occur with pricing algorithms. For example, online vendors may use algorithms to monitor and undercut rival's prices. However, the algorithm may learn that there is no incentive to undercut prices, as the rival's algorithm simply undercuts in turn. This can lead to monopolistic pricing tendencies or could lead to perverse pricing.⁴⁰⁵ It may be that algorithmic tacit collusion needs to be made illegal to protect some of our valued institutions.

Contract Law and Liability: If damage is caused by any action or decision taken by an AI acting autonomously, legal issues may arise in relation to allocation of liability in relation to that damage. Legal intervention may be required to clarify how liability should be allocated, for example whether the responsible party is an AI developer, manufacturer, the owner or even the AI itself.

Personhood: There are already examples of inanimate things possessing legal personality, such as corporations and rivers. Some commentators argue that we need laws that establish the

personhood or otherwise of AI entities. This would allow sophisticated AI to be given the status of "electronic persons", making the AI responsible for any damage it causes and limiting the liability of its owner. In the European Union, the European Parliament has recommended that the European Commission create specific legal status for robots.⁴⁰⁶ However, the European Commission refused to adopt this recommendation in its AI strategy.⁴⁰⁷ Giving AI legal personality is a controversial question and is unlikely to gain traction, at least in the short term.

AI and Privacy Legislation: The Privacy Act 1993 does not explicitly contemplate AI, but organisations using personal information in the context of AI must comply with the Act and any applicable requirements in specified codes of practice such as the Health Information Privacy Code 1993.

Unlike GDPR, the new Privacy Bill currently making its way through Parliament does not include any requirement for algorithmic transparency for AI applications or provide individuals with the right to object to automated decisions. The Privacy Commissioner's submission on the Bill included a recommendation to address algorithmic transparency but to date that has not been included in the draft Bill. While it is possible that this may be addressed in a revised version of the Bill, any such significant amendment appears unlikely at this stage.

AI and Freedom of Expression: Following the Christchurch terror attacks, there was consensus at a political and industry level to address the dissemination of extremist content on the Internet through the Christchurch Call Agreement, led by New Zealand and Prime Minister Ardern (along with France and President Macron). This Agreement emphasises the need for research and development of technical solutions like AI to detect and prevent the upload of such content. AI can be used to automatically flag and filter the dissemination of extremist content, but legal issues may arise where such filtration becomes a form of censorship and encroaches on freedom of expression.

AI and Autonomous Weapons: AI has a range of applications in military and defence contexts, particularly in Lethal Autonomous Weapon Systems (LAWS). State use of this militarised AI would be

governed by areas of international regulation such as international humanitarian law or the United Nations Convention on the Law of the Sea. There are some efforts in the United Nations and disarmament community to propose a new regulatory regime for LAWS, but to date these have made slow progress and are not supported by most major international powers. While New Zealand has not publicly announced spending or investment in militarised AI, it might have an important role to play in shaping the international rules around this new and dangerous application of AI.

Autonomous Vehicles: The New Zealand government is encouraging the testing of autonomous vehicles in New Zealand. There are currently no legal barriers in New Zealand to the testing and operation of autonomous vehicles. However, there are a range of legal issues relating to the operation of autonomous vehicles that need to be considered, including fault allocation in the event of a crash and rules to ensure public safety. The adoption of operating standards may also need to be considered to ensure the widest range of technologies can be used in New Zealand, as well as the allocation of radio spectrum to ensure spectrum availability and to prevent interference.

AI and Freedom of Speech

AI is used to prioritize and recommend content online and this may require new approaches to civil liberties and freedom of speech

In recent years, digital technologies have – more than any time in the past – facilitated the free distribution of ideas. While generally seen in a positive light, these technologies have also been exploited to propagate hate and propaganda. Although not essential to these processes, AI has the potential to influence and amplify what is said, and how it is received. These issues are particularly salient to New Zealand in the wake of the tragic 15 March 2019 Christchurch Mosque attacks and the associated issues of content distribution and censorship.

Although AI was used to identify and remove 1.5 million copies of the terrorist's video from YouTube, a sizeable number of videos went undetected, partly due to modifications made to the video by those uploading.

New advances in AI technology are needed to more accurately identify problematic content. The scale of this task is simply too vast for humans to undertake.



HOW YOUTUBE USES AI TECHNOLOGY TO DETECT VIOLENT EXTREMISM⁴⁰⁸

YouTube has developed automated systems that aid in the detection of content that may violate our policies. Once potentially problematic content is flagged by our automated systems, human review verifies whether it indeed violates our policies. If it does, the content is removed and is used to train our machines for better coverage in the future. The account that posted the content generally receives a strike, and multiple strikes leads to account termination. With respect to the automated systems that detect extremist content, our teams have manually reviewed over two million videos to provide large volumes of training examples, which help improve the machine learning flagging technology.

Machine learning now helps us take down extremist content before it has been widely viewed. Our significant investment in fighting this type of content is having an impact: well over 90% of the videos uploaded in September 2018 and removed for Violent Extremism had fewer than 10 views.

– Extract from YouTube policy on violent extremism

In the UK, Facebook representatives have asked for government guidance regarding their obligations to report crimes identified on their platform.⁴⁰⁹

In the wake of the Christchurch attacks, Brad Smith, President of Microsoft, called for the development of a 'principled, comprehensive and effective' industry wide approach. He identified three key areas for action:

1. Prevention
2. Effective response (perhaps with an industry wide 'major event protocol')
3. Fostering a healthier online environment.⁴¹⁰

The tragic events of March 15th also led to the Christchurch Call summit in Paris (see box).

The USA did not sign the Christchurch Call agreement, citing concerns about freedom of speech and conflicts with its Constitution.⁴¹¹ It is increasingly apparent that contemporary society will need a discussion about civil liberties in the age of AI. In the USA, legal scholar Tim Wu argues that ‘the First Amendment [guaranteeing freedom of speech] is obsolete’.⁴¹² Wu further states that the foundational assumptions permitting freedom of speech are that information is scarce, and that the government is the main threat to speech (through enforcement).^{412a} In recent years, a fundamental change has taken place in how we communicate: it is no longer speech that is scarce, it is attention.

This has important implications in how we conceptually view freedom of speech. Speech (including hate, trolls and bots) is now used to freely harass and suppress other speech. A range of speech control techniques are used online, often automatically, to censor and degrade speech. For example, flooding media channels with large quantities of misinformation or spam content can detract from the content of interest without actually blocking it. In addition, attacking the author of the post with degrading comments can discourage further posts. These techniques and others like them are known as reverse censorship. Twitter CEO Jack Dorsey has explained how Twitter is investigating methods for measuring the “health” of conversations.⁴¹⁵ Again, AI again may become part of the solution to enable such measurement.

THE CHRISTCHURCH CALL – STOPPING EXTREMIST CONTENT ONLINE

Initial efforts to remove terrorist content from social media were made with the Christchurch Call hosted by New Zealand’s Prime Minister, Jacinda Ardern in Paris.

Although key decision makers such as President Donald Trump and Facebook’s Mark Zuckerberg were absent, others pledged to help eliminate terrorist content from social media.⁴¹³ Eight tech companies (Amazon, Facebook, Dailymotion, Google, Microsoft, Qwant, Twitter, YouTube), 17 countries including New Zealand, and the European Commission signed the Christchurch Call to Action, an agreement to work to stop the distribution of extremist content online. The focus is on collaboration with the goal of suppressing hate and extremist terror content, as well as impeding the radicalisation processes these platforms can facilitate.

The pledge includes expressly forbidding extremist content, enabling users to flag content, improving technology to detect and remove content, appropriate checks for live streaming. Concrete actions include Facebook’s decision to restrict people who have broken rules, such as sharing terrorist content, from using its live streaming feature.⁴¹⁴ The tech companies have pledged to assess how their business models may be leading users toward extremist online content and to make changes to prevent radicalisation. The AI Forum’s working group on ‘Ethics, Law and Society’ has published an accessible overview of some of the key issues.⁴¹⁶



PHOTO: Participants at the Christchurch Call conference, Paris, 2019.

AI and Privacy

Privacy is important because it provides common good.

“Big data analytics and artificial intelligence increasingly enable states and business enterprises to obtain fine-grained information about people’s lives, make inferences about their physical and mental characteristics and create detailed personality profiles,”⁴¹⁷ reported the UN High Commissioner for Human Rights.⁴¹⁸

AI technologies raise major new privacy issues. Potential threats include, unjustified surveillance, mass accumulations of personal data, the ability to reidentify deidentified data, bias, fairness and potential discrimination on the basis of data. These issues are often compounded by the lack of public understanding about uses of AI for purposes such as profiling, tracking and identification of individuals. While AI has been used to identify people who wish to remain anonymous and also to infer personal information about people on the basis of legitimately obtained data, such inferences can be used to make consequential decisions about individuals. In a striking example of privacy invasion, the Chinese Government uses facial recognition technology to identify Uighurs (a largely Muslim minority population) on the basis of facial traits and track their movements.⁴¹⁹

Databases of facial recognition data exist and your face is likely to be in them! They are owned by data brokers, social media companies and governments. Often these map detailed points on faces and store that data rather than an image. In June 2019, Microsoft announced that it had deleted its image database containing 10 million faces which was being used to train facial recognition systems.⁴²⁰ The FBI has scanned images of millions of Americans arrest mugshots and driver’s licences. Meanwhile, USA border security collects the fingerprints of non-US citizens. However, the House Committee on Oversight and Government Reform was told that one in seven searches of this facial data produced only innocent candidates, even though the search target existed in the database.⁴²¹

In a recent opinion piece, Jon Evans argues that privacy is important because it provides common

good, such as preventing the accumulation of imbalances in knowledge and therefore power. When personal data accumulates in massive quantities it has three key effects: First, accumulated knowledge about individuals subdues individual dissident thought. Second, a power imbalance emerges between those that have privacy and those that don’t. Third, accumulated private data can be used to sway public opinion. Mass privacy is good for society and individual choices to sell that out could erode these benefits.

“We may conclude that while individually, our privacy may usually be mostly meaningless, collectively, it is a critically important commons,”⁴²² Evans suggests.



TensorFlow

DEVELOPING AI TECHNOLOGY FOR BETTER PRIVACY

Google has been increasing research on federated learning, a new approach to machine learning that allows developers to train AI models without private personal data ever leaving the local device. This work is now open sourced as TensorFlow Federated.⁴²³

AI and Digital divides

Investment in education and communities is needed for equality of opportunity and digital inclusion

The continued existence of digital divides, where some New Zealanders lack a baseline of skills or digital access, could result in limited employment opportunities or lack of access to productivity enhancing or leisure services. This inequality of opportunity may further compound the previously detailed trend towards economic inequality.

Basic digital skills will become more important for workers as AI is integrated into more and more workplaces, leading to the automation of some tasks, and the creation of other entirely new ones.

In the Technological Change and the Future of Work issues paper, the Productivity Commission notes, “digital skills are increasingly important as nearly all forms of work are becoming more digitalised”.⁴²⁴ Education and training will be an essential part of ensuring that New Zealand workers can adapt to an increasingly digital workplace. Also, it’s not just in the workplace and other groups such as senior citizens are needing to upskill as well.⁴²⁵

There is currently no comprehensive data about, or agreed measures for, the extent of digital divides or digital inclusion in New Zealand. In the 2013 Census, 82 percent of the population had household internet access, compared to 67 percent in 2006.⁴²⁶ Research indicates that the communities and demographics most likely to be lacking the access or skills required to take advantage of the digital world include seniors, people with disabilities, people living in rural communities, families with children living in low socioeconomic communities, and Māori.⁴²⁷ There is a concern that New Zealand firms are often unable to find skilled workers for technical positions.⁴²⁸ The OECD reported

in 2016 that nearly half of New Zealand adults (45.3 percent) score at or below Level 1 (the lowest level) in problem solving in technology-rich environments.⁴²⁹

Ideally, increased education and investment in communities will lead to greater digital inclusion, an outcome in which all New Zealanders have equitable opportunities to participate in society using digital technologies.⁴³⁰

A Note on Terminology: There is a wide range of terminology in use to describe universal digital skills and access, including digital divides, digital literacy, “digital capability” and digital inclusion. The term digital divide has generally been used to refer to variability in the ability of different groups to access and use the internet and other technology. The Government has chosen the term ‘digital inclusion’ for its work programme, whereas other groups, including InternetNZ, primarily use ‘digital divide(s)’.

FIGURE 37: High-level timeline for Government action toward digital inclusion

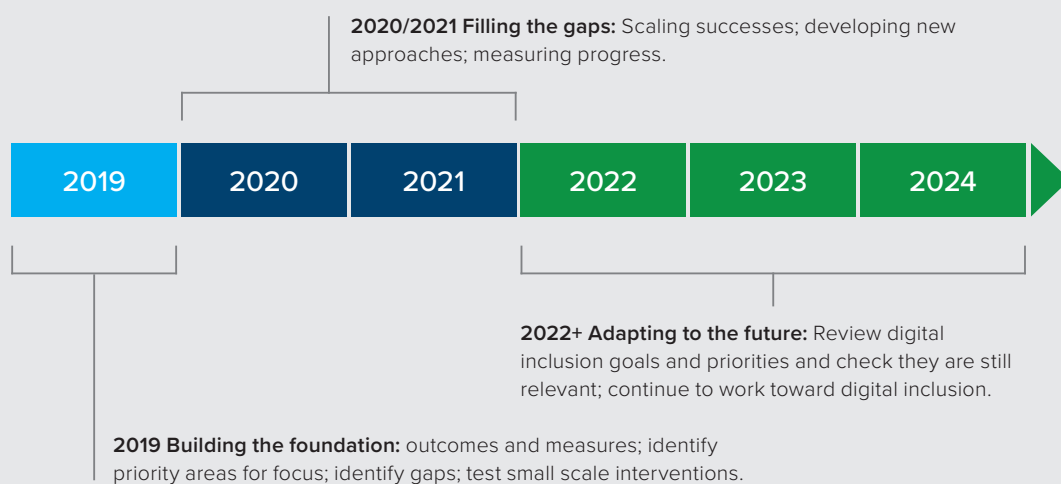


FIGURE: Adapted from The Digital Inclusion Blueprint, Te Mahere mō te Whakaurunga Matihiko. Published by the Department of Internal Affairs, 2019. Shared under a Creative Commons CC BY 4.0 license.

