February 2022 | New Zealand Thematic Research

Artificial Intelligence

A guide to AI and implications for New Zealand

FORSYTH BARR

Neil Jacobstein
Chair, AI & Robotics , Singularity Univ.
NZ: +64 (0) 21 0230 6475
US: +1 415 650-7603

Aaron Ibbotson
aaron.ibbotson@forsythbarr.co.nz
+64 9 368 0024

Andy Bowley
andy.bowley@forsythbarr.co.nz
+64 4 495 8246

Artificial intelligence (AI) represents one of the largest opportunities for industry since the advent of electricity.

AI enables us to create, improve, and externalise amplifiers for our minds and intellect, but it can also be dangerous and carries real risk. This has already had a profound effect on virtually every field of human endeavour — from medicine, to finance, to transportation. Going forward, AI will disrupt and augment current industrial processes and fundamentally transform how businesses operate. This report sheds light on what AI is and what it isn't, what competitive advantages it could deliver, its practical and ethical risks, and how to harness the power of AI responsibly. We take a closer look at the main listed sectors and how AI could impact these businesses over the short and medium term. The report also includes the inaugural Forsyth Barr Corporate AI survey. The survey provides on the ground insight about the gaps between the potential for applications of AI and the realities of current implementations in the New Zealand companies surveyed.

Table of contents

A guide to AI and implications for New Zealand

Artificial intelligence (AI) represents one of the largest opportunities for industry since the advent of electricity. Electricity allowed humans to harness electromechanical power to transform almost every area of our social, business, and personal lives. AI enables us to create, improve, and externalise amplifiers for our minds and intellect, however it can also be dangerous and carries real risk. This has already had a profound effect on virtually every field of human endeavour — from medicine, to finance, to transportation. AI aided the unprecedented economic growth of Google, Tesla, Facebook, Amazon, and Netflix. Going forward, AI will disrupt and augment current industrial processes, and fundamentally transform the possibilities for business and humanity itself. Wise management will require both boldness and informed caution.

In recent years AI has become a hot topic in the media — from applications breakthroughs to concerns about ethics and privacy. There is a tendency to conflate exposure to AI in the news with knowledge about how AI works and where it can be applied successfully in business and government. However, the more you know about the underlying technology and its applications, the more you realise that we are at the beginning, not the end of the opportunities and risks ahead.

This report seeks to clarify:

- What AI is and what types of AI have been developed
- What competitive advantage AI can confer and what industry applications already exist
- What classes of vulnerabilities surround the application of AI
- What are the breakthrough AI opportunities and risks for New Zealand companies
- What are the key ethical concerns with AI, how can AI projects best be managed
- How mature and competitive is New Zealand's use of AI
- What global competition exists in AI today
- How can organisations do due diligence around AI
- What are some relevant references for exciting new research and development directions in the field.

AI in New Zealand

Our proprietary survey of New Zealand listed corporates' use of AI provides on the ground insight about the gaps between the potential for applications of AI and the reality of current implementations. There is much room for increased adoption of the technology and the resulting opportunities for growth.

The challenge for New Zealand is fourfold: (1) learn the essentials about AI and machine learning, (2) identify the most promising NZ companies applying AI effectively, (3) identify some of the key international AI R&D results and best practices for managing AI responsibly, and (4) create new AI business opportunities and partnerships.

New Zealand's relatively small economy is both a weakness and a strength. The key weakness is the relatively small AI research and development community. However, that also means that nimble NZ companies can harvest the world's AI R&D — which is typically published immediately in the open literature. This is an enormous opportunity to get the return on investment associated with commercialising the last 5% of powerful AI innovations. Our purpose in writing and delivering this report is partly to empower New Zealand companies to make the right AI investment and regulation decisions, or at least figure out who they can trust to do it.

Artificial intelligence: The branch of computer science that deals with giving computers the ability to perform aspects of intelligence, including problem solving, visual and auditory pattern recognition, and sensory motor control.

Machine learning: The branch of artificial intelligence that deals with giving computers the ability to learn aspects of intelligent behaviour, typically by exposing algorithms to data.

Deep learning: Is a class of powerful machine learning algorithms that utilise multiple layers of a hierarchy of pattern recognition elements or nodes to extract higher level features from initial data input, and passes the partial results up the layers of the hierarchy until the target or problem is recognised.

Algorithm: A set of steps or processes that are performed by a computer to solve a problem or complete a task.

Neural networks: Machine learning algorithms that use networks of artificial neurons or nodes modelled loosely on the biological neural networks of animal brains.

Our report in pictures

Figure 1. Deep learning is a type of machine learning which is a type of Artificial Intelligence

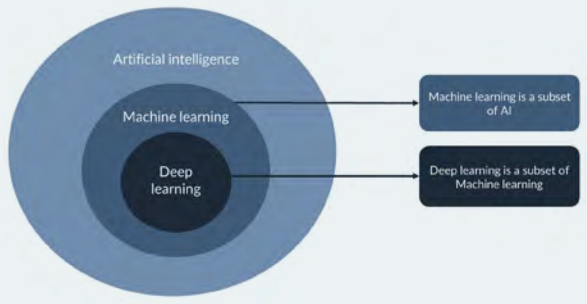

Source: Forsyth Barr analysis

Figure 2. AI can provide organisations with an enhanced OODA loop, providing faster, more effective decisions

Source: Forsyth Barr analysis

Figure 3. More global AI start-up funding in 1H21 than in any previous full year. This growth in AI investment is likely to continue for the foreseeable future

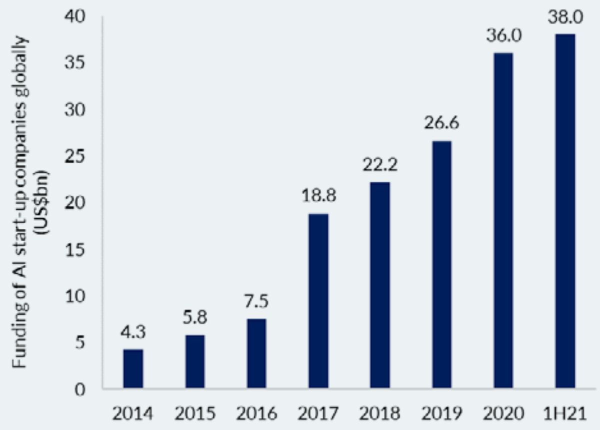

Source: Forsyth Barr analysis, Statista

Figure 4. AI is solving ever more problems — DeepMind's AI can now look at your retina and predict smoking, age, and diabetes

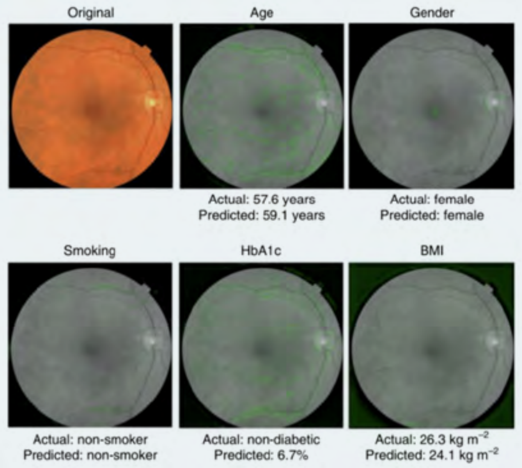

Source: UK Biobank/Google, Forsyth Barr analysis

Figure 5. Forsyth Barr NZ Corporate AI survey: Has your organisation directly invested in any AI applications? The majority have but only to a limited extent to date

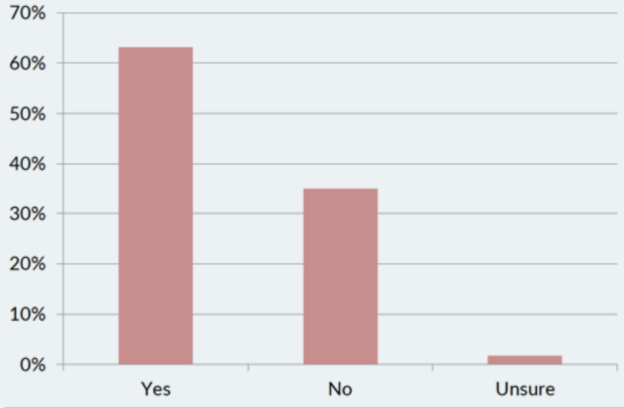

Source: Forsyth Barr analysis

Figure 6. Forsyth Barr NZ Corporate AI survey: What application categories is your organisation likely to apply AI to over the next two years?

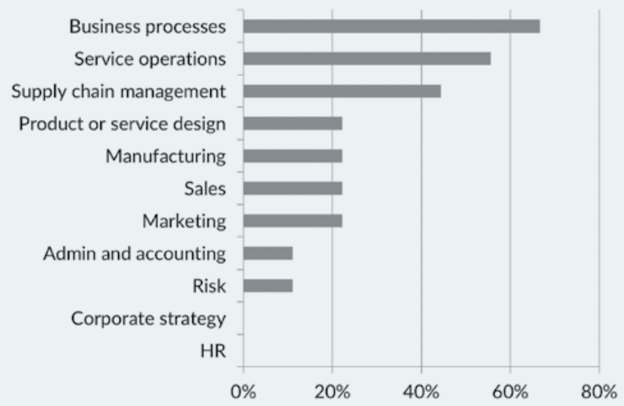

Source: Forsyth Barr analysis Note: more than one answer may apply to each respondent

What is Artificial Intelligence?

AI can be framed in different ways: (1) pattern recognition and problem solving at expert or above expert skill levels, and (2) rational or intelligent agents that do a situation assessment, consider what actions are available, and select the action that will best improve the probability of achieving its objective. These are some of the top new companies in different areas of AI and Machine Learning.

Pattern recognition and problem solving

The AI field is now over 60 years old. It has been the subject of many hype cycles, investment booms and busts, science fiction stories, and speculation about the future of super-intelligent machines. In spite of all this volatility and drama, the central business objective of artificial intelligence is to build computers that do pattern recognition and problem solving at expert or above expert skill levels.

This has been an extremely fruitful endeavour, beginning with systems in the 1960s that could recognise patterns and beat human checkers/draughts players. The late 1970s ushered in a wave of rule-based expert systems that emulated the problem solving behaviour of experts in narrow fields like oil well diagnostics, specialised disease diagnostics such as bacteremia or blood inflections, logistics, planning, case law, tax accounting, and even credible watercolor art.

Today, machine learning systems do pattern recognition and problem solving in virtually every narrow field of human endeavour.

The 1980s and 1990s produced hundreds of working rule-based applications, but they were expensive to build, difficult to maintain as the knowledge of the field changed, and they didn't learn. They had to be hand coded rule by rule, based on interviewing human experts and replicating their domain knowledge and problem solving procedures. In addition, the computers at the time were slow and underpowered for the class of problems addressed, and the data that could have informed a real time reevaluation of the problem space were mostly not yet available in convenient machine readable forms. In 1997, IBM's DeepBlue beat the then world champion, Garry Kasparov, at chess, but it was generally acknowledged to be a brute force approach to chess problem solving.

The result of these problems meant that the systems were expensive, slow to build, hard to maintain, and they did not continuously improve in pattern recognition or problem solving over time. Organisations became increasingly disenchanted with the performance of these systems, leading to a bust cycle commonly referred to as the AI Winter.

By 2009, computers were orders of magnitude faster, machine readable data were readily available on the Internet, and neural network based machine learning algorithms that had been incubating for decades began making major strides in pattern recognition. Exponentially faster computing hardware has been pivotal to progress in the field, allowing practitioners to build systems quickly with well established algorithms that could learn directly from readily available data. Suddenly, the AI Winter was over, and machine learning systems began cracking very hard problems. Examples include: detecting diabetic retinopathy, recognising virtually all the images in a large database of images (ImageNet), and beating world champions at Go.

Today, machine learning systems do pattern recognition and problem solving in virtually every narrow field of human endeavour, as Figure 7 on startups demonstrates. These systems are still addressing narrow problems. They do not exhibit the broad, deep, and subtle features we associate with human intelligence. However, they are cost effective to build, they learn from data, can continuously improve over time, and they sometimes outperform humans doing similar narrow tasks.

Intelligent agents achieve their objectives

A different way of framing AI is by thinking of these systems as rational or intelligent agents that do situation assessment, consider what actions are available, and select the action that will best improve the probability of achieving its objective. There are many agents that have been built for industry, games, or the military that have this framework. These systems often use a combination of techniques including pattern recognition, planning, optimisation, game theory, and other algorithms. What they have in common is the production of an agent that utilises each algorithm or strategy to achieve its subgoals and maximise a reward signal or outcome. Some of these systems form the basis of characters in multi-agent games, or an adversary in military competitions.

Figure 7. AI 100 class of 2021

Source: CBInsights, Forsyth Barr analysis

Figure 8. AI revenues are increasing

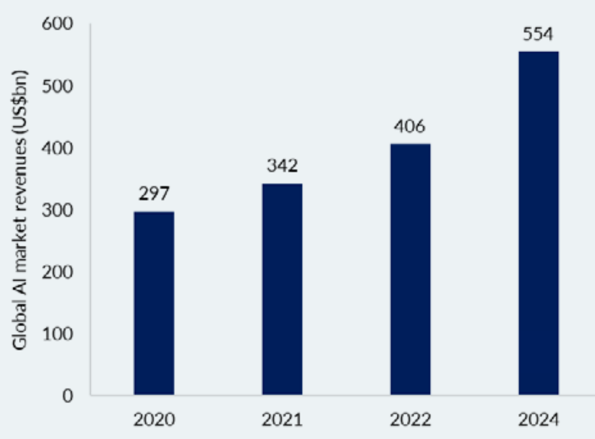

Source: Forsyth Barr analysis, Statista

Figure 9. More AI startup funding in 1H21 than in 2014–2017

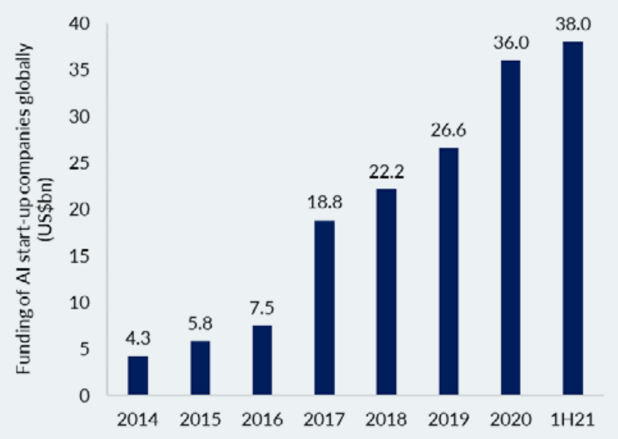

Source: Forsyth Barr analysis, Statista

What are the key types of AI?

There are many types of AI, encompassing everything from simple algorithms based on a few specific rules to large language modes that can write news articles often barely distinguishable from articles written by humans. We outline and discuss nine key types of AI below on page 15 (Figure 10).

Rules that explain but don't learn

Backchaining rules operate by working backwards from an action goal — what must be true to conclude a variety of action alternatives.

An example of traffic rules might be:

- TrafficLight is "green"
- Crosswalk is "clear"

IF TrafficLight is "red" **THEN** action is "stop"
IF TrafficLight is "green" **AND** crosswalk is "clear" **THEN** action is "proceed"

If queried for an explanation, this system could provide an obvious rationale for its "action" recommendation. Rules can also work forward (forward chaining) from a set of facts to what actions can be taken or ruled out.

Algorithms that learn but don't explain

Deep learning is the neural network machine learning algorithm that has received the most industry attention over the past decade.
Deep learning is an example of a hierarchical pattern recognition network. Each node or "neuron" in the first layer of the network is assigned an input value based on a data item, in this case, a pixel or picture dot in a photograph. Information is passed from each layer in the network to the layer above it using a channel.

Channels have a "weight" or coefficient based on training data. Each neuron has a value associated with it called its bias. The bias of a neuron is added to the weighted sum of the inputs reaching the neuron. The result is applied to an activation

function which determines if the neuron is activated or not. Every activated neuron passes information to the layer above it in the hierarchy.

Recognition occurs by passing partially recognised features up the layers of the pattern recognition hierarchy for refinement and combination with other parts of features. The hidden layers in the networks are all the layers between the input layer and the output layer. When the neural network reaches its output layer, the value of the activated node is the value of the recognition task.

These neural networks can achieve up to 98% plus accuracy recognising human faces. Practical applications may have higher error rates as part of a Know Your Customer program in a bank. Consider a bank that has 50,000 customers, and they provide customers with an incentive to allow their photo to be taken, preferably at least three times from several perspectives.

When exposed to a photo, the first layer of a multi-layered pattern recognition system may only identify vertical and horizontal lines. Layer two may recognise parts of faces, such as eyes, ears, nose, mouth, forehead, and chin. Layer three may recognise faces but not differentiate different people. Layer four may recognise feature sets that are associated with particular classes of people. Layer five may recognise specific people. Of course, it may take many more layers if the people belong to multiple population groups, or there are several members of a large family with similar age and looks.

Say the bank develops an image database with labelled images of people with two or three photos taken of them. The labelled photos are considered the "training set", the labels are the "answers", and the system proceeds to learn to differentiate and identify the people in the training set.

Then the bank tests the system's error rate by exposing the system to a batch of 50,000 unlabelled photos. Depending on the outcome, the system may need a larger and more diverse training set. The bank keeps working with the system and the training sets until the system achieves the target accuracy in identifying customers. Accuracy may well be affected by the race, sex, lighting or face orientation. This is a potential important source of bias in the recognition system that would need to be addressed.

It may take a lot of work to recognise 50,000+ heterogeneous people accurately under different poses and lighting conditions, with different glasses, haircuts, clothing, skin colour, facial features, and hats. However long the project, the deep learning algorithm that does the pattern recognition will likely be unable to explain to bank personnel how it did its task. The reason is that its methods are often in high dimensional mathematical spaces that do not lend themselves to explanation to humans. There are many research teams working on this problem but none have cracked it yet.

Supervised vs. unsupervised learning

The key distinction between these two approaches is that supervised learning makes use of labelled data and unsupervised learning does not. Supervised learning is conducted with "ground truth" of what the output values should be for the sampled data inputs in the training set.

Supervised learning is often done in the context of classification mapping of input to output labels. The objective of these systems is to learn a function that approximates the relationship between inputs and outputs observable in the data. These systems may predict 8 output data given the structure of the input data. However, if the "ground truth" associated with the input labelled data set is wrong or noisy, your mileage may vary, and the system's effectiveness will be reduced accordingly. Hence, the importance given to marshalling and cleaning input data.

Unsupervised learning systems may infer the relationships inherent in the input and output structure of the sample data. The features may be unknown or unanticipated by humans. This is one of the reasons why explanation is such a challenge in machine learning systems. In some cases, pattern recognition systems for interpreting x-ray data have been "reduced" to marking or pointing to significant discovered features without providing human interpretable text.

Common use cases for unsupervised learning are clustering, data exploration, and dimensionality

Figure 10. Key types of AI

Types of AI	Comments
Rules	Encode knowledge as facts, if-then conditions, and logical operations that can emulate the problem solving behaviour of human experts.
Supervised learning	Makes use of labelled data as examples of "ground truth" for recognising typical outputs.
Unsupervised learning	Learns patterns from unlabelled data, captured as probability densities or features useful for dimensionality reduction, clustering, or data exploration.
Reinforcement learning	Algorithms identify patterns in unlabelled data via a reward signal for prediction or task performance.
Constraint satisfaction	The process of finding a solution to a set of constraints that impose conditions on the variables that a viable solution must satisfy.
Recommendation engines	Algorithms that predict the match between customer buying history and behaviour and the most probable product candidates.
Semantic web	A vision of the current World Wide Web with machine-understandable information, together with intelligent agents that utilise this information to provide useful services.
Large language models	These systems find patterns and make inferences about the huge corpus of text parameters they contain based on text objects often from public sources.
Differentiable programming	Collection of machine learning techniques that rewrite or control at least one aspect of itself by using a neural network to optimise performance automatically against a set of objectives.

Source: Forsyth Barr analysis

Figure 11. Deep learning use

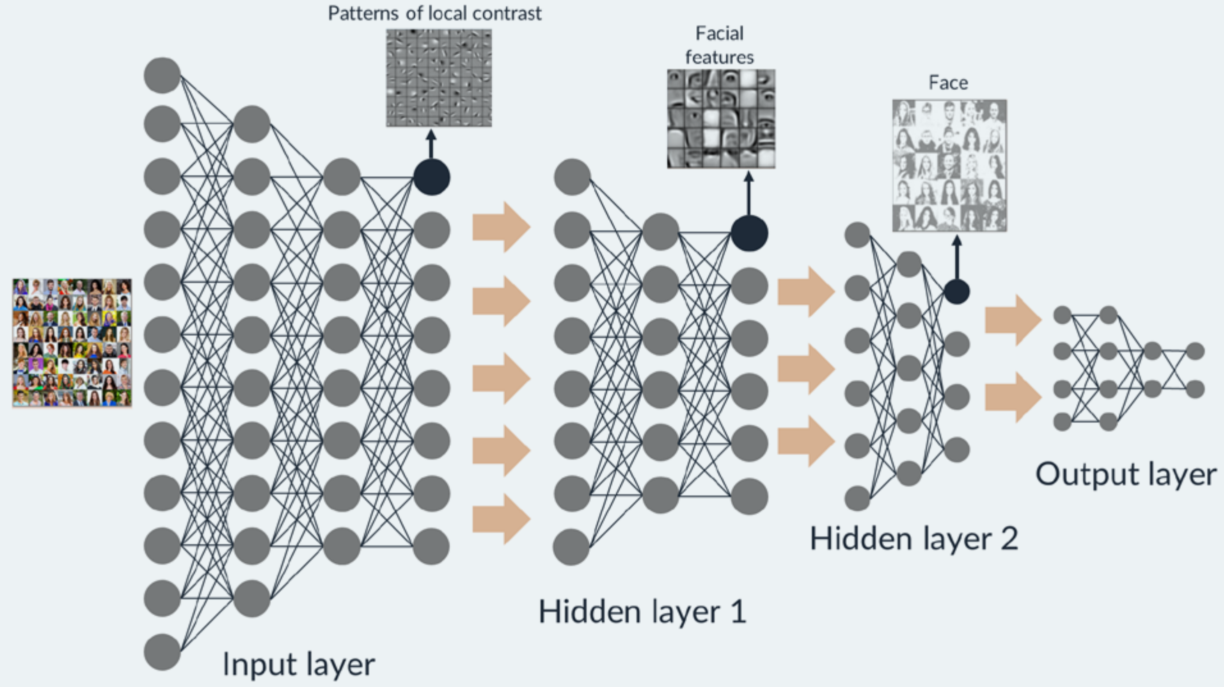

Source: Edureka, Forsyth Barr analysis

reduction. Unsupervised learning discovers relationships inherent in the data set. For example, discovering what non obvious features are common to people who may want to buy a consumer product, attend a seminar, or join a club.

Dimensionality reduction determines which features or dimensions of the data are less important for making specific predictions.

Those features may be dropped in further data analysis, which given massive data sets, could be significant for response time and cost control of the analysis.

Reinforcement learning

Reinforcement learning uses algorithms that typically do not rely only on historical data sets to learn to make a prediction or perform a task. They learn as humans sometimes do, through trial and error. The technology is scalable and often able to optimise complex decisions. It can also be wrong if the problem is not specified correctly.

The agent takes actions inside a design problem or environment and it receives "rewards" when it takes productive actions against the human specified performance objective. It continues to explore the problem space until it has maximised cumulative reward. It can learn from examples and test new and adaptive solutions much faster than humans do.

For example, Emirates Team New Zealand used reinforcement learning to simulate optimal designs of the hydrofoils critical in new sailboat performance. Using reinforcement learning, they were able to evaluate thousands of hydrofoil design concepts before engaging in costly and limited builds. In addition, the rapid simulations of new hydrofoil designs enabled the human team to develop and test new competitive sailing strategies.

Constraint satisfaction

Constraint satisfaction is the process of finding a solution to a set of constraints that impose conditions on the variables that a viable solution must satisfy. Constraint satisfaction problems may be arbitrarily complex, requiring simplifying assumptions and methods.

Examples of problems that fit into this framework are product configuration, complex auctions, and resource allocation. Constraint satisfaction may combine symbolic or logical constraints, and augment traditional numeric constraints captured in linear programming or other mathematical optimisation systems.

These problems may be solved by formulating an overall objective that must be satisfied, subject to a set of specific constraints. For example, a winning paper in the Innovative Applications of AI Conference utilised AI based constraint satisfaction techniques to enable a much more granular auction than the traditional auction where participants bid only on the price of a commodity or service contract.

The insight that Tuomas Sandholm and his graduate students at Carnegie Mellon University developed was that most auctions underspecify the actual specific interests of the respective parties. In addition to price, buyers and sellers may be interested in timing, loan guarantees, defect rates, customer service contracts, new properties of materials such as ceramics or steel, and alternative logistics arrangements such as shipping routes (Tuomas Sandholm, et. al, 2006, 2015).

Traditional auction software doesn't begin to capture this complexity, so billions of dollars of value are left on the table. This work led to Sandholm founding an auction platform that utilised granular constraint satisfaction and saved both buyers and sellers billions of dollars.

Figure 12. Supervised and unsupervised learning

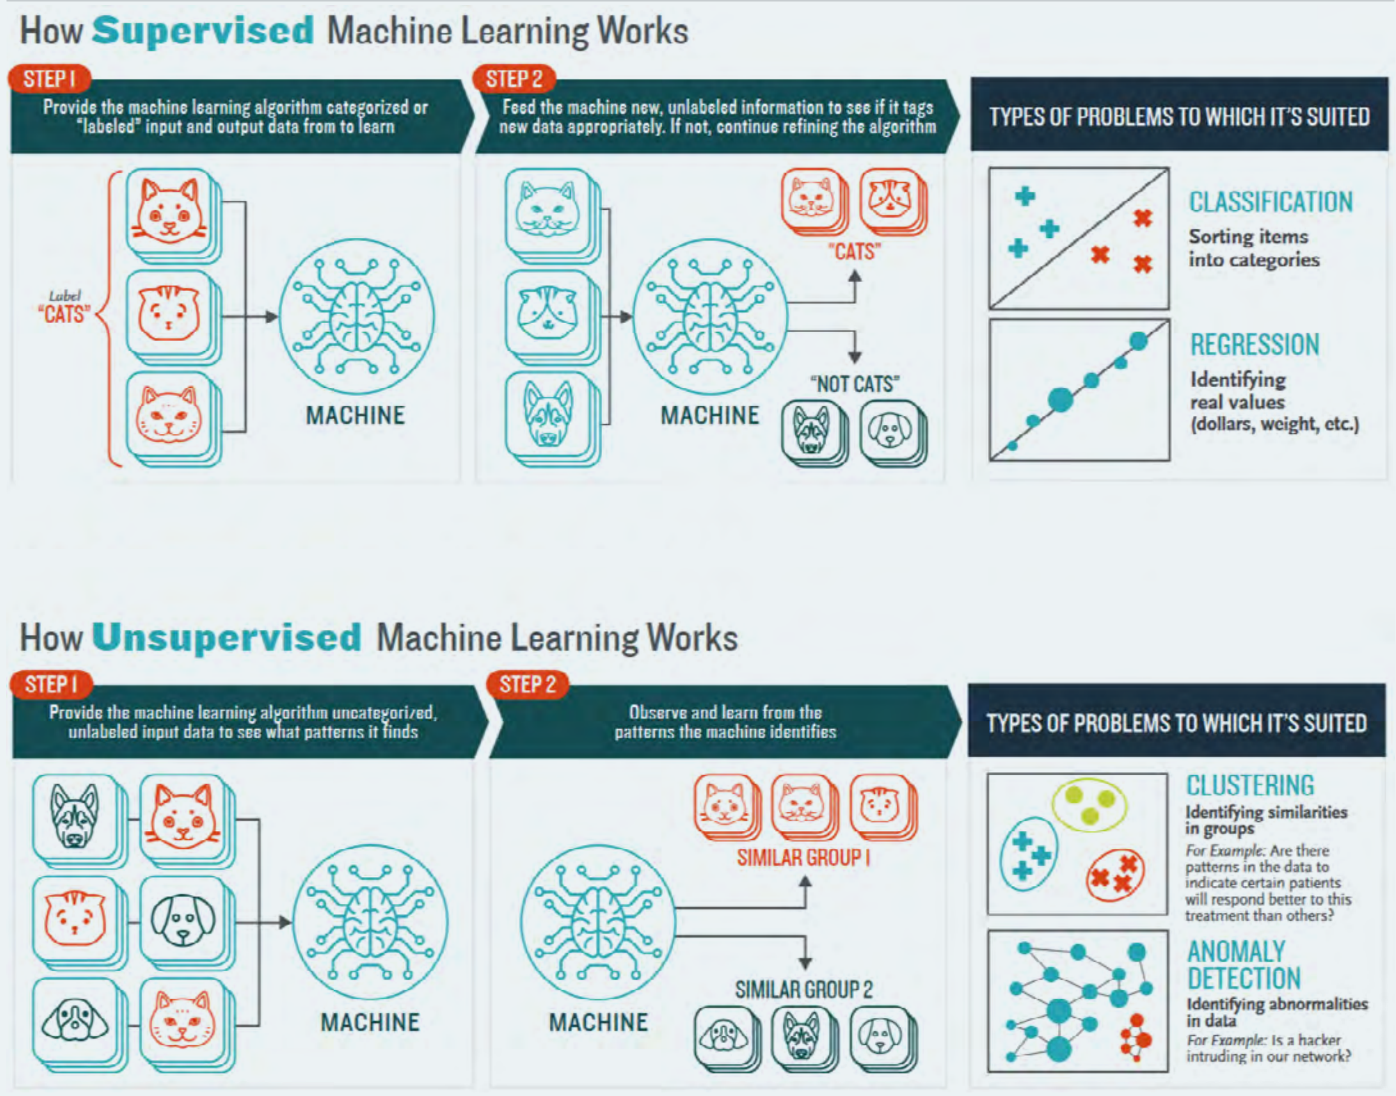

Source: Andrej Baranovskij ["Unsupervised Machine Learning Example in Keras", Towards Data Science, May 28, 2020], Forsyth Barr analysis

Recommendation engines

AI algorithms have been used in recommendation engines since the 1980s. The early systems used rule-based technology. However, these systems were modified over time to use powerful statistical techniques, as captured in the 2002 book *"Word of Mouse: The Marketing Power of Collaborative Filtering"*, by John Riedl.

Collaborative filtering algorithms perform "birds of a feather" clustering methods that match similar customers that bought similar products and predict recommendations accordingly.

Since then more advanced machine learning recommendation engines based on deep learning and other algorithms were developed by companies such as Amazon to make book and product recommendations. These systems generate significant revenue opportunities given that the recommendations are customer specific predictions.

The algorithms may know aspects of the customer desires better than the customers themselves. Tweaking recommendation algorithms is a matter of improving the predictive match between customer buying history and behaviour and the most probable product candidates.

Netflix was so keen on improving its own recommendation engine that, in 2006, it sponsored the Netflix Prize — an open data competition with a US$1m prize to improve its recommendation algorithm. The minimum objective was specifically to improve its prediction of user movie ratings by 10%. The data set at the time had over 100 million ratings of 17,770 movies from 480,189 customers. That is a small fraction of what today's numbers would be.

The contestants cooperated and competed with each other. The contest drew hobbyists, academics, and professionals, both for the money and the unprecedented data set. The rules specified that the winners could retain ownership of the algorithm and license it nonexclusively to Netflix. The improvements captured by the process of developing a winning algorithm produced changes in the Netflix predictive ratings engine that paid for the contest in spades.

In addition, one of the Netflix Prize outcomes was the discovery that blending the results you got using different methods produced a significant improvement in the predictive value of the final results. Many public improvements in prediction and recommendation algorithms came about as a result of the contest. This success likely inspired the machine learning contest platform Kaggle, which was later sold to Google.

When it was discovered that the contest user data could be de-anonymised, using background knowledge from the Internet Movie Database, a lawsuit (subsequently settled) convinced Netflix to end its algorithmic prize competitions based on user data.

Netflix eventually migrated its recommendation engine from predictions of user ratings to predictions about what movies users were likely to consume based on their actual streaming behaviour. The lesson here is that not only does the AI technology change rapidly, but the underlying problem formulation changes based on rapidly evolving business model competition.

Semantic web

The Semantic Web is a vision of the World Wide Web with machine-understandable information (as opposed to most of the current Web, being targeted at human consumption), together with services from intelligent agents that utilise this information. The original vision of the Semantic Web was described in a Scientific American journal article (Tim Berners-Lee et.al, 2001). The concept has been partially implemented over the past 20 years and it has also undergone modification — to be more inclusive of other data sources in heterogeneous data bases. There is a series of research projects pursuing this vision with a variety of metadata formats, ontologies, knowledge graphs, and related tools.

What this means for corporate and government entities interested in staying on top of the accelerating knowledge in the world is that the tools for machine interpretable real time knowledge are present in early and embryonic form today, and they are evolving rapidly.