In May 2019, Hon Dr Megan Woods, then Minister for Government Digital Services, launched the Digital Inclusion Blueprint, Te Mahere mō te Whakaurunga Matihiko, which sets a goal of ensuring New Zealanders have “convenient access to, and the ability to confidently use, the internet through devices such as computers, smartphones and tablets”.^{430a}

The Blueprint introduces four elements that are all needed in order to be digitally included: motivation, access, skills and trust. It also sets out four roles that Government will take – lead, connect, support, deliver – and highlights that many other businesses, NGOs, iwi, communities and individuals are also contributing to this work, and have been for many years. The Blueprint was released at the same time as an action plan for 2019. The 2019 report *Out of the Maze: Building Digitally Inclusive Communities* also outlines a number of suggestions to help break down barriers to inclusion. The report notes that “a common theme in all the suggestions is the need to consult with excluded people and to work in partnership with trusted community groups to ensure the problems are solved.”

AI and Inequality

AI has the potential to reduce inequality, but could widen it, if adoption is not well planned

A growing body of evidence confirms that wealth and income inequality is increasing worldwide.⁴³¹ According to the World Inequality Report 2018, since 1980 the global top one percent of earners have captured twice as much of the growth in global income as the poorest 50 percent of the population. In New Zealand in 2015, the wealthiest tenth of the population own nearly a fifth of the country’s net worth, while the poorest half of the country has less than 5 per cent.⁴³² A complex set of factors have contributed to this shift towards increased concentration of wealth, but the unequal ownership of capital and the shift towards privatisation are major contributors. As digital technologies, including AI, become further embedded in the way companies and governments function, who will see the most financial benefits?

One of the significant AI policy discussions will be understanding whether the global diffusion of AI technologies will exacerbate or mitigate this trend towards increased inequality. For instance, banks

that use AI algorithms to make decisions about which applicants are eligible for mortgage loans may end up inadvertently excluding certain groups or communities from purchasing property.⁴³³ Likewise, AI assisted autonomous vehicles may reduce the number of truck drivers needed, resulting in job losses for workers.⁴³⁴ Some research has predicted potential outcomes, especially regarding the future of work and income,⁴³⁵ but more research needs to be completed to understand the likely impacts on the New Zealand workforce.

Of course, AI encompasses a broad range of technologies and applications, and its effects will vary, depending on the specific application and the business model being implemented. Just as some uses of AI have the potential to exacerbate inequality, other uses are likely to mitigate it, creating more opportunities for workers or reduce prejudice in systems or decision making. Increased adoption of AI technologies in the workplace may create whole new classes of jobs, resulting in new opportunities for well paid employment.

When planning for an AI enabled future, it appears there is an opportunity to foster economic growth and a robust, forward looking economy while ensuring that the gains are captured by many New Zealanders, not just a few.

In Summary

AI will certainly introduce new challenges in ethics and law. Existing legislation will need to be reviewed to ensure it is fit for purpose in the new context. Similarly, society may need to revisit debates around free speech and privacy, especially now that the context these concepts emerged from has changed.

Emerging applications of AI will also have a major impact on our society, wellbeing and the resilience of democracy. The accelerating economic returns on AI could drive income and wealth inequality, with resulting inequities in opportunity. Such inequity could be exacerbated by any digital divide, where some individuals have digital skills and access to new technologies and others do not. Social policy will need to anticipate, monitor and respond appropriately to any discrepancies.

8. AI Readiness and Investment

“AI will be useful wherever intelligence is useful, helping us to be more productive in nearly every field of human endeavour and leading to economic growth.”

HARRY SHUM – Executive Vice President of Microsoft AI and Research Group and
BRAD SMITH – President and Chief Legal Officer at Microsoft

AS ILLUSTRATED IN THE CASE STUDIES WITHIN THIS REPORT, A NUMBER OF NEW ZEALAND ORGANISATIONS ARE ACTIVELY WORKING WITH AI AND BEGINNING TO YIELD RESULTS FROM THEIR INVESTMENTS. HOWEVER, THE DEPTH AND BREADTH OF AI ACROSS THESE VARIOUS INDUSTRY SECTORS IS STILL A SMALL PART OF THE ECONOMY.

Are New Zealand Organisations Prepared for AI?

In addition to collecting and managing fit-for-purpose data and employing individuals with technical and soft skill talents, organisations need to be AI ready with appropriate knowledge, policies, processes and partnerships in place.

We are seeing fast movement on this globally. A survey of more than 3,000 business executives by the MIT Sloan Management Review and Boston Consulting Group reports that leading companies such as Chevron, Allianz, and Daimler have committed deeply to AI with priorities for revenue-generating AI. High level management is involved in AI initiatives and 90 percent of executives from these companies report that AI strategies are in place.⁴³⁶

A number of recent studies have investigated the readiness of local New Zealand organisations to adopt AI.

Oxford Insights described the readiness of government agencies around the world to make use of AI solutions, and found New Zealand to be in ninth position among the 35 OECD countries in 2017⁴³⁷ and thirteenth among an expanded list of 194 countries in 2019.⁴³⁸ The report found that New Zealand’s positioning could be improved with a stronger score for innovation (linked to lifting low R&D spending), enhancing the AI startup ecosystem

and number of startups, and by improving what is relatively limited availability of data.⁴³⁹ Oxford Insights notes that New Zealand’s strong performance on the index is due to a central government creating opportunities for effective digital services. However, continued strong performance depends – at least in part – on the Government producing a strategy (and by implication funding) for AI. They recommend the strategy should, ‘take advantage of New Zealand’s small size and efficient government to create a niche for itself globally, perhaps in piloting innovative AI applications in government.’⁴⁴⁰

Callaghan Innovation’s discussion paper, *Innovation through AI* highlighted the concern that New Zealand businesses will lack the capability to make use of AI. Callaghan Innovation CEO, Victoria Crone shared her concern that New Zealand businesses are lagging behind regional competitors in Australia and Asia in adoption of new digital technologies.⁴⁴¹ The Callaghan report notes that ‘*the stakes are high, with Forrester Research predicting that globally, by 2020, businesses that are driven by the insights provided by AI, big data and the Internet of Things will “steal” \$US1.2 trillion a year from competitors who don’t embrace these opportunities*’.⁴⁴² Similarly, the **Productivity Commission** finds that New Zealand does not yet have high levels of innovation and capability, partly due to low levels of R&D, moderate uptake of new ways of operating, with generally a weak level of management practices, low returns on

AI Adoption in New Zealand

In 2018, just 1.2 percent of New Zealand organisations said they have adopted AI solutions.⁴⁴³ 62 percent told IDC that they have plans to implement the technology within the next five years (i.e. by 2023). However, in IDC's 2018 Future of Work Survey the current adoption rates are much higher.⁴⁴⁴ This is a reflection of the overall lack of awareness of what is and isn't AI. Small, relevant and practical uses will boost the overall AI investment in the next two years and beyond and this will apply to all sectors.

In IDC's 2018 APEJ Cognitive and AI Adoption Survey, about 20 percent agreed or strongly agreed with the statement, "My organisation should have in-house cognitive/AI capabilities and the skill set comprised of both technology and dedicated staff." This desire for inhouse resources suggests the current skills shortage may worsen. Organisations may need to look to alternatives to get skillsets. This could include reskilling the existing workforce or seeking staff from the gig-economy.

IDC asked organisations about the importance of AI to gain a competitive advantage in the next five years. Over half (60 percent) of the surveyed organisations agreed or strongly agreed that AI would be a critical tool for their organisation. This reaffirms the intent to deploy by 2023.

New Zealand businesses told IDC that the key benefits they expect to see from AI systems deployment is increased process automation, increased employee productivity and to uncover new insights. The top AI use cases include automating repetitive tasks, security and fraud detection and real-time inventory management.

AI SYSTEMS ADOPTION IN NEW ZEALAND

Q. Which of the following best reflects your organisation's status regarding cognitive/AI Systems?

My organisation is exploring cognitive/AI systems but has no plans to adopt for now.	37.0%
My organisation is planning to adopt cognitive/AI systems in two to five years.	44.4%
My organisation is planning to adopt cognitive/AI systems within two years.	17.3%
My organisation has adopted cognitive/AI systems.	1.2%

AI WILL BE CRITICAL FOR COMPETITIVE ADVANTAGE

Q. Within 5 years, Cognitive/AI Systems will be a critical requirement for businesses to gain a competitive advantage.

Strongly Disagree	2.5%
Disagree	10.0%
Neither agree nor disagree	27.5%
Agree	52.5%
Strongly Agree	7.5%

TOP AI USE CASES FOR NEW ZEALAND AND AUSTRALIAN BUSINESSES

Q. Please rank the top 3 business processes in which you deem internal data is amenable/usable for the adoption of cognitive/AI solutions.

- 1. Automated repetitive tasks**
(Claims, invoicing, lead generation, asset management, quality assurance, etc.).
- 2. Security and fraud detection**
(Threat/anomaly/ intrusion detection and prevention).
- 3. Real-time inventory management**
(Demand, inventory and sales forecasting).
- 4. Issue and root cause**
(Discovery).
- 5. Adaptive and automated**
(Product/service interface).
- 6. Biometrics**
(Know your customer).

SOURCE: IDC, 2019

innovation, limited international connections, limited absorptive capacity, and capital shallowness.⁴⁴⁵

IDC's Microsoft-commissioned report, *Future Ready Business Survey: Assessing Asia's Growth Potential Through AI* surveyed 1,605 business leaders and 1,585 workers at organisations with more than 250 staff across the region and included New Zealand. Respondents were decision makers involved in shaping their organisation's business and digital strategy.⁴⁴⁶ Results for New Zealand showed that 67 percent of business leaders have yet to implement plans to upskill employees for an AI enabled future.

'A significant majority of the business leaders and workers surveyed believe cultural traits that support AI initiatives, such as risk-taking, proactive innovation, as well as cross-function partnerships among teams, are not pervasive today', says CIO magazine's Divina Paredes.⁴⁴⁷

Results from **PwC's annual CEO Survey 2019**, show that 84 percent of New Zealand CEOs believe AI will transform their businesses within the next five years. In addition, 67 percent believe the impact of AI will be greater than that of the internet. Only 30 percent

have no current plans to pursue any AI initiatives and 39 percent are planning to implement AI in the next three years. However, no New Zealand CEOs reported that AI was widely used in their organisations and nor was it seen as fundamental to their businesses.⁴⁴⁸

This is important because when we examine AI Maturity Models (below), it is obvious there is a lag between foundational, approaching, aspirational and mature deployment. Planning now and implementing pilots may help minimise long-term lag in adoption.

"There is a clear divide in CEOs' belief in the ability of AI to change how we work and the level at which it's being implemented."

MARK AVERALL
– CEO and senior partner a PwC New Zealand.

FIGURE 38: IDC research New Zealand / Australia comparison



25%

New Zealand is Different

Have no immediate plans to adopt AI.



48%

60%

Strongly believe within 5 years AI Systems will be critical to gain competitive advantage.

44%

25%

Say it's important their AI solutions are based on/ incorporate open source technology.

35%

43%

Say their data scientists are seasoned business analysts with a solid background in statistics.

20%

17%

Say their data scientists are experienced software engineers with hands-on knowledge of popular machine learning and deep learning.

39%

Māori Business and Organisation Perspectives

SUCCESSFUL MĀORI BUSINESSES ARE AN INTEGRAL PART OF A SUCCESSFUL NEW ZEALAND BUSINESS ECOSYSTEM. UNLOCKING THE SCIENCE AND INNOVATION POTENTIAL OF MĀORI KNOWLEDGE, PEOPLE AND RESOURCES WILL BENEFIT MĀORI, AND AOTEAROA AS A WHOLE.⁴⁴⁹

In addition, Te Mana Raraunga has identified supporting the development of sustainable Māori digital businesses and innovations as one of the key ways through which to advance Māori data sovereignty.⁴⁵⁰

There is an opportunity for businesses and organisations to use AI for purposes that serve Māori, reflect Te Ao Māori in their business practices, boost Māori culture, and create value for Māori workers and communities. There are also opportunities to continue to incorporate te reo Māori and Te Ao Māori into technology in a way that is future looking, congruent with tikanga, and beneficial to Māori. For example, Dr Te Taka Keegan has pointed out that the future of human computer interaction is not in typing, and that because Te Reo does not currently feature in speech recognition tools, there is a need for Māori to develop this capability.

While there are many thriving Māori tech businesses, our scan of the AI business sector showed limited Māori businesses and organisations incorporating AI. This indicates that we need to find out more about the sector, and encourage development.

There is no one definition of what makes a Māori business or organisation – it might include having Māori founder or owner, embedding a Māori kaupapa/purpose and world view into their day-to-day business, being an iwi organisation, or a mixture of all of these. Within this broad definition, current Māori businesses and organisations using AI include:

Robotics Plus, which combines expertise in agricultural technology, machine vision and artificial intelligence to create robotics solutions for smart farming and forestry applications. Robotics Plus won the Māori Innovation category at the Hi Tech Awards 2019.