Figure 13. Reinforcement learning via reward signal

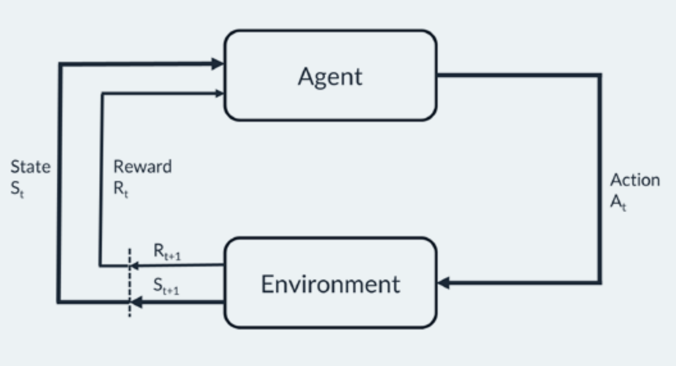

Source: Kdnuggets.com, Forsyth Barr analysis

Figure 14. Constraint satisfaction in auctions

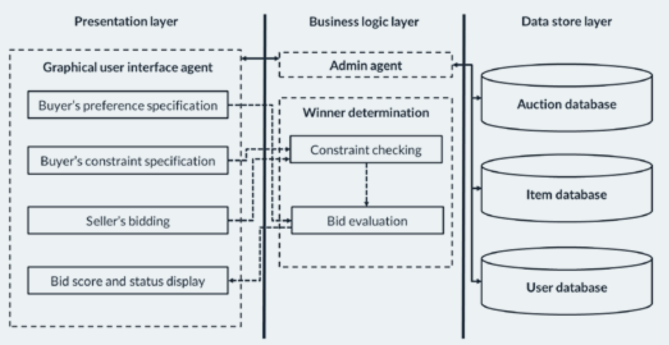

Source: scieno.cl, Forsyth Barr analysis

Figure 15. Yelp data available to machine learning

Source: Forsyth Barr analysis, Statista

Figure 16. Christmas shopping recommendation sources

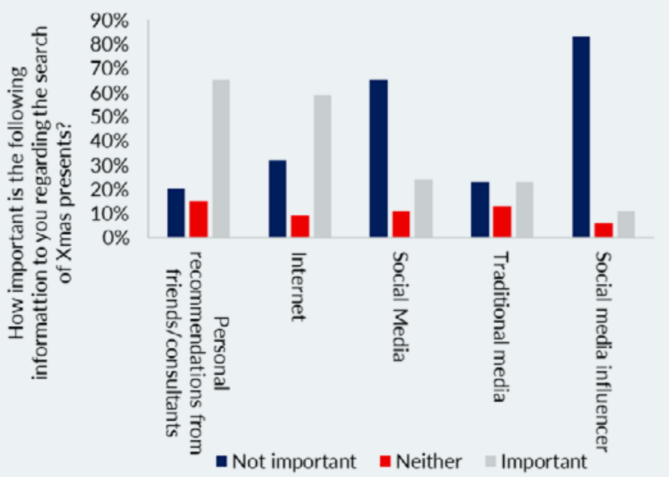

Source: Forsyth Barr analysis, Statista

There are mature commercial companies such as Semantic AI and Palantir that attempt to make this technology readily available for entities interested in "connecting the dots" of heterogeneous knowledge. The early adopters and heavy utilisers of this field of research and application are in intelligence and military services, hedge funds, future forecasting groups, and leading edge knowledge utilisation groups.

Large language models

Large language models incorporate billions of text objects from public sources like Wikipedia, ebooks, Reddit, and social media text. These large language models have grown radically in size and scope.

For example, OpenAI's Generative Pretrained Transformer GPT1 was trained with 117 million text parameters, or relationship weights between text objects. GPT2 contained 1.5 billion parameters. GPT3 was trained with orders of magnitude more data — 175 billion parameters. The GPT systems are still embryonic but they demonstrate that parameter size matters. GPT3 can do very credible sentence completion, text generation, and even image completion if it is trained on image objects vs. text objects.

These systems find patterns and make inferences about the huge corpus of text parameters they contain. The resulting system can generate and emulate human-like text in complete sentences and paragraphs. While these generated texts often appear to be the result of an intelligent agent, in reality, these systems have a very shallow understanding of the text. They have trouble with sentence comparisons and common sense.

For example, they may not be able to predict the consequences of constructing a commercial airplane with wings made of pasta. These systems are able to generate credible poems or narratives around a sporting event or other constrained situations but the results they generate are sometimes humorously off base.

For example, when asked what number comes before one million, GPT3 responded nine hundred thousand and ninety nine. The system also responded that Steve Job's location was in the Apple headquarters in Cupertino, California. That would be a plausible answer if Jobs was still alive. GPT2 generated text indicating that a human has two legs and two eyes, and a sunflower has one eye.

Even when GPT3 makes a mistake, it can often be tuned by providing the appropriate prompt — the questions used to solicit a response. GPT3 may indicate that "a pencil is heavier than a toaster", but if given the prompts first that "a kettle is heavier than a cat"and "the ocean is heavier than dust" the system may provide the correct response.

Nvidia and Microsoft recently announced their largest monolithic transformer language model to date, an AI model with 530 billion parameters they developed together, named the Megatron-Turing Natural Language Generation model (Alvi and Kharya, 2021). Dealing with such large volumes of text makes it difficult to clean the dataset of toxic language. This means that these systems can generate outputs that might be offensive (i.e. racist or sexist).

In January 2021, Google tested a 1.6 trillion parameter mixture of experts model (Fedus et. al. 2021). The rate of improvement in these large language models is blistering. However, as promising as these systems are, they are computationally intensive and still have many practical limitations.

Differentiable programming

Differentiable programming is a collection of machine learning techniques that rewrite or control at least one aspect of itself by using a neural network to optimise performance automatically against a set of performance objectives. The significance of this is that these programs incorporate the pattern recognition power of algorithms like deep learning, but they also have meta control of the problem solving or optimisation process, which make them more flexible problem solvers.

For example, the recent extraordinary performance of DeepMind's AlphaZero in rapidly learning to play world class chess, Japanese Shogi,

Figure 17. Parameter counts of several released pretrained language models

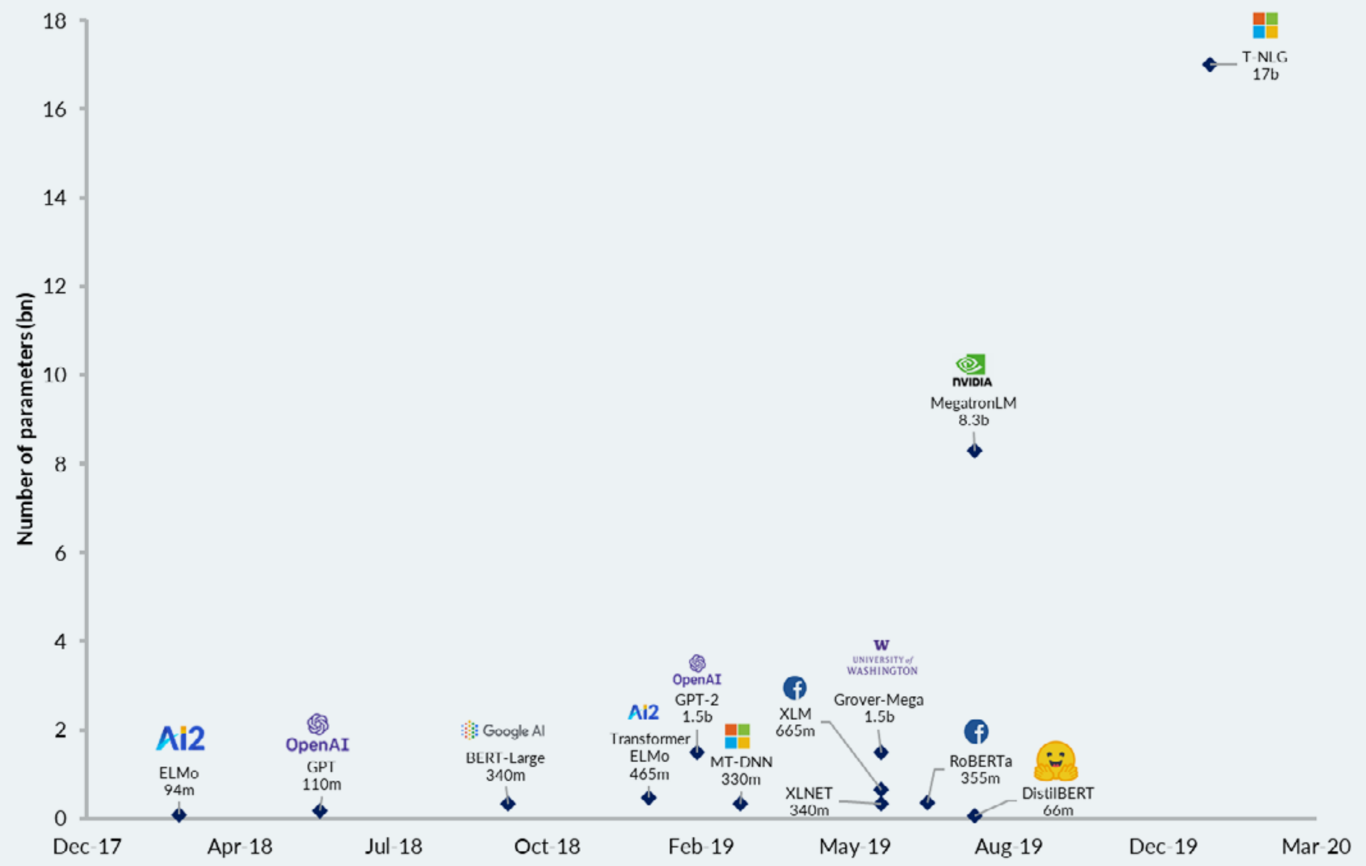

Source: Microsoft, Forsyth Barr analysis

These systems provide the basis for a new set of flexible and increasingly powerful problem solvers that can learn new behaviours continuously.

and Go. DeepMind's AlphaStar mastering the complex StarCraft II game is the result of the combination of powerful pattern recognition combined with learned strategies for optimising performance against a goal.

These hybrid systems overcome some of the known limitations of deep learning alone. New work by DeepMind researchers (Stooke, et. al, 2021) overcomes the narrow task training of previous reinforcement learning systems. It is moving in increasingly general directions for game play — including open-ended learning and rule discovery in unknown game environments. These systems provide the basis for a new set of flexible and increasingly powerful problem solvers that can learn new behaviours continuously.

Convergence of technologies

Deep learning is a very powerful technique in machine learning but it has variable ability to generalise beyond its training set data. Traditional numeric algorithms don't learn via data, but they are more general across data domains and often come with performance guarantees. The convergence of these two types of AI called Neural Algorithmic Reasoning (Veličković and Blundell, 2021) could yield systems that use machine learning to learn from data, use numerical algorithms when appropriate, are more general, and have some circumscribed performance guarantees. This is likely a productive future direction for AI systems that will yield more flexible and efficient problem solvers.

Most AI application systems are composed of a heterogenous mix of AI and non-AI software technology modules. Typically, the AI specific components comprise approximately 15% of the total system code. These systems are inevitably combinations of legacy technology, and this puts a premium on understanding and testing cross system interactions.

In addition, there is a different class of convergence which refers to powerful new combinations of technology domains amplified by AI. Examples include: synthetic genomics, AI driven materials science, flexible and stable robot balance and control systems, drones and hyperspectral sensor systems for managing crops, asset allocation and risk assessment models, and next generation security systems that act like a corporate immune system. Each of these hybrid domain application systems is more powerful than the sum of its parts.

Figure 18. AlphaZero has a near perfect score against some of the most advanced game engines in the world

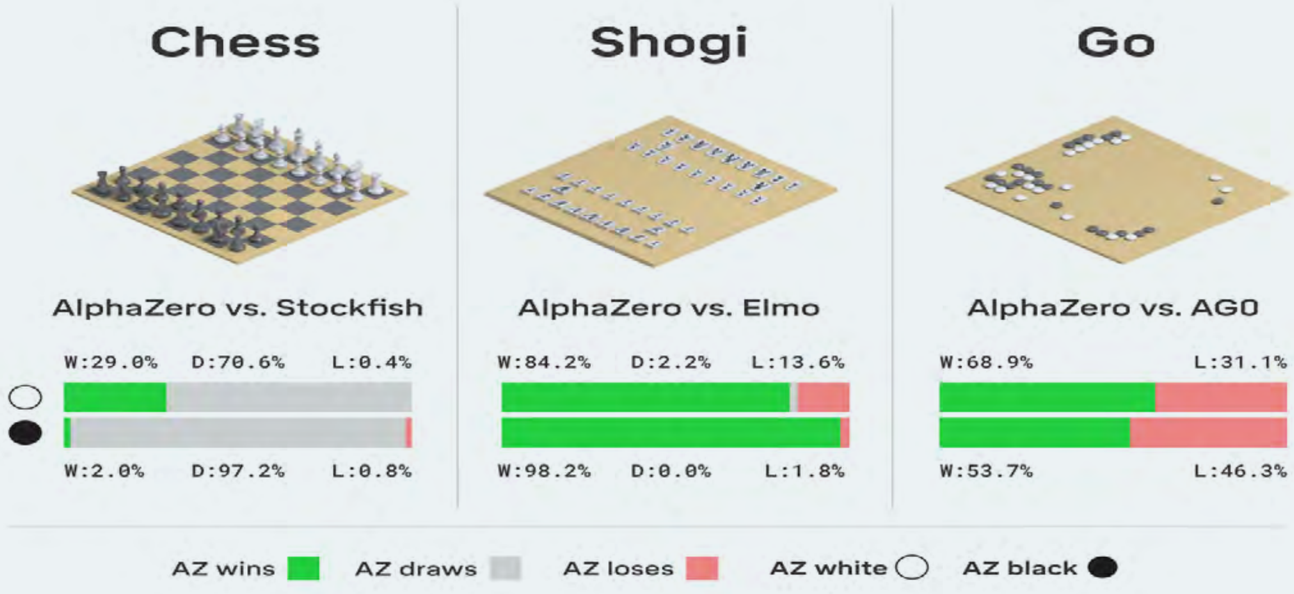

Source: DeepMind

DeepMind's AlphaZero incorporates the pattern recognition power of algorithms with the meta control of the optimisation process, making it a more flexible problem solver in learning to play chess, Shogi, and Go.

What competitive advantage can AI confer?

AI provides a range of capabilities that can confer significant competitive advantages on organisations. These advantages are application and industry specific.

There is direct linkage between AI driven competitive advantage, industry disruption, and new financial opportunities. The best opportunities are sometimes new products and services possibilities opened up by new machine learning capabilities.

For example, new tele-healthcare opportunities leverage inexpensive new consumer sensors and doctor supervised AI service avatars to keep track of patients with chronic illness (such as heart disease or diabetes) and provide near real time compliance checking and coaching. These services not only disrupt old style office visits but also offer closer health monitoring and individualised patient support.

Tele-healthcare leverages consumer sensors and AI service avatars to keep track of patients with chronic illness.

OODA loop — faster effective decisions

In nature, business, and military operations, a key competitive advantage may be framed as having an effective "OODA loop" — the ability to observe, orient, decide, and act faster than the competition. This translates to evaluating the situation and taking effective action rapidly — the opposite of analysis paralysis.

US Air Force Captain John Boyd found that when he could fly his fighter plane inside the OODA loop of his adversary, he would always win simulated fights, but if the adversary had a faster OODA loop, he would lose predictably.

Business and industry are subject to similar OODA loop competition, albeit sometimes on slightly slower time scales. Michael Lewis's 2014 book "Flash Boys: A Wall Street Revolt" documents high-frequency trading in the US financial markets (Figure 20) used as a method to front run orders placed by investors. The construction of a 1,331 km fibre-optic cable that cut straight through mountains and rivers from Chicago to New Jersey enabled competitors to reduce data transmission time from an estimated 17 to 13 milliseconds. This made a critical difference, until it was discovered as an unfair competitive market advantage.

The point is not the technology employed or the precise amount of latency reduction to improve the OODA loop. These change rapidly in an exponential world. The pattern that doesn't change is the enduring competitive advantage from rapid situation assessment, decision making, and execution. AI augments humans in this decision loop. When data transmission time differentials are minimised, competitive advantage shifts to inhumanly rapid situation assessment, decision making, and effective execution.

Note that this is not all about speed in the reaction time sense, but also about discernment: the observe, orient, and decide part of the loop. These factors are often minimised in analysis of OODA advantage but they are a critical factor — speed without effective signal interpretation and situation assessment could just make a situation worse. The competitive advantage comes from each of these factors working together accurately.

Sometimes competitive advantage is a function of recognising big patterns over a long time horizon and being able to predict outcomes that matter. Consider Moore's Law, and the competitive advantage of Apple prior to the iPhone of being able to predict that computing price performance would double every 18 months.

Apple was able to invest in a multi-year R&D project that started as totally unrealistic in terms of cost, form factor, and computing speed. That OODA loop driven bet was not made in milliseconds, but it contributed to Apple becoming a significantly bigger and more profitable company and being an early adopter of embryonic AI assistant technology — Siri.

These competitive advantages have a tendency to accumulate like compound interest over time. The good news for AI startup companies is that the inertia associated with large size and industry tenure can blind an established company to new competitive possibilities. This contributes to an impaired OODA loop.

An example can be seen in the classic competition between previous industry leader, Blockbuster Video, and startup, Netflix. Just a few years after Netflix offered to sell to Blockbuster, Blockbuster was bankrupt and Netflix was the undisputed home video industry leader. It was able to augment its innovative business model with an AI driven customer specific movie recommendation engine, which compounded its competitive advantages.

Types of AI competitive advantage

An analysis of 360 AI applications that were selected for the Innovative Applications of AI Conference indicated common sources of competitive advantage (Jacobstein, 2016):

- Augments human productivity and skills
- Discovers new patterns
- Improves prediction accuracy
- Accelerates process timing
- Solves complex problems quickly
- Improves product and service quality
- Decreases total product and service costs
- Manages corporate task and domain knowledge
- **Expands the range of the possible!**

Figure 19. AI can provide organisations with an enhanced OODA loop, providing meaningful competitive advantage

Source: Forsyth Barr analysis

Figure 20. The rise of high-frequency trading

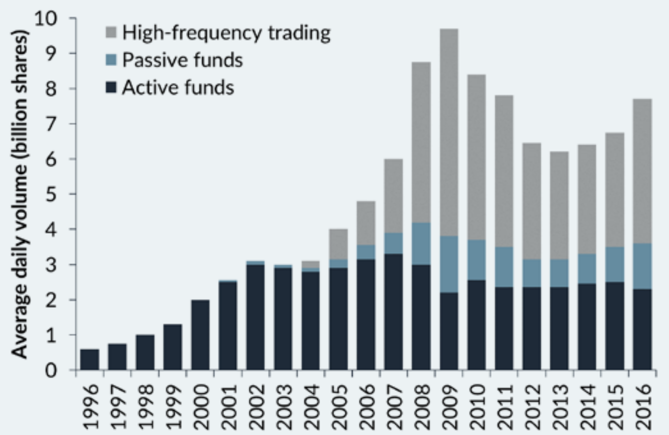

Source: Forsyth Barr analysis, Credit Suisse Trading Strategy

Figure 21. AI application sponsors

Source: Forsyth Barr analysis

Figure 22. Porter's five competitive forces

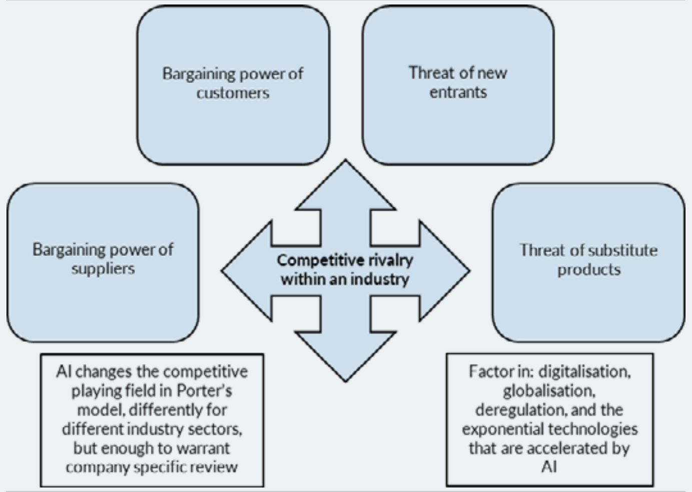

Source: Forsyth Barr analysis

Figure 23. Cyber security impact on businesses

Source: Forsyth Barr analysis, Statista

Figure 24. Cyber security market growth

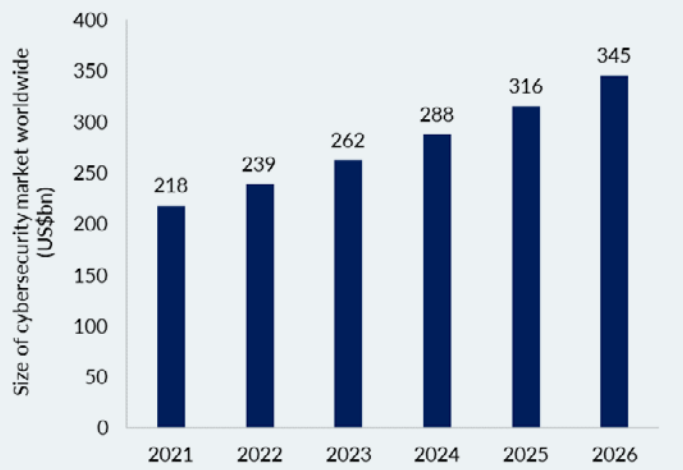

Source: Forsyth Barr analysis, Statista

Industries and governments now recognise the power of the competitive advantages that AI and machine learning can confer, and this is reflected in widespread AI application sponsorship by some of the world's largest companies.

Porter's five forces is a simple framework for assessing and evaluating the competitive strength and position of a business. The five forces are: (1) the threat of new market players, (2) the threat of substitute products, (3) the power of customers, (4) the power of suppliers, and (5) industry rivalry which determines the competitive intensity and attractiveness of a market.

To assist in determining corporate strategy, companies can analyse their competitive position with respect to each of Porter's five forces.

Porter's five forces competitive matrix

Some have argued that Porter's model is obsolete due to the increasing significance of digitalisation, globalisation and deregulation. However, these factors influence the power of the other five factors.

Exponential technologies such as AI and Robotics have also changed the playing field and will continue to do so. However, there is value in understanding businesses in terms of Porter's seminal forces, and then determining the influence of digitalisation, globalisation, deregulation, and the many exponential technologies that are significantly accelerated by AI. An example would be AI and machine learning's influence on genomics, drug discovery, and drug delivery.

AI powered cyber security

It has slowly dawned on both governments and industries that as they have adopted software and automated solutions to their critical products and services, they have simultaneously become dangerously vulnerable to cyber security risks. Artificial intelligence and machine learning algorithms have broadened the surface area and depth of security vulnerability. This used to be the subject of dinner party reflections and what if scenarios by computer specialists. Now, it is front page news and top of mind for executives and boards.

Malware is shorthand for "malicious software". It is typically AI powered code that's intended to cause damage to systems, steal data, gain unauthorised access to networks, or otherwise wreak havoc.

Malware is often powered by AI, and it refers to several variants, including:

- Viruses — software agents operating like biological systems, computer viruses attach themselves to files, replicate, and spread to other files. They may delete files, force reboots, join systems to a botnet, or enable backdoor access to infected systems.
- Worms — are like viruses but without the need for a host file. Worms infect systems directly and reside in memory, where they self replicate and spread to other systems on the network.
- Backdoors — are a method of bypassing normal authentication or encryption. Backdoors are used by attackers to secure remote access to infected systems, or to obtain unauthorised access to privileged information.
- Trojans — are agents that may disguise themselves as a legitimate application or hide within one. They open backdoors to give attackers easy access to infected systems, often enabling other malware.
- Ransomware — originally ransomware was designed to encrypt data and lock victims out of their systems until they paid the attackers a ransom in order to restore access. Modern ransomware tends to exfiltrate copies of the victim's data and threaten to release it publicly if their demands are not met. This increases pressure on victims, as stolen data often contains private information of customers and employees, sensitive financial details, or trade secrets. Ransomware can be used to find all files of a specific type locally and across the network, encrypting — and often stealing — them. The original files, recovery points, and backups are then deleted to prevent users from restoring the system on their own. Ransomware usually changes the file extension, (e.g. myFile.doc.encrypted) and thoughtfully adds a "help" file, explaining how victims can pay to recover their data. Given that ransomware is increasingly AI powered and sophisticated, it behooves individuals, corporations, and governments to be proactive about deploying the latest generation of security software.
- Phishing — attackers use deceptive AI powered communications — email, instant messages, SMS, and websites — to impersonate a trustworthy person or organisation, such as a legitimate business or government institution. New AI and machine learning techniques make this increasingly possible. Taking advantage of users' trust, attackers trick users into clicking malicious links, downloading malware-laden attachments, or disclosing sensitive personal information. A more focussed approach is "spear phishing", in which attackers target a specific individual — or a small group of individuals, such as employees in a specific role at a specific company.
- Distributed denial of service (DDOS) attacks — target servers, services, or networks to disrupt traffic flow, preventing users from accessing these resources. DDoS attacks are often intended to cause financial or reputational damage to an organisation or government body. These attacks often use AI powered bot networks of malware-infected systems — both computers and IoT (internet-of-things) devices — that the attacker controls. Attackers use these botnets against servers or networks, having each bot send repeated requests to the target's IP address. This ultimately causes the server or network to become overloaded and unavailable to normal traffic.
- SQL injection (SQLI) — Structured Query Language (SQL) is a standard language for building and manipulating databases, often used in web and other servers. AI driven SQLI attacks insert malicious SQL code into a server, manipulating it to display database information that the attacker shouldn't be authorised to access. This information may include sensitive corporate data, user credentials, and employees' and customers' information. A malicious actor (human or AI) could carry out an attack simply by submitting an SQL command into a vulnerable (not properly protected) website's search box, potentially retrieving all of the web app's user accounts.

AI is cited as both a source and potential solution to these cyber security problems. Both points of view are true. AI and machine learning algorithms are central components of virtually all modern computer malware, viruses, worms, and denial of service attacks. And computer security companies have adopted AI and machine learning techniques as the most potent tools in their defence armamentarium. This is a classic "predator prey coevolution" scenario — security systems and malware are locked in an endless and largely unavoidable war of rapid evolution of attack and defence. The practical implication of this is that both corporations and governments are now obligated to avail themselves of the most sophisticated and continuously updated AI powered cyber security systems.

AI and machine learning algorithms are central components of virtually all modern computer malware, viruses, worms, and denial of service attacks.

A wide range of deep learning applications are being adopted by industry

Machine learning is being utilised heavily across many industries. In everything from agriculture, education and fashion to credit scoring, insurance and healthcare, machine learning has become a critical success factor for many companies.

Customer relationship management

Salesforce has a leading market share in the customer relationship management (CRM) space and the resources to match. Lead prediction and scoring are among the greatest challenges for even the savviest digital marketer, which is why Salesforce is betting big on its proprietary Einstein machine learning technology. Salesforce's Einstein allows businesses that use its CRM software to utilise machine learning in analysing every aspect of a customer relationship — from initial contact to ongoing engagement touch points — to building much more detailed profiles of customers and identifying crucial moments in the sales process. This includes much more comprehensive lead scoring, more effective customer service (and happier customers), and more new business opportunities.

Education

Duolingo is a free language learning app that's designed to be fun and addicting. Although using Duolingo feels a bit like playing a game on your phone, its effectiveness is based on peer reviewed machine learning research. Using data collected from user answers, Duolingo developed a model of how long a person is likely to engage in a new lesson or remember a certain word from a prior lesson before needing a refresher. Armed with