Kōrero Māori, a project to teach machines Te Reo Māori, from Te Hiku Media, supported by Te Punaha Matatini, Te Puni Kokiri and Dragonfly Data. The project has two sister websites: kōreromāori.com, where people can record themselves speaking Te Reo phrases, to help train machine learning models, and kōreromāori.io, which offers language tools including a transcription tool and speech recognition API. All data for the project is collected under a Kaitiakitanga licence, reflecting that Indigenous people do not have a concept of private ownership of land and resources.⁴⁵² Kōrero Māori won the Aotearoa New Zealand Community award at Ngā Tohu Reo Māori, the National Māori Language Awards in 2018.⁴⁵³

Straker Translations, which offers translations to businesses and individuals worldwide. Straker Translations uses machine learning to support human translators by sorting through huge quantities of data and finding patterns in languages, and by helping to match the best translator to a job.

While there have been some exciting developments in the Māori tech sector and AI specifically, it is still a nascent field. In order for the Māori AI ecosystem to thrive, there are challenges when it comes to finding investment and taking businesses to commercialisation. Aroha Armstrong, group manager of the Māori Economy Team at Callaghan Innovation, has emphasised the need for further development of Māori business, with a focus on “building business capability, connecting networks, supporting deeper R&D, and providing co-funding.”⁴⁵⁴

Māori business networks are a way to strengthen community and provide mentorship and opportunities for collaboration. Te Tira Toi Whakangao (T3W) is a group of global Māori tech companies and Māori sector investors working to build the Māori tech ecosystem. Callaghan Innovation’s Māori Economy Team is a group of dedicated Māori business innovation advisors, providing R&D and business innovation advice, as well as connections to expertise and co-funding. Kōkiri, the first Māori accelerator programme, launched last year to provide a pathway for Māori businesses (including tech and AI businesses) to grow and hone their work.

Strategy, Maturity and Change

In order to begin their AI journey, New Zealand organisations may consider the use of maturity models and/or excellence models to frame the raw resources needed and ongoing progress in a systematic and measurable way.

For example, Gartner's simple AI Maturity Model illustrates five levels:

- Level 1 Awareness (early interest in AI)
- Level 2 Active (experimenting with AI)
- Level 3 Operational (AI is in production and delivering value)
- Level 4 Systemic (AI is used pervasively for digital processes and disruptive new business models)
- Level 5 Transformation (AI becoming part of an organisation's DNA)⁴⁵⁵

Despite the current hype, many organisations are still determining how to make AI work for them and are currently in the earliest awareness phase of the maturity model. Only a handful of

organisations are in the transformation phase.⁴⁵⁶ According to Gartner, enterprises cited finding use cases, defining strategy, security and privacy, risks and integration complexity as roadblocks to AI. Furthermore, a majority of organisations mentioned finding where to start was a concern and had not met overly optimistic timing goals.⁴⁵⁷

There are other useful AI maturity models which describe a similar readiness journey for organisation cultures. Also many companies provide AI maturity assessments – for example the Microsoft AI Readiness Survey for organisations.⁴⁵⁸

The detailed maturity model shown on the next page is from *The AI Maturity Playbook*, by Susan Etlinger for Altimeter. It illustrates how an organisation's AI maturity spans multiple dimensions: Strategy, Data Science, Product and Service Development, Organisation Culture, and Ethics and Governance.^{458a}

FIGURE 39: AI Maturity Model with Indicative Return on Investment

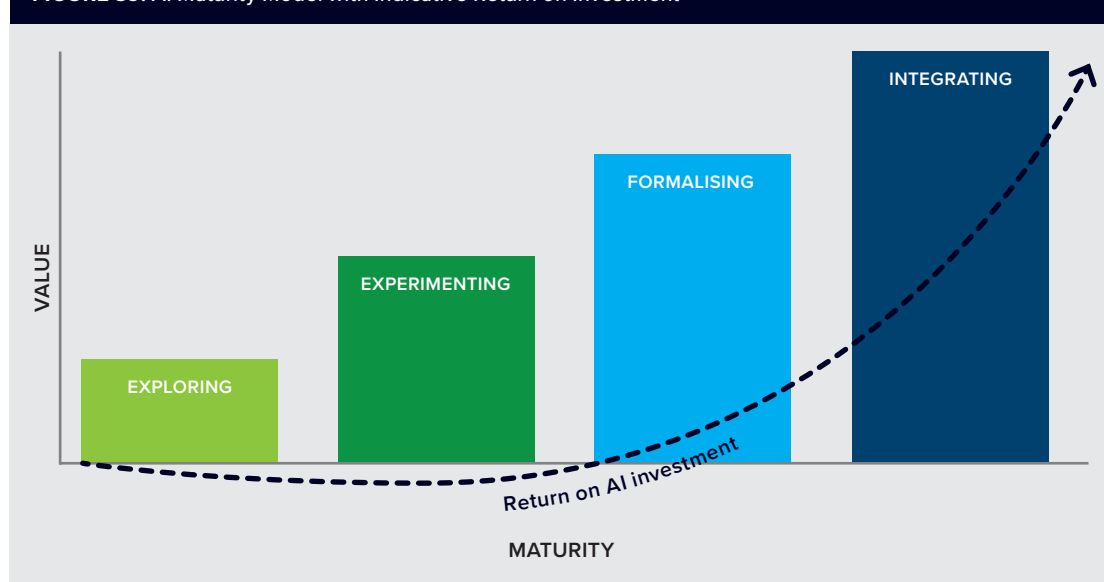


TABLE 7: Four Stages of AI Maturity

	Phase 1: EXPLORING	Phase 2: EXPERIMENTING	Phase 3: FORMALISING	Phase 4: INTEGRATING
STRATEGY	Objectives are undefined and no resource or budget has yet been allocated.	Discrete proofs of concept focus on cost reduction, productivity improvement, and/or Robotic Process Automation (RPA).	An expected part of strategic planning, focused on customer experience.	Integral to agile business and a critical component of digital transformation and competitive advantage.
DATA SCIENCE	Data is sliced, not in accessible useful form; analytics are largely descriptive and retrospective.	Organisation has committed to data strategy and is moving from descriptive into predictive analytics.	Data strategy is becoming a core competency but AI is not yet scaled across the organisation.	Organisation benefits from a compounding data advantage
PRODUCT & SERVICE DEVELOPMENT	Business cases present, but no development is underway as yet.	Organisation has begun to use APIs and internal or external resources to perform proofs of concept and pilots.	AI is becoming a critical part of product and service development.	AI is a core development competency across the organisation.
ORGANISATION & CULTURE	Seen as promising but unproven. Not yet seen as a priority at the C-level.	Organisation may have a Chief Data Office, but data science and AI projects are discrete rather than critical elements of an ensuring product, service, or business	Clear understanding of an optimised relationship with data science and AI resources, whether they are external services and platforms, a larger ecosystem, or a combination.	Organisation has a learning organisation mindset: design thinking and experimentation are valued in the culture.
ETHICS & GOVERNANCE	Emerging understanding of AI governance issues but no principles or processes present.	Organisation has identified and communicated ethical principles of AI and is implementing policies and processes to support them.	AI ethics and governance processes are formalized throughout the business.	AI ethics/ governance is embedded in corporate practice and customer experience and is part of performance evaluations and incentive programs.

WHY DO AI PROJECTS FAIL?

While every failure is different, Louise Francis at IDC outlines the top reasons for failed AI projects. These include solving the wrong problem, solving a problem that didn't need solving, or simply using the wrong tools; taking on a project that is too big or too grand; Poor data quality (includes ongoing data governance and cleansing); trying to automate everything when AI models may not be ready for it; or forgetting to regularly monitor, update and adjust the solution after the initial project is complete.⁴⁵⁹

Coordinating Organisations Through Partnerships

In addition to data, talented individuals, and AI-ready organisations, partnerships will help drive the success of AI in New Zealand. According to Callaghan 'collaboration is New Zealand's strength on the global AI stage' and New Zealand digital firms will thrive through collaborations with corporates.⁴⁶⁰ New Zealand's current Data Strategy and Roadmap (Dec 2018) identifies partnerships as one of four key focus areas. The strategy calls for partnerships across government and between public and private organisations.⁴⁶¹ It will, of course, be important to monitor the progress and results generated through such partnerships to identify success and failure factors, as it is not obvious that partnerships are the solution to every problem.

The AI Forum, with its membership drawing from across the AI ecosystem, is well positioned to facilitate future partnership initiatives.

Andrew Ng's AI Playbook calls for organisations to look outwardly as they explore AI pilots.⁴⁶² New Zealand ventures exist to support these early explorations and innovations. Aware Group, a Waikato based AI services company, offers proof of concept development to help organisations explore rapid pilots.⁴⁶³ Theta is another technology consultancy feeding the appetite of local organisations by building out AI solutions.⁴⁶⁴

Consultancies such as Deloitte can partner with New Zealand firms to develop AI strategies and solutions. Accenture provides services to help New Zealand organisations unlock the power of 'new technologies such as robotic process automation, virtual agents, cognitive robotics, machine learning, deep learning, natural language processing and video analytics'.⁴⁶⁵ Tech firms such as Microsoft and SAS are providing support services for those who wish to begin to explore AI.

PARTNERSHIP EXAMPLES

Precision Driven Health

THE PRECISION DRIVEN HEALTH PARTNERSHIP BETWEEN UNIVERSITIES, DISTRICT HEALTH BOARDS, THE GOVERNMENT AND ORION HEALTH IS EMERGING AS A KEY FACILITATOR OF INNOVATION IN HEALTHCARE IN NEW ZEALAND.

This partnership means that researchers have access to data that they otherwise would not be able to access and also that universities and the public hospital system can leverage the commercial acumen of private health technology firms.

Robotics Plus

THE UNIVERSITY OF AUCKLAND, UNIVERSITY OF WAIKATO, PLANT AND FOOD RESEARCH, AND MBIE HAVE PARTNERED TO SOLVE INDUSTRY PROBLEMS THROUGH SMART RESEARCH COLLABORATION.

The project will develop and commercialise robotic solutions to overcome the shortage of seasonal agricultural workers, for example in the kiwifruit industry.

Conclusion

THE OPPORTUNITIES ABOUND TO HARNESS THE POWER OF AI FOR A PROSPEROUS, INCLUSIVE AND THRIVING FUTURE FOR ALL NEW ZEALANDERS.

In this report we have aimed to energise and inform our national conversation on AI.

Our analysis, following months of research, dialogue, writing and iteration, is that AI technology is ready to be applied to New Zealand's goals – whether they relate to citizen Wellbeing, Sustainability or a more productive Economy.

However in the last few years New Zealand has been slow to act with agency to seize these opportunities. The pillars for success which we identify in the report – Investment, Research, Talent and Skills, Trusted Available Data, Ethics and Regulation - must all be grown together to achieve sustainable, successful outcomes.

This is not just a government responsibility – we sense that a partnering approach between the private sector, government, academia and iwi will see a diversity of voices contributing to the conversation and enable a more action-oriented approach.

Relatively modest partnership investments in flagship AI projects – for example in health, road safety, conservation of our environment and improving citizen life outcomes – promise to deliver results and also develop long term strategic capability.

It is time to continue talking – but start acting.

We look forward to taking the next steps together on New Zealand's AI roadmap.

Towards Our Intelligent Future
Te ara mō tātou atamai o āpōpō

Supporting Materials & Appendices

Understanding and Explaining AI: Machine Learning

Strictly speaking, the term “artificial intelligence” is a collective noun for a diverse set of techniques and technologies which have evolved since the 1950s. Since the 1980s, however, machine learning techniques have become the most widely used and these days AI and machine learning are often used interchangeably, although this is technically incorrect.

Machine learning techniques are categorised by the fact they can find complex and subtle patterns in data without being explicitly programmed to do so. Machine learning is often used to predict things, usually based on what has happened in the past.⁴⁶⁹ Using statistical analysis and algorithms, machine learning can quickly process large volumes of data. Machine learning can provide insights at a scale and speed that would be difficult or impossible for human brains to do alone.

For example, a supermarket manager may want to predict whether a first-time customer is likely to become a repeat customer, so they can effectively target advertising. A driverless car manufacturer will need to predict whether bright spots on the road in front of a car are more likely to be a road marking or a person, so they can program the car to avoid collisions. Machine learning can help in these and similar situations because it can sort through huge amounts of data and use it to make predictions about future customers or bright spots on the road.

So, how does this work? This section outlines the key components of machine learning; algorithms, data and models and how they fit together.

Machine learning may be viewed as a way of capturing complex relationships between recorded data and particular outcomes or realities. For example, a set of recorded image pixels with certain colour and luminance values may equate to an image of a cat — or maybe to some other four-legged creature. However, the underlying “rules” that dictate when a certain collection of image pixels equates to “cat” may be extremely difficult to manually uncover and code. Machine learning, in contrast, is a way of “learning” an approximation of such “rules” from a set of examples, something like the way a child’s mind “learns” after enough images of “cat” have been seen to reliably spot cat pictures.

ALGORITHMS

The machine learning process requires choosing a generic mathematical method for representing data, and then systematically evolving the chosen method in such a way that it comes to optimally represent a particular set of examples (a “training set”). This process of adapting a generic method to best fit specific data is known as training, and it is executed by mathematical algorithms.

An algorithm in this sense is a sequence of instructions, a “recipe” for executing a task, in this case the task of adapting a generic model into a model that represents specific data. The word “algorithm” is a general term that can be applied to all manner of systematically executed tasks; a recipe for baking a cake could certainly be regarded as an algorithm. Most frequently however algorithms are associated with computer science, simply because everything a computer does, even responding to a single mouse click, requires an exact, correct sequence of instructions.

Algorithms are core to machine learning, because they can specify a way that a generic mathematical method can be progressively evolved to “fit” a provided set of data examples, and hence to “learn” those unknown rules that — for example — determine the “cat-ness or not cat-ness” of images. The end result of simple mathematical methods and training algorithms may be explainable models that provide insight into the underlying rules, but often — like in the case of cat images — the end result is more like a black box, just as we do not yet understand how a human brain can accomplish the same task.

LEARNING STYLES

Machine learning algorithms vary in terms of their learning styles. Different learning styles work well with different kinds of data sets. Broadly, these learning styles can be divided into four types: supervised, semi-supervised, unsupervised and reinforcement learning.⁴⁷⁰

Supervised learning is the most widely adopted form of machine learning.⁴⁷¹ In supervised learning methods, each piece of training data is labelled with a corresponding output. For example, each item in a training set of image files would each be labelled ‘house’ or ‘not house’. The algorithm uses this labelled data to create a statistical model that figures out

what features divide a ‘house’ from a ‘not house’, and which can then accurately place new, unlabelled, data points into the pre-specified categories.

Unsupervised learning finds patterns or commonalities in datasets that have not been previously labelled or classified.⁴⁷² Unsupervised learning models create labels or categories by deducing patterns in the input data, rather than classifying data by predefined categories (which is the case with supervised learning). The goal is usually to cluster the data into groups with different characteristics. It may do this through a mathematical process to reduce redundancy, or it may organise data by similarity.⁴⁷³ One application for unsupervised learning is exploratory analysis, that is finding patterns or trends in large datasets where it would be otherwise impractical or impossible for a human to do so.⁴⁷⁴

Generative Adversarial Networks (GANs) are one class of AI algorithms used in unsupervised machine learning, involving a system of two neural networks learning by opposing each other in a zero-sum game framework. GANs can be used to produce authentic-looking synthetic images from large sets of training data.

As the name implies, **semi-supervised learning** uses a mix of labelled and unlabelled data inputs. The inclusion of unlabelled data can help improve the accuracy of semi-supervised learning algorithms, by providing additional context and a larger sample size.⁴⁷⁵ **Reinforcement learning** is an area of machine learning that uses a system of reward and punishment for learning how to attain a complex objective, sequentially, over time.⁴⁷⁶ Often, reinforcement learning is used in an uncertain, dynamic environment, where the conditions and the data inputs are constantly

changing. A reinforcement learning algorithm will try different ways to achieve the objective, get feedback on whether the methods were successful, and then adjust its methods until the objective is achieved.⁴⁷⁷ Gameplay AI like AlphaGo use reinforcement learning.

There are dozens of mathematical methods that can be used in machine learning algorithms, depending on the type of problem to be solved. Here are the most well known and effective methods, which can be used independently or as building blocks for other algorithms.⁴⁷⁸

Broadly speaking there are five main classes of mathematical approaches to machine learning, shown in Table 1 below. Mathematical methods for both supervised and unsupervised learning are explored in more detail below.

Choosing the best mathematical method can be a process of trial and error. Computer scientists will generally start with the simplest method they consider appropriate for the task, before inputting a collection of data and running the learning algorithm. This produces a model that will hopefully be able to accurately categorise or analyse new data.

DATA

Data is a key ingredient in machine learning. One of the reasons that machine learning has become more widely used in recent years is the greatly increased availability of suitable data. Both the data and the ability to store it has lowered in cost and become more accessible.

While there are many definitions, for the purposes of this report we define data simply as information in digital form.⁴⁷⁹ There are many types of data

TABLE 8: Mathematical Methods

MACHINE LEARNING APPROACH	EXAMPLE MATHEMATICAL METHOD(S)
Connectionist	Neural networks
Evolutionary	Genetic algorithms
Case-based	Nearest neighbour
Induction	Decision trees
Statistical	Support vector machines, Bayesian networks

Transfer Learning is a term used to describe using knowledge gained from solving one problem domain and applying it to a different but related domain. For example, knowledge gained while learning to recognise cars could apply when trying to recognise other vehicles. In practice this involves taking an existing machine learning model and modifying it to apply to the new domain, thus “transferring” the knowledge. In practice either supervised or unsupervised learning could be utilised for this.

used in machine learning, depending on the problem being solved or the opportunity being harnessed. The endless possibilities include structured and unstructured data formats: images, text, measurements, sound clips, sensor data, and information about people’s online activity including searches, clicks, posts, purchases, or streaming history. Many machine learning algorithms process numerical data, so some types of raw data, including images and text, may first need additional processing to represent them as numbers.


Data is used in three stages of the machine learning process:

- **Training:** a dataset, consisting of a large number of data instances, is fed into a computer program that runs a machine learning algorithm to produce a model. The model can then be used to classify previously unseen data. Depending on the learning style, training data will be either labelled or unlabelled.
- **Testing and validation:** data previously unseen by a model is used to check whether it is working in the way expected. It is also used to help adjust the parameters of the model to ensure it’s working as well as possible.
- **Deployment:** new data is fed into the final model to gain insight and predictions.

TABLE 9: Identifying AI Opportunities: Data Inputs and Machine Learning Outputs

In *The Business of Artificial Intelligence (2017)*⁴⁸⁰ economists Erik Brynjolfsson and Andrew McAfee present a technique to identify opportunities by mapping data inputs to machine learning outputs as illustrated in the table below.

INPUT X	OUTPUT Y	APPLICATIONS
Voice recording	Transcript	Speech recognition
Historical market data	Future market data	Trading bots
Photograph	Caption	Image tagging
Drug chemical properties	Treatment efficacy	Pharma R&D
Store transaction details	Is the transaction fraudulent?	Fraud detection
Recipe ingredients	Customer reviews	Food recommendations
Purchase histories	Future purchases behavior	Customer retention
Care locations and speed	Traffic flow	Traffic lights
Faces	Names	Face recognition

FIGURE 40: Main steps for training and deploying machine learning models


Below we look more closely at the process for building, testing and deploying machine learning models, including the important role that data plays.

MODELS

A model is the outcome of the machine learning training and testing process. Once the mathematical methods have been trained using a learning algorithm and training data set, it will produce a model that can then be tested in a validation process. The goal is to produce a model that can accurately process and categorise new, previously unseen, data. In many cases, a computer scientist will adjust the training data, or provide new and different training data, to create multiple models that each work slightly differently and produce different results.

Once a model has been produced and tested, it's time to implement the model in the real world. Companies and organisations take machine learning models, which can be custom created or purchased as a service and run them with their own data sets. At this point, they will need to determine which model is the most appropriate and whether it needs to be modified in order to provide accurate results.

Mathematical Methods for Machine Learning

MATHEMATICAL METHODS FOR SUPERVISED LEARNING

Linear and logistic regression: statistical methods that can be used to model relationships between one

or more predictor ("independent") variables, and a single dependent ("response" or "outcome") variable. In linear regression, the variables are numeric. For example, if you have a data set listing the number of litres of milk purchased by customers (predictor) and want to assess how many boxes of cereal each of these customers are likely to purchase (response). In logistic regression, there may be one or more predictor variables, and the response variables will be binary values (for example true/false, yes/no, 0/1). Logistic regression could be used to predict whether patients will suffer from lung cancer (yes/no), based on variables like cigarettes smoked per week and exposure to air pollution. Deep learning neural networks use logistic regression as a standard unit.

Decision tree learning: a type of branching graph, similar to a flowchart, that can be used to make predictions. At each branching node, a binary or either/or question is asked of the data point, which sends the data point in one of two directions depending on the answer. Additional questions may be asked, making the tree more complex. Decision trees are good for classifying data, especially when there are multiple variables that need to be considered in sequence.

Support vector machine (SVM): a method designed for two-class classification problems. The essence of SVM is the idea of separating two classes (eg. "cat" and "dog") by constructing an optimal separating "soft" margin between the classes, concentrating on only those data points in the training data ("support vectors") that are in close proximity to (ie. impact upon) this margin. Notably, SVM makes no assumptions

A sample of faces produced by a Generative Adversarial Network trained on face datasets.



SOURCE: thispersondoesnotexist.com

about the statistics of the training data, is outlier resistant, and is easily extended to encompass highly non-linear problems. A first choice, benchmark method for many classification tasks, and can be extended to multi-class and regression problems.

This contrasts with a method such as Linear Discriminant Analysis, also a two-class classifier, which focuses on obtaining maximal separation of the class means, and where all training data points directly impact the model.

Neural networks: are described on Page 31.

Ensemble methods and other "meta-learners": A powerful supervised learning approach that involves combining multiple instances of models ("learners", of types such as those above) into what are known as "meta-learners", with a supervised learning algorithm then applied to the entire meta-learner as a whole in addition to its members. There are numerous meta-learning strategies, but they are all like applying the "wisdom of the crowd" idea to machine learning.

MATHEMATICAL METHODS FOR UNSUPERVISED LEARNING

Deriving value from data via unsupervised learning may involve a number of distinct motivations and techniques. Some of the most common are:

Visualisation of complex data by reducing their dimensionality in such a way that points that are "close" in their original form are similarly close in a 2 or 3 dimensional representation – allow complex data to be accurately represented as points in 2-D plot.

Clustering data, that is arranging data into groups in such a way that within-group differences are smaller than between-group differences.

Exposing latent variables often the property of interest in a dataset is something that could not be directly measured – a trivially simple example could be Body Mass Index (BMI). This cannot be directly measured, as it is a latent or underlying variable, but it is easy to measure height and weight across a population and then construct the latent variable (BMI) as a mathematical combination of these. This general technique has a multitude of uses and the pre-eminent mathematical method is known as Singular Value Decomposition (or the essentially equivalent, Principal Component Analysis). Latent variable methods have particular advantages in explainability and model interpretation.

Reducing redundancy has a different motivation to extracting latent variables, however the same methodology may be applied. Where two or more measured variables are highly correlated they can be mathematically replaced by a smaller set of variables that are orthogonal (uncorrelated). This again can be a highly effective technique to represent complex data in just two or three dimensions in such a way that high-dimensional information is maximally preserved.

The Research Team



AI Forum NZ

The Artificial Intelligence Forum of New Zealand (Te Kāhui Atamai Iahiko o Aotearoa) aims to raise the level of awareness and capabilities of Artificial Intelligence in New Zealand.

The Forum brings together citizens, business, academia and the Government connecting, promoting and advancing the AI ecosystem to help ensure a thriving New Zealand.

The AI Forum designed, collated and edited the research.



IDC

IDC is a premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide.

IDC conducted the local and international case study research.



Adapt Research

Adapt Research provides bespoke research, analysis and writing services in health and technology. Adapt canvasses the evidence to help clients in the public or private sector answer strategic questions and produce evidence-based reports and recommendations to support decision making.



Antistatic

ANTISTATIC is a communications and research group based in San Francisco and Wellington. ANTISTATIC brings clarity to complex issues around technology and the environment, and helps amplify the voices of people driving positive social change.



Sapere

Sapere Research Group is one of the largest expert services firms in Australasia. Sapere provides independent expert testimony, strategic advisory services, data analytics and other advice to Australasia's private sector corporate clients, major law firms, government agencies and regulatory bodies.

Eltanin

Eltanin New Zealand assists organisations with AI technology strategies, roadmaps, and proof-of-concept studies. Building on 20 years experience applying machine learning and knowledge graph techniques to New Zealand primary sector science and industry.

Methodology

The AI Forum and our research partners have spent much of 2019 carrying out extensive desktop research, in-person interviews and dialogues with many AI Forum members, New Zealand AI practitioners, government representatives, academic institutions and ecosystem participants. We have also drawn on the contributions of AI Forum working groups, particularly in the sections on Talent, Skills and Capabilities.

We have attempted throughout to synthesise a comprehensive, balanced, objective and relevant view of AI in New Zealand which draws on the considerable amount of high quality local and international activity and places global perspectives into a New Zealand context.

Our research partner IDC carried out in-depth case study interviews to produce the case study library accompanying this report which has been supplemented in this report with case studies provided by AI Forum members and partners.

To obtain the information on research in New Zealand, we visited Government and institutional websites, approached funders for grant award information, and talked to a sample of institutional and industry researchers.

Corrections policy

We welcome feedback and are committed to correcting any errors in future editions. Please direct any comments on this report to our editorial team research@aiforum.org.nz.



Licence

Except where otherwise noted, all our work in this report is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. We hope you find it useful.

Original photos and images are copyright and licensed from their respective owners.

Logos are trademarks of their respective owners.

Our thanks to those authors and organisations who have licensed their shareable content which has been adapted including Nesta, Susan Etlinger at Altimeter, Gavin Starks, Open Data Institute, MIT Center for Collective Intelligence, XKCD, Department of Internal Affairs, New Zealand Data Futures Partnership, Ministry of Business, Innovation and Employment, Stats NZ.