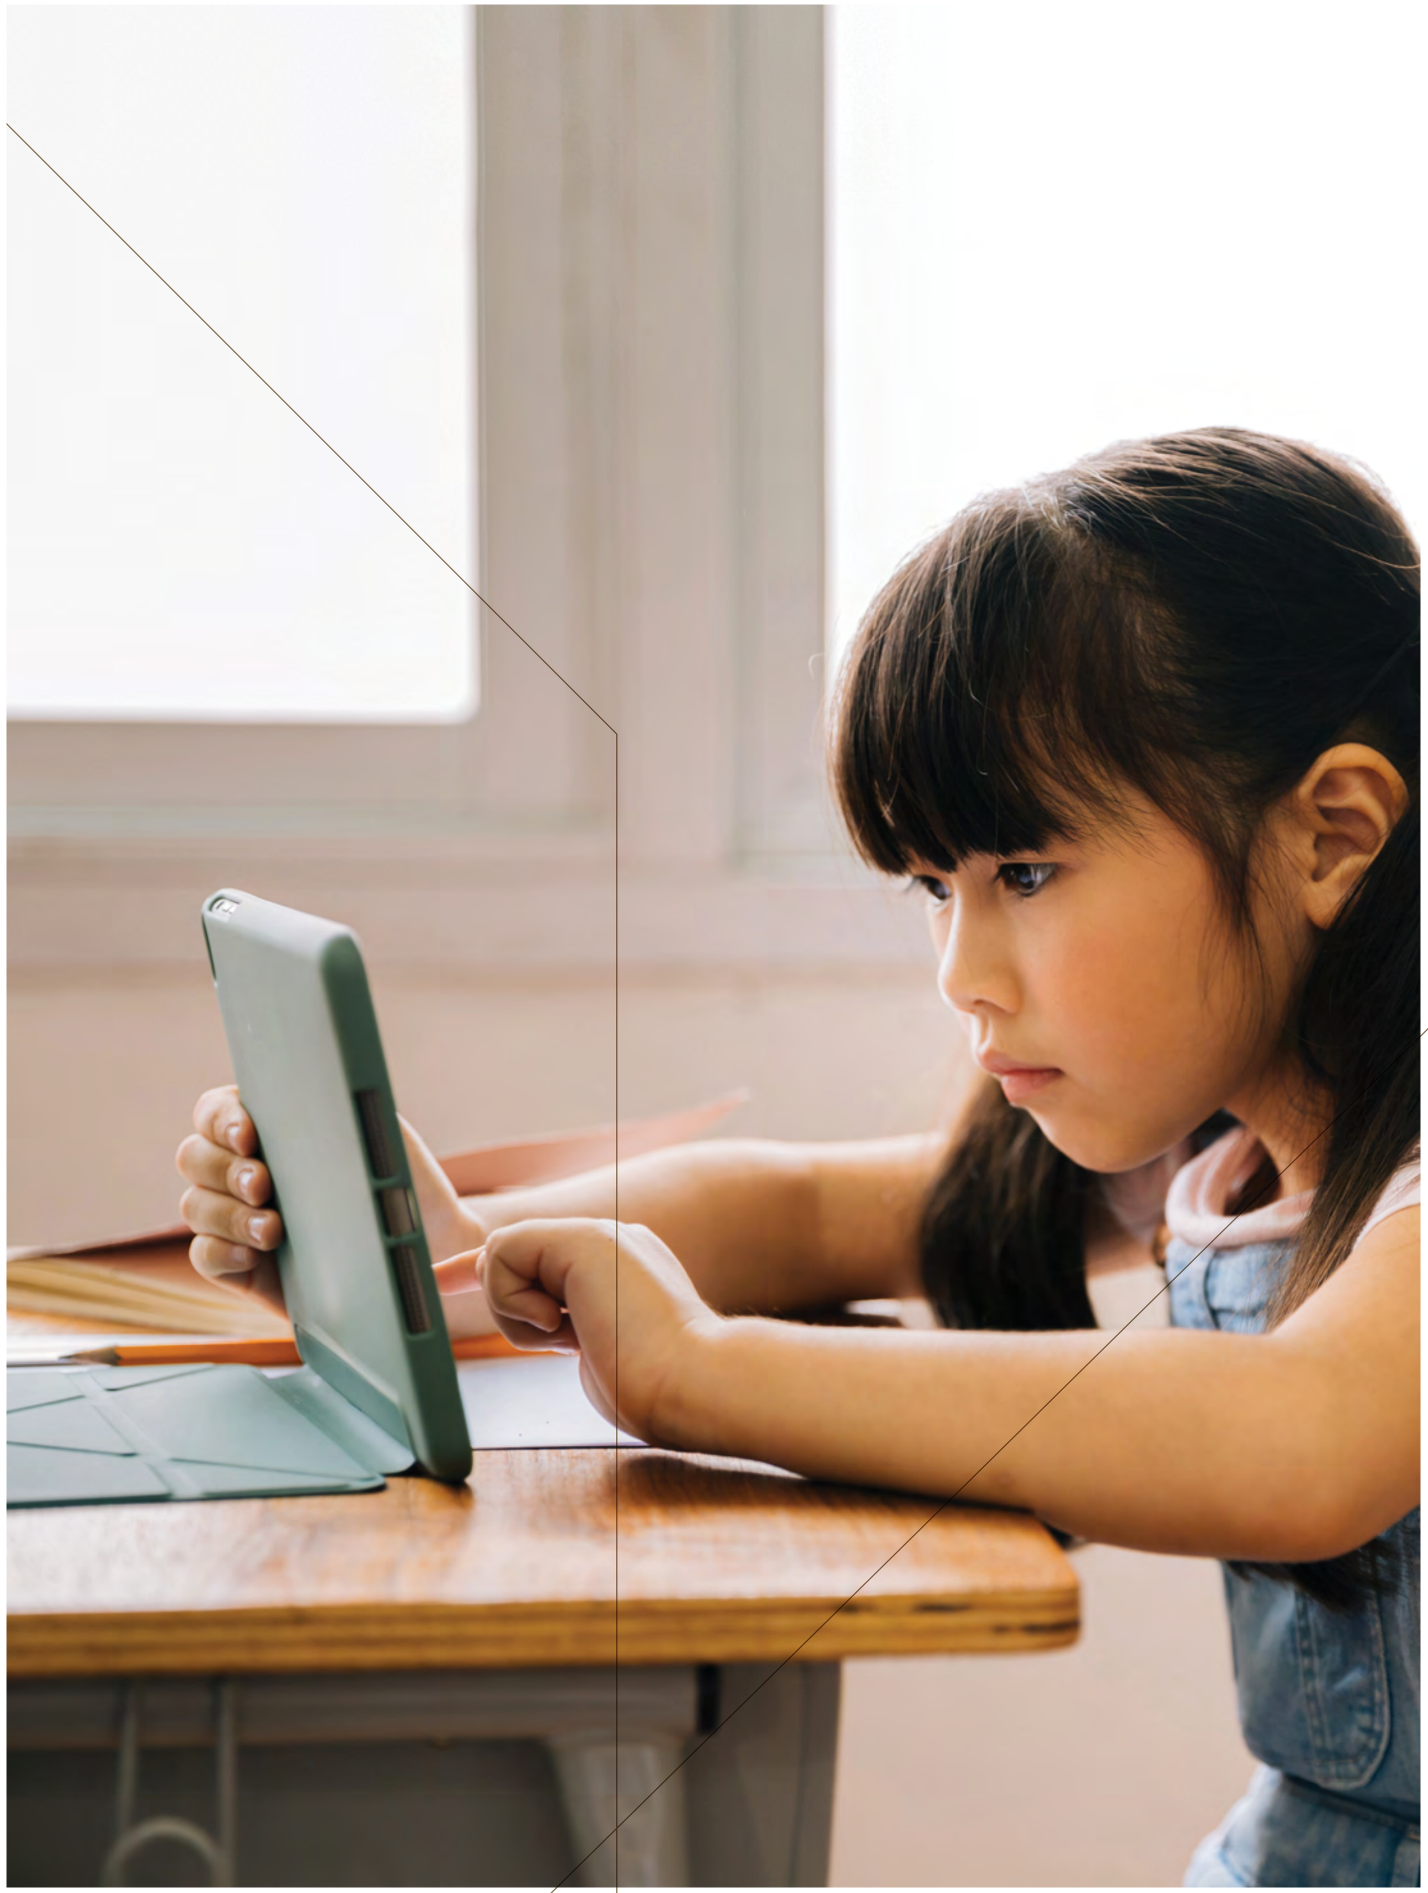

that information, Duolingo knows how to present new information in compelling ways and when to ping users who might benefit from retaking an old lesson.

Fashion

Fashion retailer Asos uses machine learning to analyse Customer Lifetime Value (CLTV). This metric estimates the net profit a business is likely to receive from a specific customer over time. In Asos' case, CLTV shows which customers are likely to continue buying products from Asos. Once that is determined, Asos can prioritise high-CLTV customers and attempt to influence them to spend more the next time around. Retailers can end up losing money on low-CLTV (with free shipping or ignored marketing promos), and this model increases the probability that Asos will turn a profit with each customer purchase.

Internet

Yelp's crowd-sourced reviews cover everything from restaurants, bars, doctors' offices, gyms, concert venues and more. Besides giving a merit star rating and a written review, Yelpers are encouraged to include pictures of the business or service they're reviewing.

Yelp hosts tens of millions of photos and uses machine learning to classify and sort them. When you look up a popular restaurant on Yelp, images are sorted into groups: menus, food, inside, and outside views. That makes it easier for people to find relevant photos rather than sorting through all of them.

Quora uses machine learning in multiple ways but the most important is to determine which questions and answers are pertinent to a user's search query. When ranking answers to a specific question, the company's machine learning algorithms take into account thoroughness, truthfulness, reusability and a variety of other characteristics in order to always give the "best" response to questions.

Financial services

Traditional credit card companies determine eligibility through an individual's FICO credit score and credit history. This can be a problem for students and the poor, who have little to no credit history. In light of that, companies like Deserve and Affirm are geared toward students and new credit card applicants. They calculate credit worthiness using machine learning algorithms that take into account a variety of other factors like current financial health and habits. Affirm has many partners, such as Walmart, for its point of sale credit scoring. Affirm IPO'd in early 2021 and has a current market cap of ~US$30bn.

JPMorgan Chase utilised machine learning in interpreting loan agreements. IKts Contract Intelligence AI system saved 360,000 hours of human labour interpreting contract clauses, freeing up support staff to focus on high value interactions with clients.

The insurance company Lemonade utilises machine learning in every stage of interaction with its customers: marketing, risk analysis, insurance contracts, adjudicating claims, customer service, and sentiment analysis. It not only provides competitive insurance at discount prices but reserves a portion of its profits to donate to its customers' favourite charities. This model has the potential to be disruptive and could set the insurance world on edge.

Agriculture

John Deere-owned Blue River Technology's 'See & Spray' system uses computer vision and machine learning to identify plants in farmers' fields. That's useful for spotting disease and weeds among acres of crops. Used in conjunction with drone technology, the 'See & Spray' rig can also target specific plants and spray them with herbicide or fertiliser. It's far more cost efficient than spraying an entire field and far better for environmental health exposure or runoff into rivers and streams.

Figure 25. Duolingo uses AI to optimise language learning through word memory

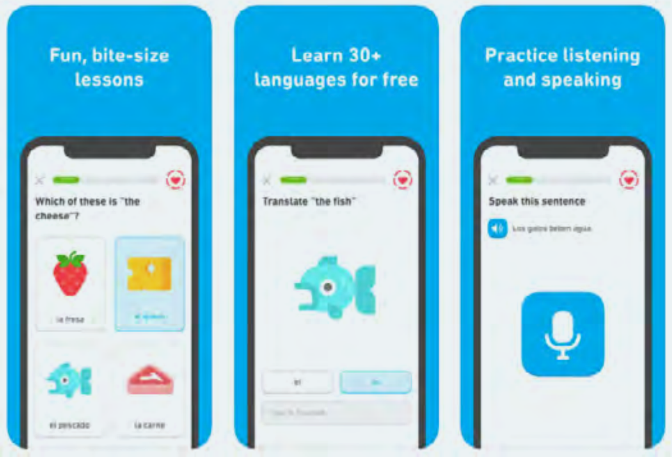

Source: Duolingo, Forsyth Barr analysis

Figure 26. Asos uses machine learning to analyse CLTV

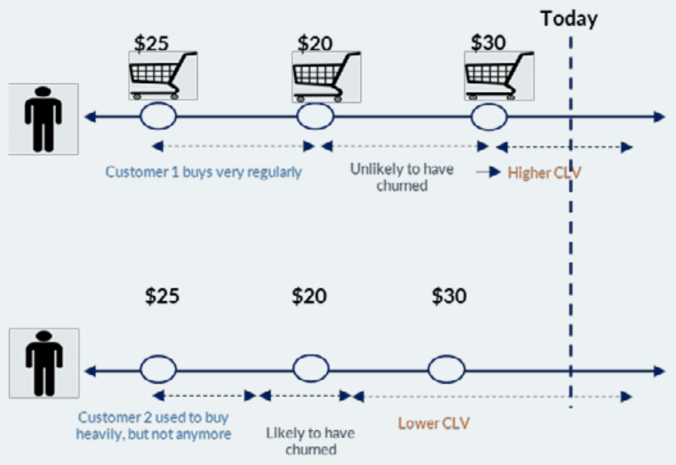

Source: Forsyth Barr analysis

Figure 27. Lemonade's rapid growth in the insurance market assisted by AI

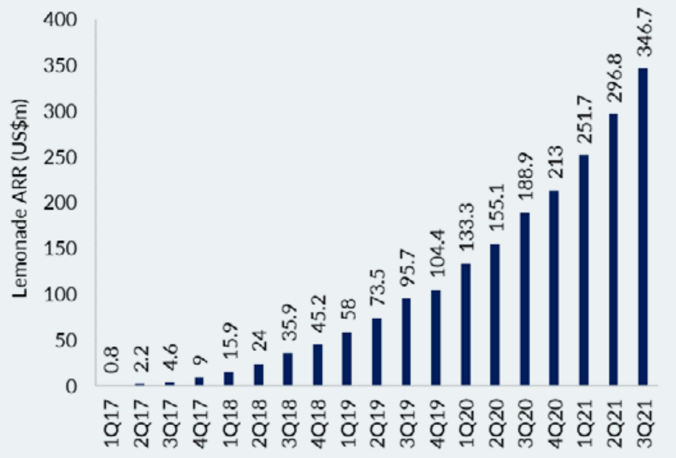

Source: Forsyth Barr analysis

Figure 28. Blue River Technology's AI powered 'See & Spray' system

Source: Blue River Technology, Forsyth Barr analysis

Figure 29. DeepMind's AI can now analyse retina images and predict smoking, age and diabetes

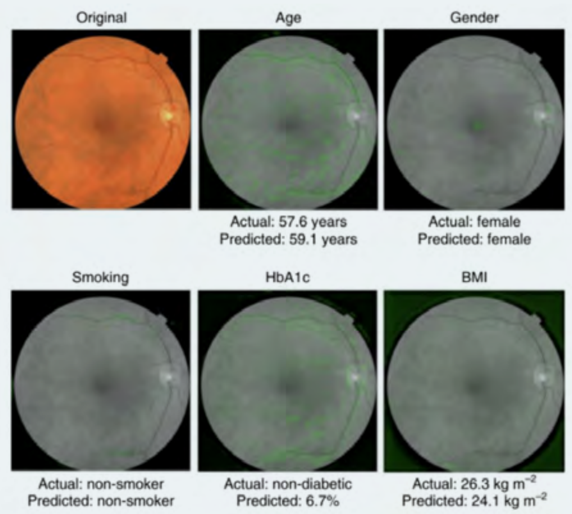

Source:UK Biobank/Google, Forsyth Barr analysis

Figure 30. World Champion Go player beaten by Google's AlphaGo AI

Source: Google Deep Mind, Forsyth Barr analysis

Consumer goods

Label Insight uses machine learning and data science to create more than 22,000 high-order attributes for retail and consumer packaged goods products. The company's "LabelSync" tool employs machine learning to give a personalised view of each food product, including ingredients, suppliers, supply and/or cold chain history, organics, fair trade, and more. This gives consumers much greater transparency and better insights into their purchases. Label Insight is setting new standards of consumer information prior to purchase.

Healthcare

In the early 1970s, Jack Meyers MD and Harry Pople built an internal medicine diagnostic system called Internist. The system was nicknamed "Jack-in-the-box" because it was presumed to include most of Jack Meyer's encyclopedic knowledge of internal medicine.

In 1977, Ted Shortliff MD, PhD published the Mycin rule based system for diagnosing blood borne bacterial diseases and recommending antibiotics. Mycin outperformed the doctors whose knowledge was emulated by the system. However, systems of this type rarely made it past the prototype or demonstration stage, except in biological instrumentation or monitoring intensive care units.

Rule based AI in medicine never really took off into mainstream use, due the large variation in human biology, the lack of a gold standard quality test, and the resulting presence of "corner cases" that had to be modified and included in extremely large and complex knowledge bases.

After the 2011 win of IBM's Watson hybrid AI system over the two world champions in the TV Quiz show Jeopardy, IBM boldly announced that Watson would next tackle oncology. However, several trials, including a very large and expensive one with MD Anderson Cancer Center in Texas, proved disappointing. AI and machine learning in medicine had a mixed track record except in pattern recognition of images. By 2009, with the improvements to the Deep Learning algorithm, machine learning was beginning to outperform the average radiologist. But this was just an example of an AI system that was doing a human task faster, better, and cheaper. What happened next was different.

In 2018, a team from the Google company, DeepMind, and Verily Life Sciences published an astonishing result in Nature Biomedical Sciences diagnosing human retina. Up until this time, expert ophthalmologists did not know it was possible to analyse a retina image and differentiate a male from a female retina, or a smoker from a nonsmoker, or estimate age, or diabetic and cardiovascular risk. The experts didn't even know what to look for, and the algorithms were not good at explanation.

This was a breakthrough example of machine learning performing a "different" task – not just the same task performed faster, better, and cheaper than humans. There is every reason to consider this a case study of the ability of machine learning to deliver results that are not just mimicking human capabilities.

In 2020 MIT researchers published a paper in Cell on the use of deep learning to screen a large database of drugs for antibiotic resistant drug potential (Stokes, et. al, 2020). The result was the identification of a new powerful and general antibiotic called Halicin. Halicin is effective against a number of antibiotic resistant bacteria. This is a stunning success story in AI augmented pharmacology.

Vulnerabilities if competitive advantages are ignored

An examination of the classes of competitive advantage possible by deploying AI and machine learning in applications makes it clear that industries are vulnerable on each one of these fronts.

The most potent source of AI competitive advantage is the ability of new entrants to fundamentally disrupt an established industry via higher quality or lower cost of products or processes. This puts a premium on established incumbent companies being proactive about experimenting with AI and machine learning, and then deploying successful results rapidly in new business models. Their choices are to disrupt themselves or be disrupted by competitors.

Gaming

Games are serious business. Global videogame revenue was estimated at ~US$180bn in 2020, according to IDC data. In contrast, the global film industry reached ~US$100bn in revenue for the first time in 2019, according to the Motion Picture Association, while PWC estimated North American sports would bring in more than US$75bn in 2020. This means that the videogame industry is a bigger moneymaker than the global movie and North American sports industries combined. Very few videogame characters and event engines are not driven by AI and machine learning algorithms today.

AI has a multi decade history with traditional games such as tick tack toe, checkers, and chess. Its influence is likely to grow with three trends: (1) increasingly realistic AI driven video game characters, (2) new uses of AI in the rapidly growing virtual and augmented reality game industries, and (3) pandemic associated changes to the traditional movie theatre industry, which will drive more compelling and interactive AI driven home movie experiences, complete with customer selected alternative outcomes — similar to games.

IBM's Deep Blue "big iron" computer was purpose built to play chess, and it scored an upset win over Gary Kasparov, the world chess= champion in 1997. In 2017, Google's DeepMind team built AlphaGo that defeated Ke Jie, the Chinese number one ranked World champion in the super complex board game of Go.

Following that win, DeepMind built AlphaZero. While AlphaGo learned Go by playing thousands of simulated matches with knowledge from human amateur and professional players, AlphaZero learned how to play world class chess, Go, and Japanese Shogi by playing against itself, starting from completely random play. In fact, AlphaZero defeated the world champion computer chess player by starting from scratch and playing against itself for four hours!

This system developed a dynamic and unconventional style of play that differed from any previous chess player. Many of its game changing concepts have since been utilised by others at the highest levels of chess.

It is important to note that Go and chess are games with perfect information — there is no uncertainty on the game board. However, the real world of people and markets are full of uncertainty, and this makes for potentially less than perfect machine learning results in complex and unpredictable real world application contexts.

In 2019, DeepMind's AlphaStar beat two world class professional players in the epically complex strategy videogame StarCraft II. This game has a lot of degrees of freedom and uncertainty. An observer to that championship game commented: "AlphaStar split its stalkers into squads and flanked MaNa's army on multiple sides…He was dumbfounded…. What we saw there, that's not human."

Although AI and machine learning have taken the gaming world by storm, there is little sign that the resulting AI systems currently even have the broad, deep, and subtle intelligence of a five-year-old human child.

Operational and ethical risks associated with AI

The technical and operational risks of AI are many and varied, and they are often linked to ethical issues and consequences. The technical and operational risks fall into a few key categories.

Misspecification — programming is about the specification of intent. However, both traditional software engineering and machine learning systems are subject to misinterpreting ambiguous intent. This is far more likely than the malicious intent attributed to the Hollywood movie "A. I. Artificial Intelligence". The consequences of misinterpretation may be trivial, financially disastrous, or an international security risk.

Humans have been dealing relatively successfully with risks in traditional software engineering for decades. The additional risk associated with machine learning is that the algorithms are more difficult to inspect, they may incorporate misspecified objectives, and the large number of potential unforeseen conditions place a higher premium on extensive testing and safeguards.

It is worth noting that pre AI business as usual concepts of operation and delivery may be the #1 financial business risk. This is because the long list of AI advantages that a competitor could employ has the potential to disrupt an incumbent non AI business of any size or longevity.