References

1. 'The Iwi Algorithm – creating cultural capital with AI', Storicom press release on Scoop.co.nz. Retrieved July 3 from: <http://www.scoop.co.nz/stories/BU1903/S00565/the-ivi-algorithm-creating-cultural-capital-with-ai.htm>
2. Indigenous Protocols and Artificial Intelligence. Dragonfly Data Science Blog, 17 April 2019. Retrieved 3 July 2019 from: <https://www.dragonfly.co.nz/news/2019-04-17-caleb-hawaii.html>
3. The Journey to AI-enable Saas. Moor Insights and Strategy, January 2019. pg 3. Retrieved from: <https://www.purestorage.com/content/dam/pdf/en/white-papers/protected/wp-moor-journey-to-ai-enabled-sales.pdf>
4. The State of AI in 2019. The Verge, 28 January 2019. Retrieved from: <https://www.theverge.com/2019/1/28/18197520/ai-artificial-intelligence-machine-learning-computational-science>
5. Müller, Vincent C. and Bostrom, Nick (2014), 'Future progress in artificial intelligence: A Survey of Expert Opinion, in Vincent C. Müller (ed.), Fundamental Issues of Artificial Intelligence (Synthese Library; Berlin: Springer). <https://nickbostrom.com/papers/survey.pdf>
6. Toby Walsh (2017), Expert and Non-Expert Opinion about Technological Unemployment, UNSW / Data61, <https://arxiv.org/ftp/arxiv/papers/1706/1706.06906.pdf>
7. Elon Musk: artificial intelligence is our biggest existential threat. The Guardian, 27 October 2014. Retrieved 31 March 2019 from: <https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>
8. Advanced Business Manager Pty Ltd, <https://www.advancedbusinessmanager.co.nz/nz/>
9. Better Language Models and their Implications. OpenAI, 14 February 2019. Retrieved 1 April 2019 from <https://openai.com/blog/better-language-models/>
10. Weapons of Math Destruction. Cathy O'Neil, 2016. Page 7. <https://algorithmwatch.org/en/project/ai-ethics-guidelines-global-inventory/>
11. What is personal information and the Privacy Act? data.govt.nz, 2018. Retrieved 29 April 2019 from: <https://www.data.govt.nz/manage-data/privacy-and-security/what-is-personal-identifiable-information-and-the-privacy-act/>
12. What is an artificial neural network? Digital Trends, 5 January 2019. Retrieved 29 April 2019 from: <https://www.digitaltrends.com/cool-tech/what-is-an-artificial-neural-network/>
13. The Hundred-Page Machine Learning Book. Andriy Burkov, 2019. Page 20.
14. The State of AI in 2019: Divergence. MMC Ventures, 2019. Page 20. Downloaded from: <https://www.stateofai2019.com/>
15. AI Day 2019 keynote speech (Auckland, New Zealand). Sean Gourley, 28 March 2019. View at https://www.youtube.com/watch?v=_wp_a-oyZyc
16. Watson Machine Learning Accelerator. IBM. Retrieved from: <https://www.ibm.com/nz-en/marketplace/deep-learning-platform>
17. 18 AI and the cloud – a match made in heaven. IT Brief, May 2018. <https://itbrief.co.nz/story/ai-and-cloud-match-made-heaven>
18. AI-fueled organizations: Reaching AI's full potential in the enterprise. Deloitte Insights, January 2019. <https://www2.deloitte.com/insights/us/en/focus/tech-trends/2019/driving-ai-potential-organizations.html>
19. Catalyst website, <https://www.catalyst.net.nz/>
20. Artificial intelligence for language recognition. <https://catalystcloud.nz/customers/dragonfly-data-science/>
21. NIWA invests \$18 million in New Zealand's future. NIWA, media release, June 2017. <https://www.niwa.co.nz/news/niwa-invests-18million-in-new-zealands-future>
22. Services | High Performance Computing Facility, NIWA <https://www.niwa.co.nz/our-services/high-performance-computing-facility>
23. <http://www.bluefern.canterbury.ac.nz/>
24. <https://www.nesi.org.nz/>
25. The IP behind the AI boom, WIPO Magazine, Feb 2019. https://www.wipo.int/wipo_magazine/en/2019/01/article_0001.html
26. The IP behind the AI boom, WIPO Magazine, Feb 2019. https://www.wipo.int/wipo_magazine/en/2019/01/article_0001.html
27. <https://cloud.google.com/vision/>
28. <https://cloud.google.com/video-intelligence/>
29. <https://azure.microsoft.com/en-us/services/cognitive-services/>
30. Microsoft launches a drag-and-drop machine learning tool, Tech Crunch, May 2019, <https://techcrunch.com/2019/05/02/microsoft-launches-a-drag-and-drop-machine-learning-tool-and-hosted-jupyter-notebooks/>
31. <https://aws.amazon.com/textract/>
32. <https://deeptmind.com/blog/capture-the-flag-science/>
33. <https://openai.com/five/>
34. Deep Learning Chipsets, Tractica, May 2019, <https://www.tractica.com/research/deep-learning-chipsets/>
35. <https://ai.google/research/teams/applied-science/quantum-ai/>
36. Qualcomm aims to take on Nvidia, Intel with new AI chips. Reuters. Apr 2019. <https://www.reuters.com/article/us-qualcomm-ftc/qualcomm-aims-to-take-on-nvidia-intel-with-new-ai-chips-idUSKCN1RM04V>
37. 2019 Tech Trends Report. Future Today Institute. 2019. <https://futuretodayinstitute.com/2019-tech-trends/>
38. 2019 Tech Trends Report. Future Today Institute. 2019 <https://futuretodayinstitute.com/2019-tech-trends/>

40. Google's former China head Kai-Fu Lee retreats from U.S. investing amid trade dispute. CNBC. May 2019. <https://finance.yahoo.com/news/google-apos-former-china-head-162741745.html>
41. OECD (2018), "Private Equity Investment in Artificial Intelligence", OECD Going Digital Policy Note, OECD, Paris, www.oecd.org/going-digital/ai/private-equity-investment-in-artificial-intelligence.pdf
42. <https://www.cbinsights.com/research/ai-unicorn-club/>
43. According to a joint PwC and CB Insights report (<https://www.cbinsights.com/research/report/venture-capital-q4-2018/>), venture capital funding of AI companies grew 72 percent in 2018, hitting a record \$9.3 billion. This continued growth comes after three years of steadily increasing investment, with an average annual increase of 28% from 2015 to 2017. (The recent boost partly reflects the frothy funding environment overall: venture investing in the U.S. hit \$99.5 billion in 2018, the highest level since 2000.)
44. However, the article concluded that the dynamics of this AI investment "bubble" appears relatively benign, having more in common with the dot.com bubble, which helped finance the internet backbone, than the housing bubble, which wreaked havoc on the household finances of millions of homeowners in the US and around the world <https://sloanreview.mit.edu/article/learning-to-love-the-ai-bubble/>
45. Richard Dellabarca, NZVIF CEO, Interview with RNZ Business 8 July 2019 <https://www.rnz.co.nz/national/programmes/businessnews/audio/2018703091/business-news-for-8-july-2019>
- 45a. <https://nextbigwhat.com/us-embargo-export-ai-products/>
46. <https://rhg.com/research/two-way-street-2019-update-us-china-direct-investment-trends/>
47. New Zealand Private Equity and Venture Capital Monitor 2018 Full year review, NZVCA/EY, May 2018, https://www.nzvca.co.nz/wp-content/uploads/2018/05/FINAL_NZ-PE-VC-Monitor-2018_Full-year-review_FINALWEB.pdf
- 47a
48. <https://www.whitehouse.gov/ai/>
49. <https://futureoflife.org/ai-policy-united-states/>
50. <https://futureoflife.org/ai-policy-china/> - see also Full Translation: China's 'New Generation Artificial Intelligence Development Plan' (2017), Graham Webster, Rogier Creemers, Paul Triolo, and Elsa Kania, New America, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>
51. <https://em.dk/english/news/2018/01-30-new-strategy-to-make-denmark-the-new-digital-frontrunner>
52. <https://www.gov.sg/news/content/channel-newsasia---singapores-artificial-intelligence-capabilities-to-get-s150m-boost>
53. <https://ai.taiwan.gov.tw/news/cabinet-plans-to-develop-the-nations-ai-industry/>
54. <https://www.computerworld.com.au/article/640926/budget-2018-government-seeks-boost-australian-ai-capabilities/>
55. <https://www.cifar.ca/ai/pan-canadian-artificial-intelligence-strategy>
56. <https://www.msit.go.kr/web/msipContents/contentsView.do?cateld=mssw311&artld=1382727>
57. <https://www.gov.uk/government/publications/artificial-intelligence-sector-deal/ai-sector-deal>
58. <https://www.gouvernement.fr/en/artificial-intelligence-making-france-a-leader>
59. https://www.ki-strategie-deutschland.de/home.html?file=files/downloads/Nationale_KI-Strategie_engl.pdf
60. https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=51625
61. <https://www.idc.com/getdoc.jsp?containerId=prUS44911419>
62. https://www.researchandmarkets.com/research/w7fgk7/the_global?w=4
63. <https://www.tractica.com/research/artificial-intelligence-market-forecasts/>
64. <https://www.marketsandmarkets.com/PressReleases/artificial-intelligence.asp>
65. Artificial Intelligence The Next Digital Frontier?, McKinsey Global Institute, June 2017, <https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx>
66. https://www.nzherald.co.nz/nz/news/article.cfm?c_id=1&objectid=11923979
67. <https://www.computerworld.com.co.nz/article/620747/ai-diagnose-sick-cows/>
68. Kleinberg et al. (2018). Human Decisions and Machine Predictions. Available at: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5947971/>
69. <https://www.scmp.com/news/hong-kong/hong-kong-law-and-crime/article/2158400/new-supercomputer-helps-hong-kong-customs>
70. <https://ai.googleblog.com/2017/03/assisting-pathologists-in-detecting.html>
71. <https://ai.googleblog.com/2018/12/improving-effectiveness-of-diabetic.html>
72. <https://news.microsoft.com/apac/features/artificial-intelligence-takes-on-ocean-trash-cleaning-up-the-worlds-beaches-with-the-help-of-data/> and refer to the Sustainable Coastlines case study in this report.
73. <https://ai.googleblog.com/2019/03/a-summary-of-google-flood-forecasting.html>
74. <https://www.2030vision.com/news/2030vision-uniting-to-deliver-technology-for-the-global-goals>
75. <https://ai.google/social-good>
76. <https://globalfishingwatch.org/map/>

77. <https://www.microsoft.com/en-us/ai/ai-for-earth>
78. Tackling Climate Change with Machine Learning, Various Authors, June 2019 <https://arxiv.org/pdf/1906.05433.pdf>
79. Other than headline grabbing calls by high profile personalities such as Elon Musk to "regulate AI before it's too late" which invoke the "existential threat" of phantom superintelligence
80. Calo, R. 2017. Artificial Intelligence Policy: A Primer and Roadmap. Available at: <https://ssrn.com/abstract=3015350>
81. <https://becominghuman.ai/the-regulation-of-artificial-intelligence-a-case-study-of-the-partnership-on-ai-ctc22526c19f>
82. <https://www.quora.com/Should-artificial-intelligence-be-regulated/answer/Xavier-Amatriain>
83. <https://www.loc.gov/law/help/artificial-intelligence/index.php>
84. <https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html>
85. <https://www.cnn.com/2019/07/17/tech/cities-ban-facial-recognition/index.html>
86. <https://inews.co.uk/news/technology/alphabet-google-chairman-tech-firms-self-regulate-government-too-slowly/>
87. AI Now Report 2018, AI Now Institute, December 2018, https://ainowinstitute.org/AI_Now_2018_Report.pdf
88. <https://anz.co.nz/resources/5/e/5e24f728-9426-45d8-9072-b83a5a47f86b/From-term-deposits-to-turtles.pdf>
89. <https://www.stuff.co.nz/business/business-top-stories/108681909/anzs-artificial-intelligence-entity-jamie-took-12000-customer-inquiries-in-its-first-100-days>
90. <https://www.forbes.com/sites/zakdoffman/2019/05/03/china-new-data-breach-exposes-facial-recognition-and-ethnicity-tracking-in-beijing/#38160a8634a7>
91. <https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html>
92. <https://www.nature.com/immersive/d41586-019-01398-x/index.html>
93. Principled Artificial Intelligence, A Map of Ethical and Rights-Based Approaches, Fjeld et al, Draft, July 2019 <https://ai-hr.cyber.harvard.edu/primp-viz.html>
94. Dave Heiner, Microsoft, AI-DAY 2019 view at <https://youtu.be/IFrUjIDuF5U>
95. Principles set to regulate AI research, applications, China.org.cn, June 18 2019 http://www.china.org.cn/china/2019-06/18/content_74895077.htm
96. <https://www.vox.com/future-perfect/2019/4/4/18295933/google-cancels-ai-ethics-board>
- 96a. <https://www.montrealdeclaration-responsibleai.com/>
97. Report of the Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression. United Nations. August 2018. <https://undocs.org/A/73/348>
98. <https://techvibes.com/2019/03/08/canadas-new-federal-directive-makes-ethical-ai-a-national-issue>
99. <https://www.fastcompany.com/90172734/this-logo-is-like-an-organic-sticker-for-algorithms-that-arent-evil>
100. Our work, Partnership on AI website, <https://www.partnershiponai.org/about/#our-work>
101. <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>
102. <https://legalinstruments.oecd.org/api/print?ids=648&lang=en>
103. <https://www.japantimes.co.jp/news/2019/06/08/business/g20-ministers-kick-talks-trade-digital-economy-ibaraki-prefecture/#.XQoK1Fwzbb0>
104. <https://productivity.govt.nz/inquiry-content/3960?stage=2>
105. <https://www.data.govt.nz/about/government-chief-data-steward-gcdfs/new-data-advisory-group-expression-of-interest-now-open/>
106. <https://www.digital.govt.nz/digital-government/strategy/progress-towards-a-digital-government-strategy/>
107. <https://www.data.govt.nz/assets/Uploads/Algorithm-Assessment-Report-Oct-2018.pdf>
108. <https://sia.govt.nz/investing-for-social-wellbeing/what-you-told-us/introduction/>
109. https://www.sftichallenge.govt.nz/sites/default/files/2018-09/Ma%CC%84ori_Data_Futures_Report.pdf
110. <https://www.mbie.govt.nz/business-and-employment/economic-development/digital-economy/ministerial-advisory-group/>
111. <https://www.mbie.govt.nz/business-and-employment/economic-development/industry-policy>
112. <https://www.digitalskillsforum.nz/a-heartfelt-thank-you-to-all-involved-in-the-digital-skills-hui-2019/>
113. <https://www.mbie.govt.nz/science-and-technology/science-and-innovation/funding-information-and-opportunities/investment-funds/strategic-science-investment-fund/funded-infrastructure/data-science-investment/>
114. <https://www.perc-inc.nz/>
- 114a <https://www.lawfoundation.org.nz/?p=10082>
115. <https://www.brainbox.institute/report>
116. <https://aiforum.org.nz/2018/10/12/explore-new-zealands-ai-ecosystem/>
117. <https://tin100.com/tinwire-tech-news-roundup-29-april-19/>
118. <https://www.ai-day.com/>
119. <https://futureoflife.org/lethal-autonomous-weapons-pledge/>
120. <https://idealogy.co.nz/tech/2017/08/new-myob-report-new-zealands-start-scene-thriving-needs-help-truly-thrive>
121. <https://www.pwc.co.nz/pdfs/2019pdfs/startup-magazine-2019.pdf>
122. <https://www.nzvif.co.nz/>