The success of AI and machine learning in the marketplace and in social media has revealed a cluster of ethical issues that have disproportionate impact on different social groups. Some of these issues are significant, such as bias in data sets, bias baked into algorithms, privacy issues, covert manipulation of buying behaviour, misinformation and disinformation, military and cyber security risks.

These issues were not front and center in the AI community until a few years ago. Many of the

There is additional risk associated with machine learning.

standards proposed in the literature have been passionate and self-righteous but not necessarily technically well informed. These considerations will impact startups, large companies, and governments. Investors should anticipate a new set of significant ethical benchmarks and standards that will be both voluntarily adopted and imposed on the AI community.

Examples include the need to inform consumers when their data are being analysed (typically by machine learning algorithms) and used to make recommendations, and explicit representations about the origins and veracity of claims made on websites. These benchmarks and standards will provide both competitive advantage to early adopters and potentially significant economic and negative consequence risk to companies that flaunt the new and emerging standards (Kissinger, et.al., 2021).

The "Trolley Problem" is a corner case in AI ethics

Discussions of AI ethics inevitably include consideration of the so called "Trolley Problem". This problem presents a runaway trolley car with an AI at a switching station that can shunt the trolley car to hit one innocent bystander, or a group of people crossing the tracks. This problem puts a human or an AI in the no win position of making an impossible or extremely uncomfortable ethical choice. The many variants of this problem with self driving cars (sacrifice the passenger or the mother with a baby carriage on the road) have generated more heat than light on the issues. The fact remains that this problem gets a lot of attention but is a statistical corner case in actual driving conditions.

Closer in and more practical is the consideration that human drivers kill 1.3 million of their fellow humans globally every year. This is not a corner case. Given that two of the leading causes of highway death are driving while intoxicated and distracted driving, it is entirely relevant. AIs embedded in robot cars do not drink or get high, and they are not subject to distraction. Insurance companies have begun to understand the economic consequences of this.

While it is true that self driving vehicles are not infallible and in rare instances they have made decisions that resulted in human deaths, such incidents are rare, and will decrease over time. Most of these systems are not fully autonomous. Further, the underlying AIs can improve continuously and learn from each other's experience and mistakes — a feat that few humans manage. When possible, both AI and human judgement should be represented in the decision calculations.

AI bias

A significant change in the AI community and society at large in recent years has been new and increased scrutiny of AI bias. It is now recognised that bias is unambiguously present in AI systems. This is partly due to training data being based on sampling a biased society. There are different classes of bias including: race, gender, ethnicity, religion. These are not mere inconveniences. The consequences are real harms to people — stereotypes, negative assessments that result in job and scholarship loss, loan denials, and harsher treatment in the criminal justice system.

Detecting bias in AI systems requires systematic review utilising AI & human judgement. For the past five years, undetected bias has carried moral implications as well as new and serious legal liability.

Many of these issues and others were addressed at an Aspen Institute Roundtable on Artificial Intelligence (Bollier, 2019). This roundtable discussed transparency in use of data, disclosure, user privacy, diversity, bias, trust, accountability when systems fail, governance, AI regulation, and complementarity between AI systems and human workers.

Military and covert AI

The recently released Final Report on Artificial Intelligence by the US National Security Commission (2021) was the result of a very broad and well-informed study of the near term economic, ethical, and military implications of global AI competition. This study has deep implications for all nations and corporations.

AI has already become a key strategic technology for international corporate and military

Figure 31. Top reasons for ethical issues in AI use

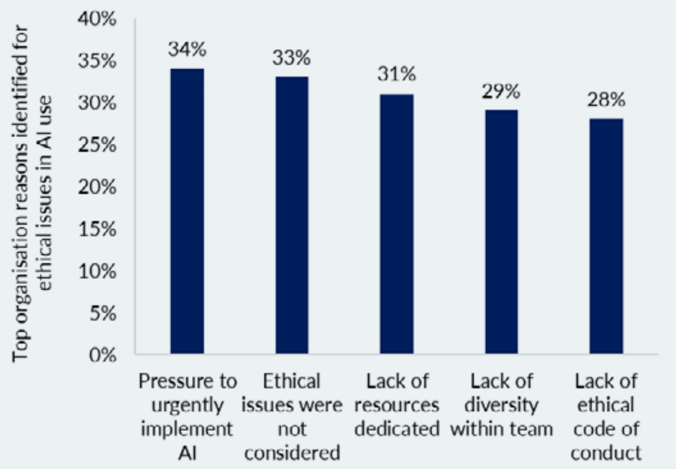

Source: Forsyth Barr analysis, Statista

Figure 32. Top AI risks considered by organisations

Source: Forsyth Barr analysis, McKinsey

Figure 33. A version of the autonomous car "trolley problem"

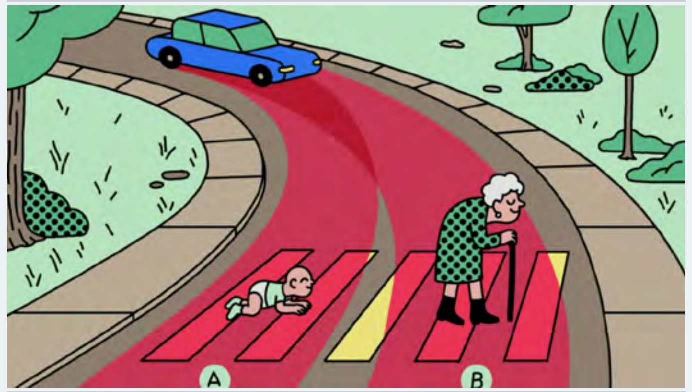

Source: technologyreview.com, Forsyth Barr analysis

Figure 34. AI bias increasingly being understood as an issue

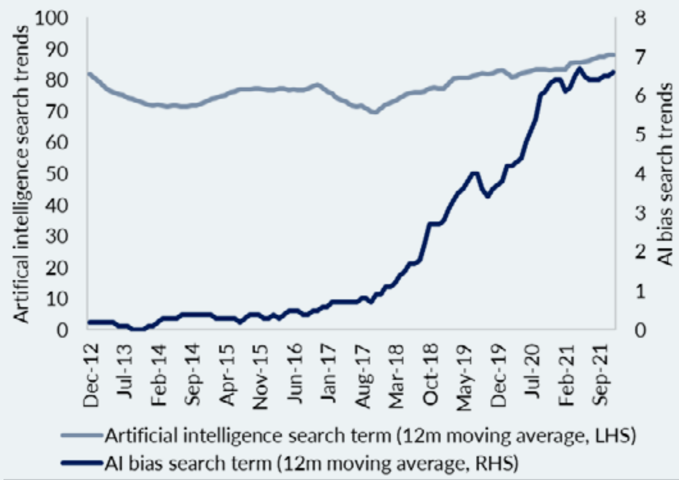

Source: Google Trends, Forsyth Barr analysis

Figure 35. China now leads the US in AI paper citations - a measure of influence

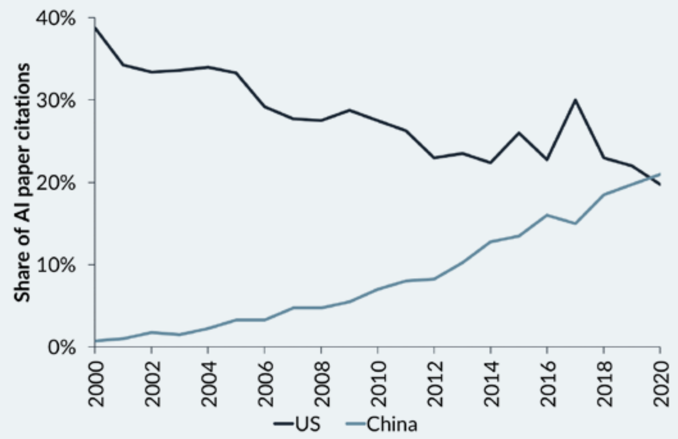

Source: Stanford University's AI Index Report, Forsyth Barr analysis

Figure 36. The US leads China in private AI investment, but China has signficant public investments in AI

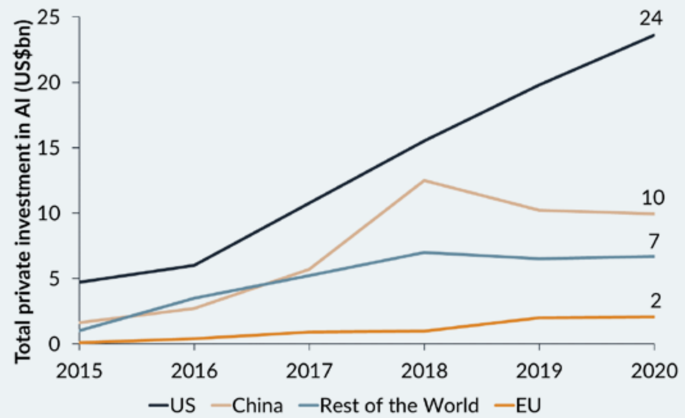

Source: Stanford University's AI Index Report, Forsyth Barr analysis competition. Although the competition is at every level in new algorithms, talent, and speed of technology adoption, from the point of view of an individual company, the risk can be near term and specific to them.

Where the US was once the undisputed leader in AI technology, China now invests more than the US, and not only publishes more AI technical papers but it leads the world today in technical paper citations — a direct measure of influence. The US has historically been reluctant to develop a national AI technology plan, but China is aggressively pursuing the goal of becoming the world's AI superpower.

It is vital that companies and nations increase their investments in AI. This is both to increase their competitive position, but also to defend against competitive threats to specific critical businesses.

New international cooperation is necessary for the rapid development of both industrial and military AI technology standards and ethical norms for key applications. In spite of this conclusion in the NSCAI report, some members of the AI community are categorically against using AI in military systems. However, a cursory review of military systems over the past 20 years indicates that the practice is already widespread, although with early vs. advanced and autonomous AI.

Many remaining issues center on the classes of AI technology deployed, and the international rules for ethical and safety tested military systems. Current ethical rules of engagement specify that potentially lethal decisions require a human in the loop. However, even in defensive situations, the velocity of military systems may, in some cases, make this objective aspirational. This may lead to attempts to build in an imperfect approximation of human values.

It is the uncertainty associated with 100% compliance that drives some of the need for these advanced systems. Focussing on defensive AI systems is not a solution to the problems, since defensive systems may have to be at least as powerful as offensive systems in order to be effective. Consider an AI augmented drone dogfight — other than explicit markings, it would be difficult to tell them apart.

The NSCAI report summarises: “The NSCAI Final Report presents an integrated national strategy to reorganise the government, reorient the nation, and rally our closest allies and partners to defend and compete in the coming era of AI-accelerated competition and conflict”. It is important to recognise that there are no simple or easy solutions, and that the "coming era" is here and now today.

AI's lack of broad judgment

When the general public learns that an AI beat the world's Go or Chess champion they tend to attribute broad rather than narrow intelligence to AI. The Go champion doesn't know what a dog, orange juice, or ethics are. It is incumbent on the AI community to communicate just how limited and narrow AI is today. However, it won't stay that way.

AI's potential to misinterpret objectives

Independent of the strength or narrowness of AI, it is always possible to misinterpret objectives set by humans. Even humans often misinterpret each other. This is a real risk and it is unethical not to address it systematically. AI systems are complex and have incompletely characterised risk — even with some testing. At a minimum, all but the most narrow and restricted AI systems connected to life support systems need to be subjected to an extensive testing protocol and run in an experimental sandbox environment before permitting access to semi-restricted or open and operational environments.

AI's potential to modify objectives

Values drift has been addressed in the AI literature as a condition where an AI system's objectives diverge over time from the humans that created them. There is an important distinction that is often overlooked in evaluating the risk of AIs modifying their own objectives. It is one thing to have an ability that outperforms a natural system. It is another thing to want to emulate it in all respects.

Canadian psychologist and science author, Steven Pinker, wrote: *“The scenario [that robots will become superintelligent and enslave humans] makes about as much sense as the worry that since jet planes have surpassed the flying ability of eagles, someday they will swoop out the sky and seize our cattle. The ...fallacy is a confusion of intelligence with motivation — of beliefs with desires, inferences with goals, thinking with wanting. Even if we did invent superhumanly intelligent robots, why would they want to enslave their masters or take over the world? Intelligence is the ability to deploy novel means to obtain a goal. But the goals are extraneous to the intelligence: being smart is not the same as wanting something"* .

In spite of this, although we do not have AI systems today that modify their own objectives, in the decades ahead it is not outside the realm of possibility that they could gain this capability — independent of the motivation to achieve specific objectives. This could lead to a divergence of alignment between human objectives and emerging machine intelligence objectives.

A potential source of future drift (currently science fiction) could center on climate change. People know that they are "working on" climate change challenges — even if these efforts are suboptimal and currently on a failure trajectory. AI planning or resource management systems may formulate climate control objectives that are considerably harsher than humans are willing to accept. No malevolence is required for this. Some of these systems may be connected to critical infrastructure such as the electrical grid. Thus, this needs to be considered from a safety and system monitoring point of view. There are already scores of AI researchers around the world considering how to constrain these more capable future systems effectively and reliably. These strategies include built in redundant and encrypted safeguards, and multiple layers of outbound isolation from information and power grids.

AI application developers' ethics

AI systems are ultimately designed and built primarily by humans, even when those humans are augmented by AI. Corporates and governments want accountability for the safe and ethical conduct of AI systems. When it comes to responsibility, no matter how complex the code, customers metaphorically want "one neck to choke" if something goes wrong. Currently, AIs have no legal standing, so that neck will be the human(s) in charge of the entity supplying the AI. Ignorance of the detailed operation of the code will not provide cover. Like the owner of an occasionally vicious dog, the owner is responsible for the dog's behaviour, even if it is not clear why the dog bit a stranger. So the humans supplying AI services and software need to be proactive about the design, application, and testing of the ethical operation of their systems.

The above considerations assume that the developers and AI service suppliers are honest — and most are. Today, most standard AI and computer science programs include an ethics course. While the cases covered vary with the non standard course material, one non technical financially driven case is often overlooked in the plethora of technical corner cases. If an AI developer's boss asks the coder to do something dishonest or clearly illegal in developing code or representing its capabilities and risks, it is the coder's responsibility to say "No", even at the risk of losing their job.

All of these potential ethical issues have prompted some to call for professional standards and government regulation of AI (AI Forum, 2020). There are issues around how operational and practical professional AI guidelines actually are, but a good working assumption is that reasonable guidelines are better than none. However, there are bound to be entities that don't participate in professional AI guidelines or regulatory regimes — no matter how accepted or reasonable. Further, even for those entities that do sign up for professional standards or regulations, there may be defectors that opt out when deemed convenient or necessary for them.

For this reason, there are two medium term potential risks:

1. Regulating AI into non competitiveness
2. Innovating AI into amoral disasters

Addressing these twin risks requires a nuanced best practices approach to AI that honors both ethical and technological imperatives.

Managing AI projects – roles of AIs and humans

There is a tendency for AI development teams tasked with automating a service currently provided by humans to assume that the AI system must do all the end-to-end subtasks of the new service. However, although machine learning systems may outperform humans on data driven analytical tasks, they may be quite weak in dealing with difficult customer service problems.

Rather than try to solve the "last mile" of customer service, it might be wiser to constrain the AI project specification to those tasks that machine learning does exceptionally well and leave to humans the more difficult parts of the "high touch" interactions with consumers. This may not be necessary in the long run, but in the short run, it could avoid project expenses and customer aggravation.

Internal vs. external project trade-offs

Large companies are often faced with a difficult decision at the beginning of an AI project. If they don't have deep internal AI expertise, they may be tempted to release their mission critical project out to bid by an external contractor. Even if the contractor does have deep AI expertise, this strategy leaves the company vulnerable to inflated project costs and long term dependency on the external development team for system maintenance and extension.