123. <https://www.stuff.co.nz/business/112029338/sir-stephen-tindall-simplicity-the-icehouse-and-100-million-for-startup-companies>
124. https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12150765
125. https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12167379
126. https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12155304
127. https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12160371
128. https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12209262
129. https://www.shareholder.anz.com/sites/default/files/190501_2019_half_year_result_-_for_disclosure_-_final.pdf
130. <https://www.iag.com.au/sites/default/files/Documents/Results%20%26%20reports/FY18-management-presentation-slides.pdf>
131. <http://investors.sparknz.co.nz/Investor-Centre/?page=Announcements>
132. <http://www.scoop.co.nz/stories/PA1805/S00031/government-will-move-quickly-on-ai-action-plan.htm>
133. <https://aiforum.org.nz/our-work/working-groups/wg1-forging-a-coordinated-ai-strategy-for-new-zealand/>
134. https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12130986
135. <https://www.mbie.govt.nz/science-and-technology/science-and-innovation/funding-information-and-opportunities/investment-funds/strategic-science-investment-fund/funded-infrastructure/data-science-investment/>
136. <https://www.gov.uk/government/publications/artificial-intelligence-sector-deal/ai-sector-deal>
137. https://www.washingtonpost.com/climate-environment/2019/05/06/one-million-species-face-extinction-un-panel-says-humans-will-suffer-result/?utm_term=.a350f71399bc
138. <https://treasury.govt.nz/information-and-services/nz-economy/living-standards/our-living-standards-framework>
139. <https://www.budget.govt.nz/budget/2018/economic-fiscal-outlook/budget-2019-focus-on-wellbeing.htm>
140. Another notable international effort to go beyond GDP is GDP-B by MIT's Prof Erik Brynjolfsson <https://qz.com/1582202/mits-erik-brynjolfsson-is-redesigning-gdp-for-the-21st-century/>
141. <https://treasury.govt.nz/publications/tp/living-standards-framework-introducing-dashboard-html#section-4>
142. <https://www.stats.govt.nz/indicators-and-snapshots/indicators-aotearoa-new-zealand-nga-tutohu-aotearoa>
143. <https://wellbeingindicators.stats.govt.nz/>
144. <https://treasury.govt.nz/sites/default/files/2019-01/dp19-01.pdf> page ii
145. <https://treasury.govt.nz/sites/default/files/2019-01/dp19-01.pdf> page 5
146. From page 6: https://www.sftichallenge.govt.nz/sites/default/files/2018-09/Ma%CC%84ori_Data_Futures_Report.pdf People's Report on the 2030 Agenda and Sustainability Development Goals, <https://www.sdg.org.nz/peoples-report/>
147. The People's Report on the 2030 Agenda and Sustainable Development Goals – 2019, [sdg.org.nz](https://www.sdg.org.nz), 2019, <https://www.sdg.org.nz/peoples-report/>
148. <https://www.accenture.com/nz-en/insight-artificial-intelligence-future-growth>
149. <https://www.accenture.com/nz-en/insight-artificial-intelligence-future-growth>
150. Haskel, Jonathan; Westlake, Stian. *Capitalism without Capital: The Rise of the Intangible Economy* (p. 7). Princeton University Press. Kindle Edition.
151. Rogers, Everett; 2003, *Diffusion of Innovations* 5th Edition, Free Press; Kindle Edition
152. *Technology Adoption in US Households*, Our World In Data, <https://ourworldindata.org/technology-adoption>
153. <https://www.pwc.co.uk/economic-services/assets/macro-economic-impact-of-ai-technical-report-feb-18.pdf>
154. <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>
155. <https://www.stats.govt.nz/information-releases/gross-domestic-product-march-2019-quarter>
156. AI Forum NZ, *Artificial Intelligence Shaping a Future New Zealand*, 2018 p.52 https://aiforum.org.nz/wp-content/uploads/2018/07/AI-Report-2018_web-version.pdf
157. Christodoulakis C, et al. *Barriers to Adoption of Information Technology in Healthcare*. 2016. http://www.cs.toronto.edu/~christina/documents/ACM_CASCON2017.pdf
158. <https://edtechmagazine.com/higher/article/2016/05/higher-education-technology-paradox>
159. AI Forum NZ, *Artificial Intelligence Shaping a Future New Zealand*, 2018 https://aiforum.org.nz/wp-content/uploads/2018/07/AI-Report-2018_web-version.pdf
160. Dominik Hartmann, Miguel R. Guevara, Cristian Jara-Figueroa, Manuel Arístarán, César A. Hidalgo, *Linking Economic Complexity, Institutions, and Income Inequality* <http://www.sciencedirect.com/science/article/pii/S0305750X15309876>
161. Calo, Ryan, *Artificial Intelligence Policy: A Primer and Roadmap* (August 8, 2017). Available at SSRN: <https://ssrn.com/abstract=3015350> or <http://dx.doi.org/10.2139/ssrn.3015350>
162. <https://www.cifar.ca/cifarnews/2018/12/06/building-an-ai>

- world-report-on-national-and-regional-ai-strategies
163. <https://www.whitehouse.gov/ai/>
 164. <https://www.whitehouse.gov/articles/accelerating-americas-leadership-in-artificial-intelligence/>
 165. <https://www.scmp.com/business/china-business/article/2115935/chinas-xi-jinping-highlights-ai-big-data-and-shared-economy>
 166. <https://epicproject.eu/>
 167. <https://cordis.europa.eu/news/rcn/131729/en>
 168. <https://legalinstruments.oecd.org/api/print?id=648&lang=en>
 169. <http://www.oecd.org/going-digital/ai/oecd-initiatives-on-ai.htm>
 170. <https://www.acc.co.nz/assets/im-injured/ef79338f63/claims-approval-technical-summary.pdf>
 171. <https://www.acc.co.nz/assets/im-injured/ef79338f63/claims-approval-technical-summary.pdf>
 172. <https://www.hinz.org.nz/news/435335/Algorithm-automates-ACCs-claims-process.htm>
 173. Sustainable Coastlines, A Global Litter Data & Solutions Platform, https://docs.google.com/presentation/d/1erolGpiKsdwQa7UUFc3zOojrA7N-ljALVdn7bpqrv4/edit#slide=id.g528c974e09_0_352
 174. <https://www.turing.ac.uk/research/research-programmes/public-policy>
 175. <https://www.health.govt.nz/your-health/conditions-and-treatments/diseases-and-illnesses/breast-cancer>
 176. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4379114/>
 177. <https://www.health.govt.nz/system/files/documents/publications/new-zealand-cancer-plan-2015-2018-dec14-v3.pdf>
 178. <http://news.mit.edu/2019/using-ai-predict-breast-cancer-and-personalize-care-0507>
 179. <https://www.ncbi.nlm.nih.gov/pubmed/30834436>
 180. <https://arxiv.org/abs/1606.05718>
 181. <https://www.noted.co.nz/tech/how-artificial-intelligence-is-powering-efforts-to-detect-breast-cancer/>
 182. Topol, E. (2019). Deep Medicine. Basic Books.
 183. <https://phys.org/news/2018-11-artificial-intelligence-highway-safety-las.html>
 184. <https://www.stuff.co.nz/business/industries/105069920/whats-causing-new-zealands-fatal-crashes>
 185. <https://www.transport.govt.nz/assets/Import/Uploads/Research/Documents/b67f729bf5/Social-cost-of-road-crashes-and-injuries-2018-update.pdf>
 186. <https://www.doc.govt.nz/news/media-releases/2017/wilding-pines-control-work-nears-million-hectare-mark/>
 187. <https://www.mbie.govt.nz/assets/436ecb3be9/strategic-science-investment-fund-investment-plan.pdf>
 188. https://www.cifar.ca/docs/default-source/ai-society/buildinganaiworld_eng.pdf
 189. <https://www.mbie.govt.nz/assets/436ecb3be9/strategic-science-investment-fund-investment-plan.pdf>
 190. WIPO Magazine, The IP behind the AI Boom, February 2019, https://www.wipo.int/wipo_magazine/en/2019/01/article_0001.html
 191. <https://www.fool.com.au/2019/04/30/quarterly-update-is-the-volpara-share-price-in-the-buy-zone/>
 192. AI Forum NZ, Artificial Intelligence Shaping a Future New Zealand, 2018 https://aiforum.org.nz/wp-content/uploads/2018/07/AI-Report-2018_web-version.pdf
 193. <https://www.nature.com/news/ai-talent-grab-sparks-excitement-and-concern-1.19821>
 194. <https://royalsociety.org.nz/what-we-do/funds-and-opportunities/marsden/awarded-grants/>
 195. <https://www.mbie.govt.nz/science-and-technology/science-and-innovation/funding-information-and-opportunities/investment-funds/>
 196. <https://www.mbie.govt.nz/science-and-technology/science-and-innovation/funding-information-and-opportunities/investment-funds/>
 197. <https://www.sftichallenge.govt.nz/>
 198. https://www.sftichallenge.govt.nz/sites/default/files/2018-12/Future%20Strategy%20SfTI%2020%20June%202018.FINAL_November.pdf
 199. <https://www.lincolnagritech.co.nz/news/workforce-robots-on-the-horizon/>
 200. <https://www.cawthron.org.nz/coastal-freshwater/news/2017/precision-farming-aquaculture-project-announced/>
 201. <https://www.mbie.govt.nz/science-and-technology/science-and-innovation/funding-information-and-opportunities/investment-funds/regional-research-institutes/>
 202. <https://www.xerra.nz/>
 203. <https://www.esr.cri.nz/our-services/consultancy/forensic-r-and-d/>
 204. <https://www.esr.cri.nz/assets/Uploads/ESR-SCI-2018-2023-WEB.pdf>
 205. <https://www.agresearch.co.nz/assets/Uploads/AgResearch-Annual-Report-2018-for-web.pdf>
 206. <https://whatonearth.co.nz/big-data>
 207. <https://www.aut.ac.nz/aut-artificial-intelligence-initiative/ai-research-projects-around-aut>
 208. <https://csda.aut.ac.nz/news/talking-predictive-risk-modelling-in-denver-rhema-vaithianathan-at-the-2018-alliance-conference>
 209. <http://www.lincoln.ac.nz/Research/Research/RC/CSBII/?sti=5>
 210. https://www.massey.ac.nz/massey/explore/research/artificial-intelligence/artificial-intelligence_home.cfm

211. <https://www.cs.auckland.ac.nz/en/about/ourresearch/researchgroups.html#3c6b7220fabd18a0415571b8d0691609>
212. <https://www.stuff.co.nz/technology/107776015/robots-to-train-vineyard-and-orchard-workers-to-do-a-better-job>
213. <https://unidirory.auckland.ac.nz/profile/m-witbrock>
214. <https://www.victoria.ac.nz/research/strengths/research-focus/artificial-intelligence> and <http://ecs.victoria.ac.nz/Groups/ECRG/>
215. <https://royalsociety.org.nz/research/the-future-of-artificial-intelligence-evolutionary-computation/>
216. <https://www.cs.waikato.ac.nz/~ml/>
217. <https://www.capterra.com/machine-learning-software/compare/170400-171057-171134/Azure-Machine-Learning-Studio-vs-TensorFlow-vs-Weka>
218. <https://thoughtexperiment.co.nz/2018/07/31/nz-merino-using-artificial-intelligence-to-monitor-sheep-well-being/>
219. <https://www.callaghaninnovation.govt.nz/grants>
220. <https://rd.callaghaninnovation.govt.nz/data-and-iot/data-science/data-discovery/>
221. <https://www.mbie.govt.nz/assets/9c551e0eb9/marsden-fund-assessment.pdf>
222. <https://www.mbie.govt.nz/assets/436ecb3be9/strategic-science-investment-fund-investment-plan.pdf>
223. <https://www.mbie.govt.nz/science-and-technology/science-and-innovation/funding-information-and-opportunities/investment-funds/strategic-science-investment-fund/funded-infrastructure/data-science-investment>
224. <https://minterellison.co.nz/our-view/joint-venture-to-drive-artificial-intelligence-in-legal-services>
225. <https://aigrant.org/>
226. <https://www.openphilanthropy.org/focus/global-catastrophic-risks/potential-risks-advanced-artificial-intelligence>
227. <https://www.nesi.org.nz/>
228. <https://www.mbie.govt.nz/assets/9c551e0eb9/marsden-fund-assessment.pdf>
229. <http://www.montrealintechology.com/elementai-raises-historic-137-5-million-series-a-round/>
230. <https://nouvelles.umontreal.ca/en/article/2017/05/15/100m-for-a-new-quebec-research-cluster-in-ai/>
231. <https://www.businessinsider.com.au/google-is-backing-a-150-million-artificial-intelligence-institute-in-toronto-2017-3?r=US&IR=T>
232. <https://www.callaghaninnovation.govt.nz/grants/rd-tax-incentive>
233. https://www.parliament.nz/en/pb/bills-and-laws/bills-proposed-laws/document/BILL_80869/taxation-research-and-development-tax-credits-bill
234. Why robots will grow your food in the future, TechWeek'18. TechWeek NZ YouTube video, 15 January 2018. Retrieved from: <https://www.youtube.com/watch?v=rIK3-46KCJ8>
235. Robotics Plus wins innovation award for apple packer. The Packer, 18 June 2019. Retrieved from: <https://www.thepacker.com/article/robotics-plus-wins-innovation-award-apple-packer>
236. Tauranga robotics firm produces industry-changing automated log scaler. Stuff.co.nz, 6 June 2019. Retrieved from: <https://www.stuff.co.nz/business/113259538/tauranga-robotics-firm-produces-industry-changing-automated-log-scaler>
237. Why robots will grow your food in the future, TechWeek'18. TechWeek NZ YouTube video, 15 January 2018. Retrieved from: <https://www.youtube.com/watch?v=rIK3-46KCJ8>
238. Robotics Plus and its agricultural robotics recognised at technology awards. Robotics & Automation News, 26 June 2018. Retrieved from: <https://roboticsandautomationnews.com/2019/06/26/robotics-plus-and-its-agricultural-robotics-recognised-at-technology-awards/23960/>
- 238a Yamaha sinks \$12m into Kiwi ag-tech company Robotics Plus. Stuff.co.nz, 22 November 2018. Retrieved from: <https://www.stuff.co.nz/business/farming/108782227/yamaha-sinks-12m-into-kiwi-agtech-company-robotics-plus>
239. Open data, transparency and power – role of the Virtual Health Information Network. Teng et al. 2017. <https://blogs.otago.ac.nz/pubhealthexpert/2017/06/14/open-data-transparency-and-power-role-of-the-virtual-health-information-network/>
240. <https://catalogue.data.govt.nz/dataset/kiwifruit-power-analyses>
241. <https://www.turing.ac.uk/>
242. <https://www.domo.com/learn/data-never-sleeps-5>
243. <https://www.mmventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf>
244. <https://www.forbes.com/sites/bernardmarr/2017/06/09/why-ai-would-be-nothing-without-big-data/#846239a4f6d0>
245. <https://nationalinterest.org/blog/buzz/america-and-chinas-great-ai-competition-what-driving-it-48677>
246. <https://www.cio.com/article/3309058/5-ways-industrial-ai-is-revolutionizing-manufacturing.html>
247. <https://www.amazon.com/Age-Surveillance-Capitalism-Future-Frontier/dp/1610395697>
248. <https://www.data.govt.nz/assets/Uploads/data-strategy-and-roadmap-dec-18.pdf>
249. <https://freedex.org/wp-content/blogs.dir/2015/files/2018/10/AI-and-FOE-GA.pdf>
250. <https://freedex.org/wp-content/blogs.dir/2015/files/2018/10/AI-and-FOE-GA.pdf>
251. <https://www.technologyreview.com/s/612588/its-time-for-a-bill-of-data-rights/>
252. <https://techcrunch.com/2019/02/10/privacy-is-a-commons/>
253. <https://www.amazon.com/Age-Surveillance-Capitalism-Future-Frontier/dp/1610395697>

254. https://www.washingtonpost.com/opinions/mark-zuckerberg-the-internet-needs-new-rules-lets-start-in-these-four-areas/2019/03/29/9e6f0504-521a-11e9-a3f7-78b7525a8d5f_story.html?utm_term=.db62c4e285da
255. <https://blogs.microsoft.com/on-the-issues/2019/05/20/gdprs-first-anniversary-a-year-of-progress-in-privacy-protection/>
256. <https://www.amazon.com/Age-Surveillance-Capitalism-Future-Frontier/dp/1610395697>
257. <https://colonizedbydata.com/>
258. https://www.parliament.nz/resource/en-NZ/52SCJU_ADV_77618_JU65196/afe12605c1393eccc42afec3340d35a73eff9062
259. <https://www.privacy.org.nz/the-privacy-act-and-codes/privacy-law-reform/privacy-commissioner-s-review-of-the-privacy-act/>
260. <https://www.privacy.org.nz/blog/welcoming-the-privacy-bill/>
261. <https://www.kensingtonswan.com/news-updates-and-events/you-call-that-a-fine-this-is-a-fine/>
262. <https://www.amcham-shanghai.org/en/article/data-governance-qa-samm-sacks>
263. <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/china-data-governance-regime-timeline/>
264. <https://www.amcham-shanghai.org/en/article/data-governance-qa-samm-sacks>
265. <https://futuretodayinstitute.com/2019-tech-trends/>
267. <https://www.cio.co.nz/article/659521/does-new-zealand-copyright-regime-need-updating-digital-age/>
268. <https://www.scmp.com/tech/innovation/article/2174564/facial-recognition-catches-chinas-air-con-queen-dong-mingzhu>
269. <https://www.wired.co.uk/article/china-social-credit-system-explained>
270. https://www.washingtonpost.com/politics/2019/03/21/what-do-people-china-think-about-social-credit-monitoring/?noredirect=on&utm_term=.59a1838f3ede
271. <http://www.fao.org/e-agriculture/news/webinar-farmersrights-data-and-ownership-issues>
272. <https://agfundernews.com/congress-wades-farm-data-ownership-debate.html>
273. <https://www.agdatatransparent.com/the-privacy-and-security-principles-for-farm-data/>
274. <http://www.farmdatacode.org.nz/>
275. <https://medium.com/@aglawpaul/farm-data-ownership-and-protection-b872441f53ba>
276. Wachter and Mittelstadt (2019). <https://robotic.legal/wp-content/uploads/2018/09/SSRN-id3248829.pdf>
277. <https://www.theguardian.com/technology/2019/jan/20/shoshana-zuboff-age-of-surveillance-capitalism-google-facebook>
278. <https://www.businessinsider.com.au/the-incredible-story-of-how-target-exposed-a-teen-girls-pregnancy-2012-2?r=US&IR=T>
279. <https://www.data.govt.nz/assets/Uploads/data-strategy-and-roadmap-dec-18.pdf>
280. <https://trusteddata.co.nz/wp-content/uploads/2017/08/Background-Trusted-Data.pdf> NOTE: DFP is ceased - the website content etc will be transitioned
281. http://archive.stats.govt.nz/browse_for_stats/snapshots-of-nz/integrated-data-infrastructure.aspx
282. http://europa.eu/rapid/press-release_IP-19-1893_en.htm
283. <https://static1.squarespace.com/static/58e9b10f9de4bb8d1fb5ebbc/t/5913020d15cf7dde1df34482/1494417935052/Te+Mana+Raraunga+Charter+%28Final+%26+Approved%29.pdf>
284. Tahu Kukutai, in https://www.sftichallenge.govt.nz/sites/default/files/2018-09/Ma%CC%84ori_Data_Futures_Report.pdf
285. <https://www.waikato.ac.nz/news-opinion/media/2019/from-populations-to-peoples>
286. <http://www.aigi.com.au/wp-content/uploads/2018/07/Communique-Indigenous-Data-Sovereignty-Summit.pdf>
287. <http://press-files.anu.edu.au/downloads/press/n2140/pdf/book.pdf#page=25>
288. <https://www.tpk.govt.nz/docs/teritorito/tpk-drtahukukutai-2017.pdf>
289. <https://www.privacy.org.nz/assets/Uploads/Principles-for-the-safe-and-effective-use-of-data-and-analytics-guidance3.pdf>
290. <https://www.otago.ac.nz/caipp/index.html>
292. https://www.sftichallenge.govt.nz/sites/default/files/2018-09/Ma%CC%84ori_Data_Futures_Report.pdf
293. <https://www2.deloitte.com/content/dam/Deloitte/nl/Documents/risk/deloitte-nl-risk-gdpr-vision-approach.pdf>
294. <http://fortune.com/2019/01/21/france-fines-google-57-million-for-gdpr-violations/>
295. <https://gdpr.report/news/2019/05/20/big-tech-firms-preparing-their-legal-teams-for-data-breach-challenges-reports-reveal/>
296. <https://www.cnet.com/news/microsoft-wants-a-us-privacy-law-that-puts-the-burden-on-tech-companies/>
297. <https://www.privacy.org.nz/further-resources/knowledge-base/view/482>
298. <https://www.simpsongrierson.com/articles/2018/gdpr-checklist-for-new-zealand-organisations>
299. <https://www.gov.uk/government/publications/unlocking-digital-competition-report-of-the-digital-competition-expert-panel>
300. https://www.researchgate.net/publication/318177824_Personal_control_of_privacy_and_data_Estonian_experience

- 300a <https://decodeproject.eu/>
301. <https://www.nesta.org.uk/blog/new-ecosystem-trust/>
302. <https://theodi.org/article/huge-appetite-for-data-trusts-according-to-new-odi-research/>
303. <https://theodi.org/>
304. <https://theodi.org/article/huge-appetite-for-data-trusts-according-to-new-odi-research/>
305. <https://theodi.org/wp-content/uploads/2019/04/Datatrusts-economicfunction.pdf>
306. <https://theodi.org/wp-content/uploads/2019/04/General-legal-report-on-data-trust.pdf>
307. <https://www.gov.uk/government/publications/unlocking-digital-competition-report-of-the-digital-competition-expert-panel>
308. <https://www.data.govt.nz/assets/Uploads/fact-sheet-data-roadmap-12422-oct-18.pdf>
309. <https://www.data.govt.nz/assets/Uploads/fact-sheet-data-roadmap-12422-oct-18.pdf>
310. <https://www.mmventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf>
311. <https://www.data.govt.nz/assets/Uploads/data-strategy-and-roadmap-dec-18.pdf>
312. <https://www.partnershiponai.org/the-partnership-on-ai-launches-multistakeholder-initiative-to-enhance-machine-learning-transparency/>
313. <https://www.forbes.com/sites/forbestechcouncil/2018/10/26/how-do-we-address-the-reproducibility-crisis-in-artificial-intelligence/#36fe1cf17688>
314. <https://transparencee.org/analysis/data-standards-what-are-they-and-why-do-they-matter/>
315. <http://dimensionsofdataquality.com/thestandard>
316. <http://dimensionsofdataquality.com/thestandard>
317. <http://dimensionsofdataquality.com/>
318. <https://www.health.govt.nz/our-work/ehealth/digital-health-sector-architecture-standards-and-governance/health-information-standards>
319. <https://www.stats.govt.nz/assets/Uploads/Data-leadership-fact-sheets/fact-sheet-data-standards-feb-2018.pdf>
320. <https://www.stats.govt.nz/assets/Uploads/Data-leadership-fact-sheets/Principles-safe-and-effective-data-and-analytics-May-2018.pdf>
321. <http://www.farmdatastandards.org.nz/data-standards/>
322. <https://www.cio.co.nz/article/658398/do-your-homework-first-data-algorithms-human-checkpoints/>
323. <https://www.cio.co.nz/article/658398/do-your-homework-first-data-algorithms-human-checkpoints/>
324. http://archive.stats.govt.nz/browse_for_stats/economic_indicators/productivity/measuring-govt-productivity/8-data-availability-health-care.aspx
325. <https://www.data.govt.nz/assets/Uploads/data-strategy-and-roadmap-dec-18.pdf>
326. <https://hbr.org/2019/01/which-countries-are-leading-the-data-economy>
327. <https://www.xerra.nz/>
328. <https://www.data.govt.nz/open-data/open-government-data-programme/open-data-nz/>
329. <https://sia.govt.nz/>
330. <https://sia.govt.nz/how-we-can-help/data-exchange/>
331. <https://www.privacy.org.nz/assets/Files/Reports/2017-04-04-Inquiry-into-MSD-collection-of-individual-client-level-data-OPC-report.pdf> and https://www.nzherald.co.nz/news/article.cfm?c_id=1&objectid=11832905
332. <https://dataventures.nz/#what>
333. <https://dataventures.nz/>
334. <https://www.mmventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf>
335. https://aiforum.org.nz/wp-content/uploads/2018/07/AI-Report-2018_web-version.pdf
336. <https://nztech.org.nz/wp-content/uploads/2014/10/Digital-Skills-for-a-digital-nation-online.pdf>
337. https://www.educationcounts.govt.nz/statistics/tertiary_education/retention_and_achievement
338. <https://www.mmventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf>
339. <https://www.mmventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf>
340. <https://www.nature.com/news/ai-talent-grab-sparks-excitement-and-concern-1.19821>
341. <https://www.nature.com/articles/d41586-018-07730-1>
342. <https://jfgagne.ai/talent-2019/>
343. <https://minterellison.co.nz/our-view/joint-venture-to-drive-artificial-intelligence-in-legal-services>
344. <https://www.mmventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf>
345. <https://www.telegraph.co.uk/technology/2019/07/17/uk-ai-gender-diversity-crisis-number-female-scientists-drops/amp/>
346. <https://jfgagne.ai/talent-2019/>
347. <http://reports.weforum.org/global-gender-gap-report-2018/assessing-gender-gaps-in-artificial-intelligence/>
348. <https://www.wired.com/story/artificial-intelligence-researchers-gender-imbalance/>
349. <https://www.telegraph.co.uk/technology/2019/07/17/uk-ai-gender-diversity-crisis-number-female-scientists-drops/amp/>
350. <https://aischool.microsoft.com/en-us/home>
351. https://www2.deloitte.com/content/dam/insights/us/articles/Tech-Trends-2019/DI_TechTrends2019.pdf