On the other hand, building internal AI development capabilities comes with its own risks that include: delays and difficulties with recruiting AI talent, a tendency for internal development

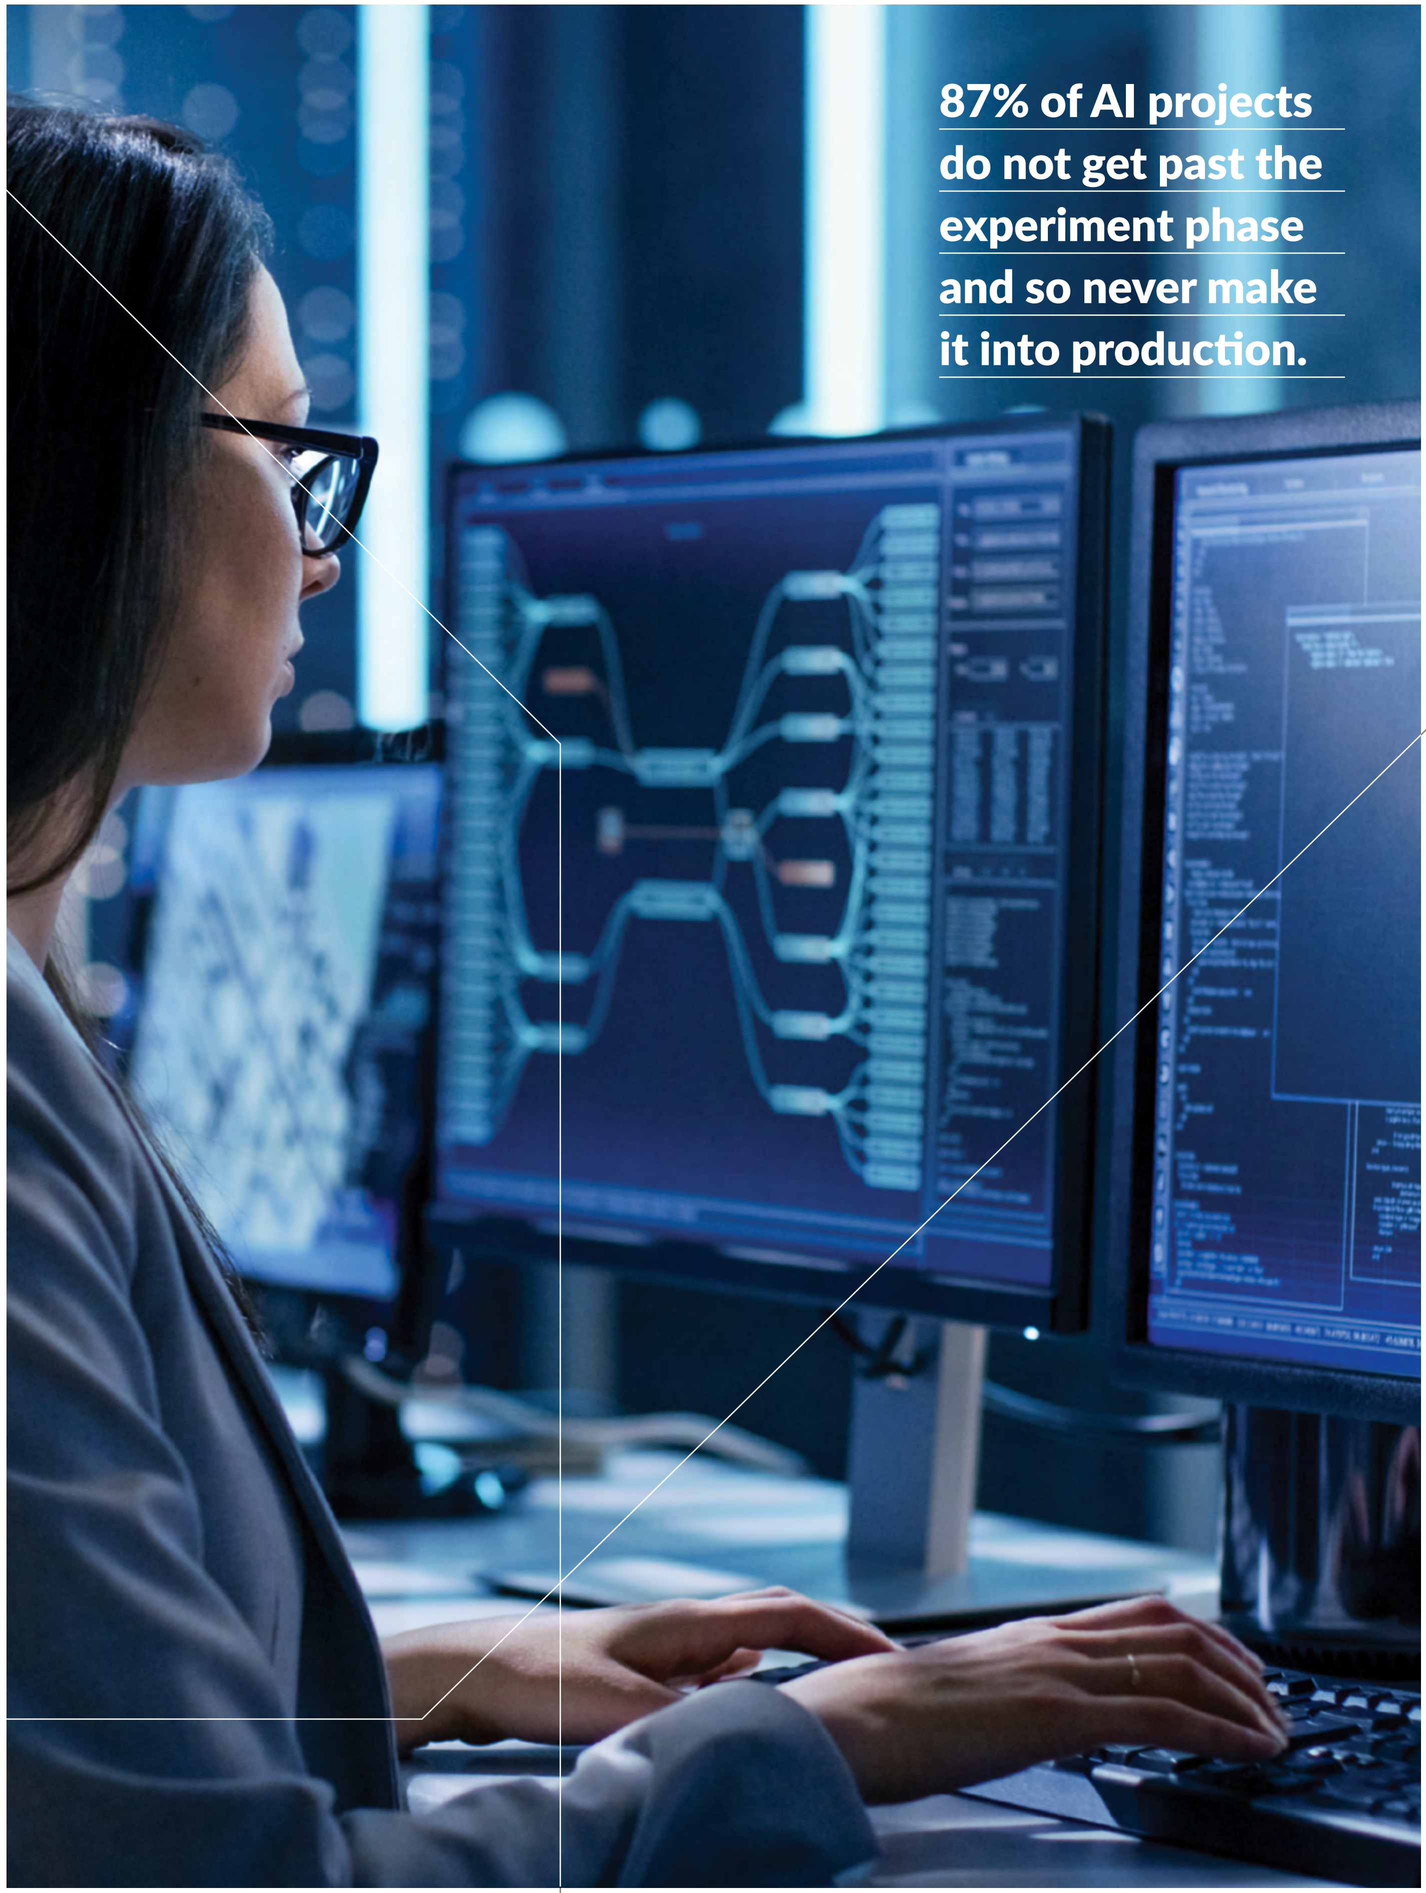

87% of AI projects do not get past the experiment phase and so never make it into production.

teams to build every system component from scratch, limited team experience on heterogeneous projects, difficulties managing the internal team productively and holding them to industry standards, and problems replacing critical members of a relatively small team.

Given risks with external contractors and internal development teams, it is prudent to hire an experienced internal AI project leader to assign the components of the project to an internal team that will provide basic data and interface systems, and an external team that specialises algorithms for the company's specialised needs. Depending on the company and project size, a small AI experienced Advisory Board may review project plans and progress to ensure sound decision making and help avoid costly mistakes.

Andrew Ng and Neil Jacobstein discussed these factors and job risk during a 2017 Wall Street Journal sponsored public event in San Francisco (Jacobstein and Ng, 2017) .

Buy most — only build for truly specialised needs

If a company has a generic need for particular system components, such as customer management services provided by cloud based commercial systems (such as Salesforce and its competitors), the build vs. buy decision may tilt towards buying that component. However, if the company has truly specialised needs, the balance may tip towards building and owning that layer of an application stack.

Ultimately, these decisions need to be made by experienced project leaders, keeping the long term interests of the company in mind.

On premise vs. cloud computing decisions

AI applications drive some unique tradeoffs between hosting solutions on traditional business infrastructure on premise vs. hosting solutions in a remote cloud computing provider. The stack of machine learning solution software can be computing intensive, complex to setup, and challenging to maintain. Cloud solutions such as AWS from Amazon, Google Cloud, or Azure from Microsoft come with scalable AI development and deployment frameworks that respond to distributed global use on demand.

Given the centralised aspects of cloud AI solutions large data test harnesses and security layers can be easier to provision and manage. Each large application is different, and the subcomponents of the solution may be distributed in a principled manner across cloud, on premise, and hybrid infrastructure. However, it is typically unwise to build AI application solutions from scratch without evaluating the potential for major acceleration by assembling solutions out of pretested framework modules.

AI applications framework

Enrique Dans observed in a July 2019 Forbes article on AI that "87% of projects do not get past the experiment phase and so never make it into production". That dismal percentage could be improved dramatically if AI application development teams took advantage of the decades of experience embedded in application development frameworks. These frameworks codify what tends to work in software, but a large remaining challenge is getting the strategy, team, training, and internal culture right for implementing AI applications.

The most important element for companies that want to increase their use of AI for competitive advantage is to ensure that their business operations are instrumented and open to change. Utilising AI in an ossified business environment does not automatically lead to productive outcomes. Businesses are well advised to seek outside consultation and follow up with workshops to prepare their internal business culture for AI and exponential technological change.

Figure 37. AI applications framework

1. Develop competitive AI strategy with measurable results
2. Create pilot projects — leverage AI and human strengths
3. Build on machine learning application platforms, not from scratch
4. Create hybrid internal/external AI teams
5. Provide free AI and application specific training
6. Develop a sustainable data & applications community
7. Manage AI applications responsibility: use guidelines

Source: ©Jacobstein, 2020

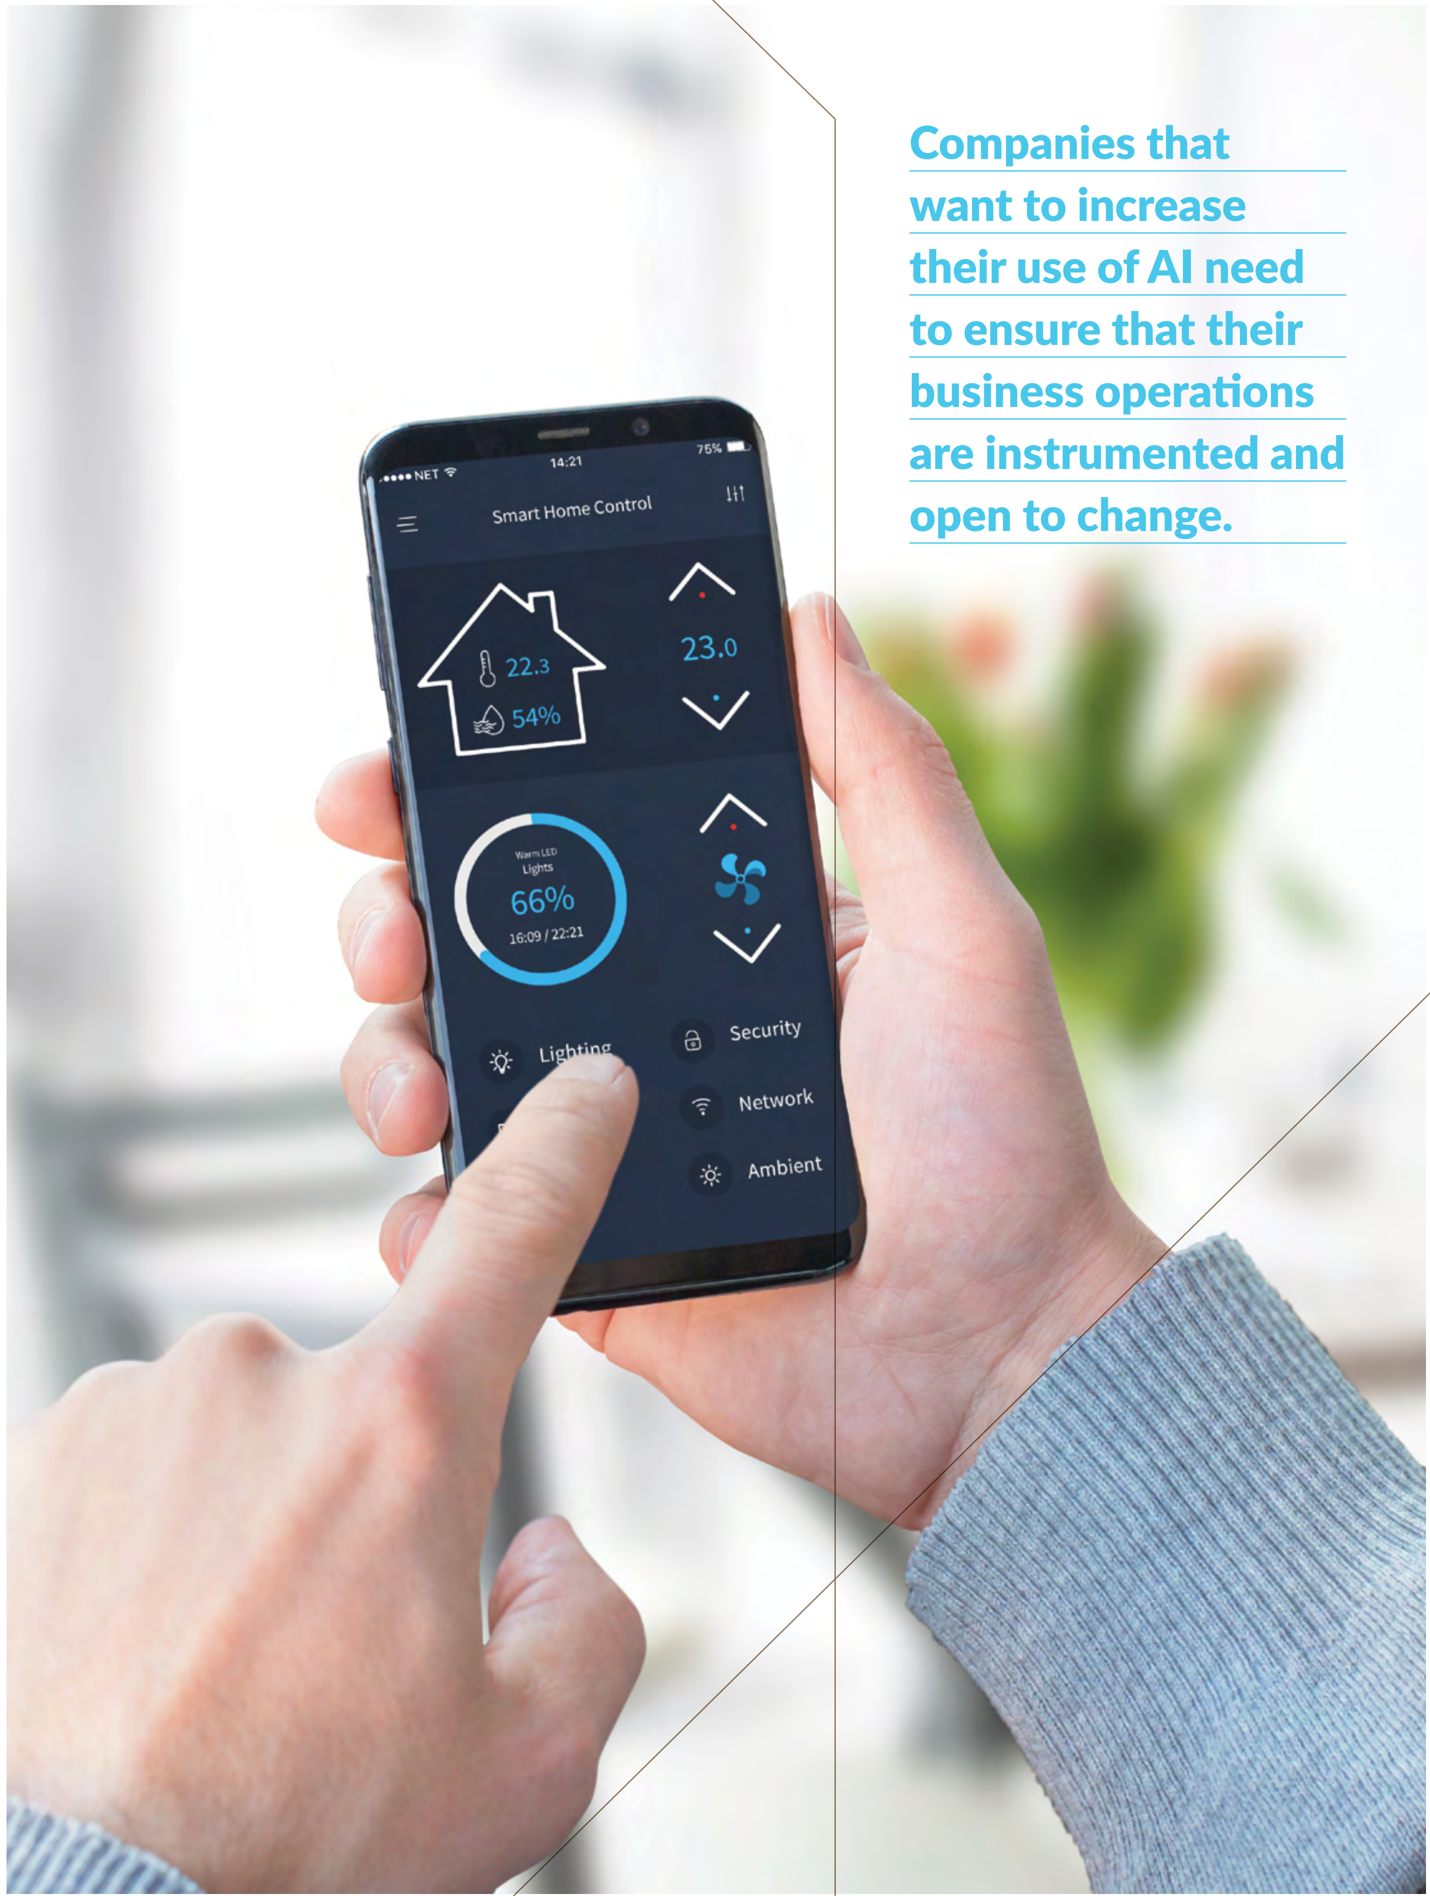

Companies that want to increase their use of AI need to ensure that their business operations are instrumented and open to change.