352. <https://landing.ai/ai-transformation-playbook/>
353. <https://landing.ai/ai-transformation-playbook/>
354. <https://www.mmcventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf>
355. https://www.pwc.co.nz/publications/ceo_survey/2019/ceo-survey-2019.pdf
356. <https://www.tensorflow.org/>
357. <https://lobe.ai/>
358. <https://www.theverge.com/2018/6/12/17452742/deep-learning-ai-learn-lobe-made-easy-coding>
359. <https://research.hackerrank.com/developer-skills/2019>
360. <https://www.newzealandnow.govt.nz/move-to-nz/new-zealand-visa/work-visa/skilled-migrant-visas>
361. <https://www.newzealandnow.govt.nz/move-to-nz/new-zealand-visa/work-visa/temporary-work-visa>
362. <https://www.forbes.com/sites/samshead/2019/02/20/uk-government-to-fund-ai-university-courses-with-115m/#54c354f430dc>
- 362a. <https://medium.com/syncedreview/south-korea-aims-high-on-ai-pumps-2-billion-into-r-d-de8e5c0c8ac5>
363. <https://www.whitehouse.gov/ai/ai-american-worker/>
364. <https://www.politico.eu/article/finland-one-percent-ai-artificial-intelligence-courses-learning-training/>
365. <https://fcai.fi/fcainews/2019/5/29/finland-offers-ai-training-to-inmates>
- 365a. Al Singapore website, <https://www.aisingapore.org/industryinnovation/ai4k/#>
366. <https://nztech.org.nz/wp-content/uploads/2014/10/Digital-Skills-for-a-digital-nation-online.pdf>
367. <http://elearning.tki.org.nz/Teaching/Curriculum-areas/Digital-Technologies-in-the-curriculum>
368. <http://elearning.tki.org.nz/Teaching/Curriculum-areas/Digital-Technologies-in-the-curriculum>
369. <https://www.immigration.govt.nz/new-zealand-visas/options/start-a-business-or-invest/i-want-to-invest-or-do-business-in-nz/the-global-impact-visa>
370. <http://nzcurriculum.tki.org.nz/The-New-Zealand-Curriculum/Technology>
371. <https://cs-field-guide-dev.appspot.com/en/chapters/>
372. <https://www.stuff.co.nz/national/education/94163443/shake-up-for-school-curriculum-revealed>
373. <https://www.weforum.org/agenda/2018/06/teaching-ai-in-schools-could-equip-students-for-the-future>
374. <https://www.education.govt.nz/our-work/consultations/recent-consultations/digital-technology-consultation/how-weve-strengthened-digital-technologies-and-hangarau-matihiko-content-in-the-curriculum/>
375. <https://www.stuff.co.nz/national/education/108203269/digital-curriculum-deadline-unrealistic-teachers-union-fears>
376. Artificial Intelligence and Machine Learning in the Education Sector - Global Trends and Application in New Zealand, X4 Consulting, February 2018, https://docs.wixstatic.com/ugd/e3f055_fdbba28bccda47adbc5adab7e964513a.pdf
377. <https://www.newerit.co.nz/is-your-school-2020-digital-curriculum-ready/>
378. <https://shadowtech.nz/highlights/>
379. <https://futurefive.co.nz/story/partnership-brings-ai-maths-tutor-to-nz-schools>
380. <https://www.kickstarter.com/projects/1740916372/the-complete-guide-to-artificial-intelligence-for>
381. <https://www.teachermagazine.com.au/articles/ai-classroom-activity-facial-recognition>
382. <https://www.teachermagazine.com.au/articles/ai-classroom-activity-machine-learning>
383. <https://thoughtexperiment.co.nz/2018/11/12/auckland-university-robotics-lab/>
384. <https://medium.com/syncedreview/chinese-publisher-introduces-ai-textbooks-for-preschoolers-b95e1a89cfa0>
385. <http://www.chinadaily.com.cn/a/201806/11/WS5b1de85fa31001b82571f4ca.html>
386. <http://www.cnbc.cmu.edu/2018/05/17/teaching-ai-for-k-12-students-national-guidelines/>
387. https://www.nzherald.co.nz/nz/news/article.cfm?c_id=1&objectid=11689381
388. <https://nztech.org.nz/2019/04/10/nztech-inform-achieving-diversity-and-inclusion-in-tech/>
389. <https://123tech.nz/>
390. <https://www.cio.co.nz/article/658398/do-your-homework-first-data-algorithms-human-checkpoints/>
391. <https://research.hackerrank.com/developer-skills/2019>
392. <https://www.forbes.com/sites/janakirammsv/#3fc2047b2872>
393. https://info.hackerrank.com/rs/487-WAY-049/images/HackerRank_2019-2018_Developer-Skills-Report.pdf
394. <https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf>
395. <https://www.forbes.com/sites/bernardmarr/2018/08/13/the-4th-industrial-revolution-is-here-are-you-ready/#dea6b27628b2>
396. <https://research.hackerrank.com/developer-skills/2019>
397. <http://www.anu.edu.au/research/innovation/all-stories/one-year-in-3a-institute-to-teach-its-new-applied-science>
398. <https://acola.org/hs4-artificial-intelligence-australia/>
399. <https://standards.ieee.org/industry-connections/ec/ead-v1.html>
400. <https://link.springer.com/article/10.1007/s11023-018-9482-5>
401. <https://www.geekwire.com/2019/microsoft-will-adding-ai-ethics-standard-checklist-product-release/>

402. <https://www.theguardian.com/technology/2018/mar/19/uber-self-driving-car-kills-woman-arizona-tempe>
403. <https://www.eturbonews.com/248194/boeing-737-maxs-disasters-could-impair-global-artificial-intelligence-development/>
404. <https://www.healthnewsreview.org/toolkit/tips-for-understanding-studies/medical-devices/>
405. <https://www.law.ox.ac.uk/business-law-blog/blog/2019/02/tackling-algorithmic-facilitated-tacit-collusion-proportionate-way>
406. http://www.europarl.europa.eu/doceo/document/A-8-2017-0005_EN.html?redirect
407. <https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe>
408. <https://transparencyreport.google.com/youtube-policy/featured-policies/violent-extremism?hl=en>
409. <https://www.theguardian.com/media/2019/apr/24/mps-criticise-tech-giants-for-failure-to-report-criminal-posts-twitter-facebook-google-youtube>
410. <https://blogs.microsoft.com/on-the-issues/2019/03/24/a-tragedy-that-calls-for-more-than-words-the-need-for-the-tech-sector-to-learn-and-act-after-events-in-new-zealand/>
411. https://www.nzherald.co.nz/nz/news/article.cfm?c_id=1&objectid=12231337
412. <https://knightcolumbia.org/content/tim-wu-first-amendment-obsolete>
- 412a. <https://knightcolumbia.org/content/tim-wu-first-amendment-obsolete>
413. <https://www.radionz.co.nz/news/world/389049/pm-arrives-at-christchurch-call-but-trump-and-zuckerberg-absent>
414. <https://www.itv.com/news/2019-05-15/facebook-live-rules-tightened-following-christchurch-attack/>
415. <https://www.youtube.com/watch?v=xP4E17cB4oY>
416. <https://aiforum.org.nz/2019/05/23/reaction-to-the-christchurch-call-from-the-ai-forums-ethics-law-and-society-working-group/>
417. <https://privacyinternational.org/blog/2537/profiling-and-automated-decision-making-artificial-intelligence-violating-your-right>
418. <https://privacyinternational.org/topics/artificial-intelligence>
419. <https://www.nytimes.com/2019/04/14/technology/china-surveillance-artificial-intelligence-racial-profiling.html>
420. <https://www.bbc.com/news/technology-48555149>
421. <https://privacyinternational.org/examples-abuse/860/databases-facial-images-proliferate>
422. <https://techcrunch.com/2019/02/10/privacy-is-a-commons/>
423. <https://blog.google/technology/safety-security/privacy-everyone-io/> and <https://www.tensorflow.org/federated>
424. The Productivity Commission. (2019). Technological Change and the Future of Work issues paper. Retrieved from: https://www.productivity.govt.nz/sites/default/files/Technological%20change%20and%20the%20future%20of%20work__Issues%20paper%20FINAL.pdf
425. <https://www.stuff.co.nz/national/politics/112836958/prebudget-announcement-77m-for-supergold-card-digital-literacy-training>
426. StatsNZ reference: http://archive.stats.govt.nz/Census/2013-census/profile-and-summary-reports/ethnic-profiles.aspx?request_value=24705&tabname=Motorvehicles, andphones, fax, andInternet
427. <https://www.digital.govt.nz/dmsdocument/113-digital-inclusion-blueprint-te-mahere-mo-te-whakaurunga-matihiko> see pg11
428. https://www.productivity.govt.nz/sites/default/files/Technological%20change%20and%20the%20future%20of%20work__Issues%20paper%20FINAL.pdf see pg38
429. <http://www.oecd.org/skills/piaac/Skills-Matter-New-Zealand.pdf>
430. Digital Inclusion Research Group. (2017). Digital New Zealanders: the Pulse of Our Nation. Retrieved from: <https://www.wairoadc.govt.nz/assets/Document-Library/Long-Term-Plan-2018-28/Supporting-Documents-2018-28/Library-digital-new-zealanders-the-pulse-of-our-nation-may-2017.pdf>
- 430a. <https://www.digital.govt.nz/digital-government/digital-transformation/digital-inclusion/digital-inclusion-blueprint/>
431. <http://www.oecd.org/social/inequality.htm> - Although some recent research asserts that global interpersonal income inequality has actually declined, see <https://www.sciencedirect.com/science/article/abs/pii/S0305750X19300968>
432. <https://wir2018.wid.world/part-2.html>
433. <https://www.brookings.edu/research/credit-denial-in-the-age-of-ai/>
434. <https://www.wsj.com/articles/self-driving-technology-threatens-nearly-300-000-trucking-jobs-report-says-1536053401>
435. For instance see <https://www.thebritishacademy.ac.uk/sites/default/files/AI-and-work-evidence-synthesis.pdf>, <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>, and http://www3.weforum.org/docs/WEF_Future_of_Jobs_2018.pdf
436. <https://www2.deloitte.com/insights/us/en/focus/tech-trends/2019/driving-ai-potential-organizations.html>
437. <https://www.oxfordinsights.com/government-ai-readiness-index/>
438. <https://www.oxfordinsights.com/ai-readiness2019>
439. <https://www.oxfordinsights.com/government-ai-readiness-index/>
440. <https://www.oxfordinsights.com/ai-readiness2019>
441. <https://www.newsroom.co.nz/2018/03/26/100058/nz-firms-way-behind-adopting-digital-technologies>

442. <https://www.callaghaninnovation.govt.nz/sites/all/files/ai-whitepaper.pdf>
443. APEJ Cognitive and AI Adoption Survey. IDC. 2018.
444. Asia/Pacific Future of Work Survey 2018. IDC. 2018.
445. <https://www.productivity.govt.nz/inquiry-content/3960?stage=2>
446. <https://www.incleanmag.co.nz/2019/03/artificial-intelligence-to-double-the-rate-of-innovation/>
447. <https://www.cio.co.nz/article/658779/firms-urged-address-looming-skills-challenge-an-ai-enabled-workforce/>
448. https://www.pwc.co.nz/publications/ceo_survey/2019/ceo-survey-2019.pdf
449. <https://www.mbie.govt.nz/science-and-technology/science-and-innovation/agencies-policies-and-budget-initiatives/vision-matauranga-policy/>
450. <https://press-files.anu.edu.au/downloads/press/n2140/pdf/book.pdf?referer=2140>
452. <https://koreromaori.io/>
453. Then you need a human translator to work on these structures to bring the text to life – to give it sense, tone, more structure and style, to make the new translation appear natural and fluent in the new language.
454. <https://www.newsroom.co.nz/@comment/2018/06/25/128744/the-rise-of-mori-tech-business>
455. <https://www.gartner.com/smarterwithgartner/the-cios-guide-to-artificial-intelligence/>
456. <https://www.gartner.com/smarterwithgartner/the-cios-guide-to-artificial-intelligence/>
457. <https://www.gartner.com/smarterwithgartner/the-cios-guide-to-artificial-intelligence/>
458. <https://docs.microsoft.com/en-us/learn/modules/build-an-ai-ready-culture/5-apply-ai-maturity-model-assessment> and <https://discover.microsoft.com/ai-readiness-and-your-organization/>
- 458a. <https://www.prophet.com/2018/12/the-ai-maturity-playbook/>
459. <https://www.theta.co.nz/news-blogs/theta-news/the-real-world-of-artificial-intelligence-in-new-zealand/>
460. <https://www.callaghaninnovation.govt.nz/sites/all/files/ai-whitepaper.pdf>
461. <https://www.data.govt.nz/assets/Uploads/data-strategy-and-roadmap-dec-18.pdf>
462. <https://landing.ai/content/uploads/2018/12/AI-Transformation-Playbook-v8.pdf>
463. <https://www.awaregroup.com/proof-of-concepts/>
464. <https://www.accenture.com/nz-en/topic-accenture-government-artificial-intelligence>
465. <https://www.theta.co.nz/solutions/innovation-lab/artificial-intelligence/>
469. Introduction to Machine Learning. Algorithmia, 19 March 2018. Retrieved 28 April 2019 from: <https://blog.algorithmia.com/introduction-to-machine-learning/>
470. What are the types of Machine Learning? Towards Data Science, 5 December 2018. Retrieved on 28 April 2019 from: <https://towardsdatascience.com/what-are-the-types-of-machine-learning-e2b9e5d1756f>
471. WIPO Technology Trends 2019 Artificial Intelligence. World Intellectual Property Organisation, 2019. Page 147. Retrieved from: https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf
472. The Hundred-Page Machine Learning Book. Andriy Burkov, 2019. Page 91.
473. WIPO Technology Trends 2019 Artificial Intelligence. World Intellectual Property Organisation, 2019. Page 147. Retrieved from: https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf
474. A Tour of the Most Popular Machine Learning Algorithms. Machine Learning Mastery, 25 November 2013. <https://machinelearningmastery.com/a-tour-of-machine-learning-algorithms/>
475. Supervised vs. Unsupervised Learning. Towards Data Science, 23 March 2018. <https://towardsdatascience.com/supervised-vs-unsupervised-learning-14f68e32ea8d>
476. WIPO Technology Trends 2019 Artificial Intelligence. World Intellectual Property Organisation, 2019. Page 147. Retrieved from: https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf
477. Artificial Intelligence: What Is Reinforcement Learning - A Simple Explanation & Practical Examples. Forbes, 28 September 2018. Retrieved on 28 April 2019 from: <https://www.forbes.com/sites/bernardmarr/2018/09/28/artificial-intelligence-what-is-reinforcement-learning-a-simple-explanation-practical-examples/#218a0bd139ce>
478. The Hundred-Page Machine Learning Book. Andriy Burkov, 2019. Page 21.
479. Oxford English Dictionary online (Non-technical definition)
480. The Business of Artificial Intelligence. Erik Bynnjolfsson and Andrew McAfee, Harvard Business Review. July 2017. <https://hbr.org/cover-story/2017/07/the-business-of-artificial-intelligence>



AI for Wellbeing
AI for Sustainability
AI for the Economy