How competitive is New Zealand in applying AI?

New Zealand punches above its size class in applying AI and machine learning. A New Zealand team from Auckland University developed the R statistical and machine learning programming language. The R system was the leading machine learning application language several years ago. It has since been partially surpassed worldwide in popularity and support packages by Python and other systems. However, New Zealand has been a fast follower of AI development components and applications.

Adoption of AI is still at an early stage worldwide, in spite of the dramatically increasing power of the technology. The price performance of AI is being driven by three factors: (1) advances in computing speed and storage per unit price, (2) continuous improvement of machine learning algorithms, and (3) online access to massive and relevant data. Internationally, adoption of AI and robotics on average is still in the low single-digit percentages. However, it would be unfortunate for New Zealand companies to squander this potential competitive opportunity.

It appears that China and the US have recognised the economic and strategic significance of AI and machine learning. They and other countries are currently engaged in a type of arms race competition for talent, technology, data scale, and application effectiveness.

China's vision and use of AI

China has not been shy in announcing its use and commitment to AI. It invests more in AI annually than the US. It patents, publishes, and is cited more frequently (refer to Figure 35 earlier). The biggest innovations in AI are still not coming out of China, but that is likely to change over the next decade.

Taiwanese computer scientist, Kai-Fu Lee, and Chinese science fiction writer, Chen Qiufan, recently published "AI 2041: Ten Visions for our Future in 2021". The ten visions were projections of AI technology maturing over time and plausible narratives around the forecasts. The forecast topics were:

1. Narrow deep learning insurance applications and externalities.
2. Deep learning finance applications and the externalities of narrow objective functions.
3. Predator prey cat-and-mouse cycles in visual Deep Fakes and fake detection.
4. Video fakes and fake detection, facial pattern recognition, multiple uses of computer vision in security.
5. Natural language processing, AI in education.
6. AI companions in education, Artificial General Intelligence, Consciousness.
7. Contactless love post pandemic, automation acceleration, precision digital medicine; living healthier longer.
8. Games and immersive AR/VR, brain computer interfaces.
9. Autonomous vehicles, smart cities, ethical issues.
10. Quantum computing, Bitcoin security, Autonomous weapons, Universal basic income.

The key issue is not whether the forecasts of Kai-Fu Lee and Chen Qui-fan are correct or not. Rather, they are intended to prepare Chinese society and the wider world to deal with the many implications of AI in the future. Given the huge discrepancy in size and scale with other leading AI powers such as the US and China, how is New Zealand doing on its own terms, at least as reflected by leading listed New Zealand corporates as determined in our 2021 Corporate AI survey?

Inaugural Forsyth Barr Corporate AI Survey — five key takeaways

With the increasing importance of AI both as a tool for success and as potential source of risk we have conducted an AI survey of NZ's mid and large cap listed universe. The response rate was over 90% and we received responses from almost 60 different companies. Our survey suggests a familiarity with AI at a basic level, but there is no sense of urgency with regards to deepening knowledge and increasing its use. Overall, the survey respondents were generally positive about AI applications, but most had not seen major results yet. There was limited indication of a keen sense of urgency that can sometimes be seen in Chinese, American, Israeli, and other corporations.

We walk away with a number of broad conclusions. The majority of companies do not have a defined AI strategy. The impression overall is one of curiosity rather than confident application. The majority of companies have invested directly in AI, but the investments are modest, generally below NZ$1m p.a. Expectations of AI going forward is overwhelmingly positively skewed. The risks around complexity, ethics and competition is dwarfed by a belief that AI will positively contribute to revenues and drive costs down.

#1 NZ corporates are dipping their toes but not diving into AI

Given the strategic significance of AI and the intense international competition for talent and application results, one of the most important survey findings was that only one in seven stated that they had a defined strategy for AI, and a further 30% answered that they had it "to some degree". The limited presence of a defined AI strategy contrasts with the over 60% of respondents that have invested in AI. To us, this suggests that NZ Inc is still in the trial and error phase and companies haven't properly figured out what to do with it yet. This impression is further emphasised by the limited resources dedicated to AI. Only one in five companies have more than one person dedicated to AI and even if we included external spend and look at total investments in AI, the vast majority spent less than NZ$1m. That is certainly not enough to be transformative, in our view.

Figure 38. What is your role within your organisation?

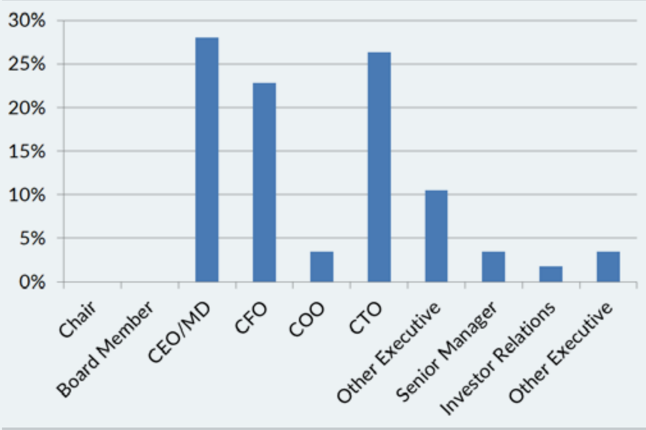

Source: Forsyth Barr analysis

Figure 39. Does your organisation have a defined artificial intelligence (AI) vision and strategy?

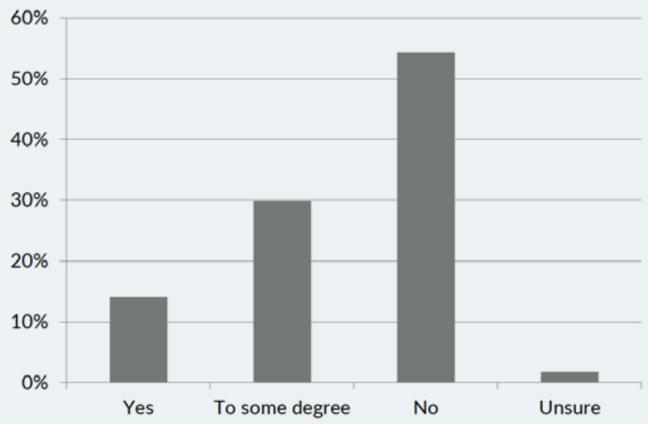

Source: Forsyth Barr analysis

Figure 40. Has your organisation directly invested in any AI applications?

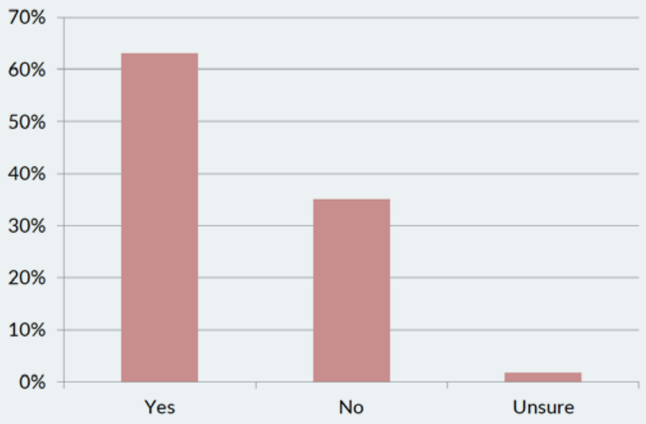

Source: Forsyth Barr analysis

Figure 41. How much did your organisation spend on AI in your last financial year?

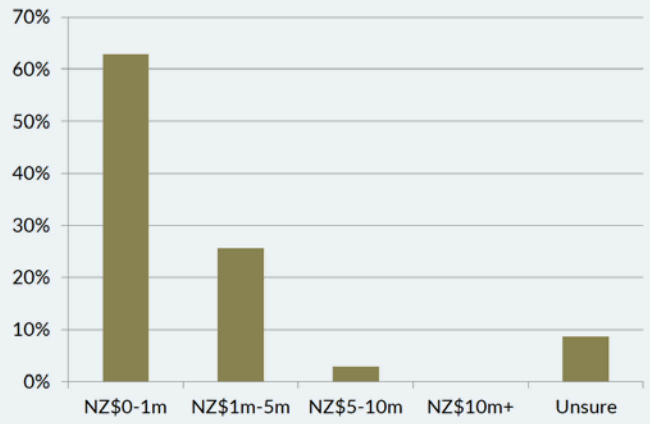

Source: Forsyth Barr analysis

Figure 42. How many full time equivalents are dedicated to AI within your organisation?

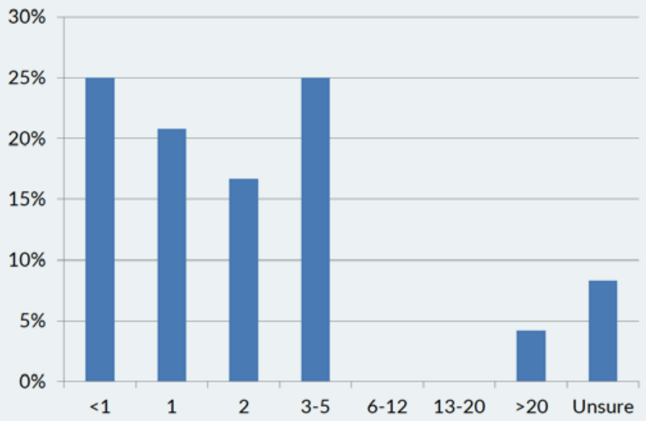

Source: Forsyth Barr analysis

Figure 43. How do you think AI will impact your industry over the next 3–5 years?

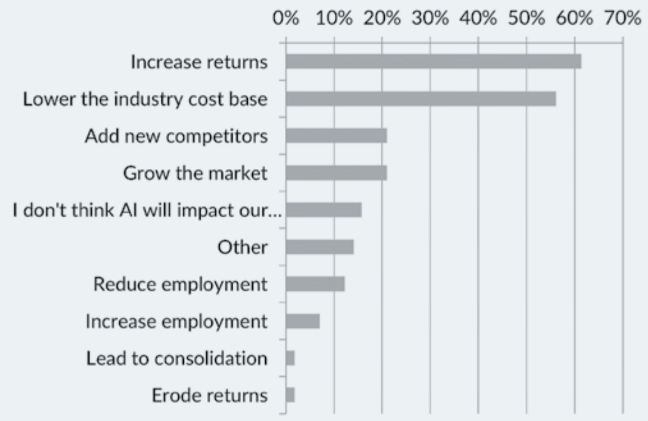

Source: Forsyth Barr analysis

#2 A positive experience of AI so far and an expectation of more to come

While investments so far are modest, the majority of respondents have invested at least something in AI, and the experience so far is overwhelmingly positive. Close to two thirds had had an overall positive experience and not a single one stated that AI so far had had a net negative impact on their business. Looking ahead, the respondents are even more positive, with the majority of respondents expecting AI to drive both increased returns and lowered cost base while only one company stated that it expects AI to erode returns. We view it as encouraging the corporates see AI as a "force for good". However, we caution that it also points to many corporates underestimating the risks of disruption. Only a very small sub sample of NZ companies could be considered leaders within AI. There is a risk is that a global digitally native company disrupts domestic companies.

#3 Plain vanilla AI applications within SG&A dominates; but some are innovative and critical

We asked the corporates that used AI what applications they had adopted, and we were surprised and encouraged by the breadth of the answers. Although (likely off the shelf) applications within areas such as sales, marketing and customer services dominated, "hard" applications featured more prominently than we had anticipated. Several companies listed areas such as supply chain management, product design and risk as areas where they had applied AI, indicating a relatively advanced use of the tools available. Slightly less encouraging was that only very few companies declared that senior management regularly use AI as a support for decision making. This is worth noting given AI's potential to improve the speed and quality of evidence-based decision making. However, as the event velocity of the business world increases exponentially, their feelings about the optionality of computer assistance may change significantly. Consider that many corporates within healthcare and finance simply could not do their jobs without assistance from AI.

#4 Lack of AI talent is the biggest problem encountered

When asked what problems their organisation encountered developing or utilising AI, only one in four stated that they had not encountered any material problems. The most common response by far was delays and difficulties with recruiting AI talent. A broad lack of skilled software engineers has been named as an impediment for growth, for instance by Plexure, so this is not that surprising. However, what we found more surprising was the limited experience of inflated project costs and also that the least common answer was lack of board support. The latter may be a consequence of the very senior nature of our respondents.

#5 Significant risks associated with AI may be underestimated

Overall, our survey results suggest a positive, almost naive approach to AI. It is expected to deliver higher revenues, lower costs and the risk of disruptive competition is seen as low. At the same time there appears to be a lack of urgency regarding capturing these opportunities amongst most respondents. Maybe the most significant indication that NZ corporates are still early in the discovery process concerning AI's risks and opportunities is a lack of awareness around the risks of misuse of AI.

Only around 10% of respondents are concerned about data bias in their organisation. We take the view that even very sophisticated users of AI would benefit from being somewhat concerned about bias, or at the very least be unsure. Furthermore, several companies suggested that bias in the data was "inapplicable"; this may indicate incomplete knowledge of the problem. Data bias is not the only risk. Companies need to consider testing and ethical guidelines for reducing the risk of AI having unanticipated consequences, and the need to adapt human processes so AIs tend to augment rather than replace humans at work.

Figure 44. Has the introduction of AI had an overall positive or negative impact on your business?

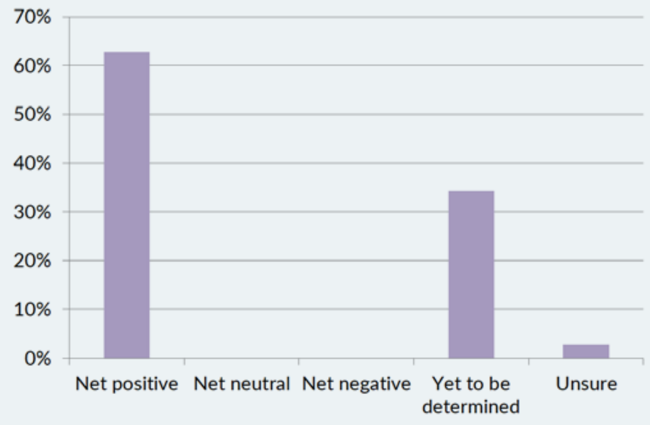

Source: Forsyth Barr analysis

Figure 45. How else has the introduction of AI had an impact on your business?

Source: Forsyth Barr analysis

Figure 46. What application categories is your organisation likely to apply AI to over the next two years?

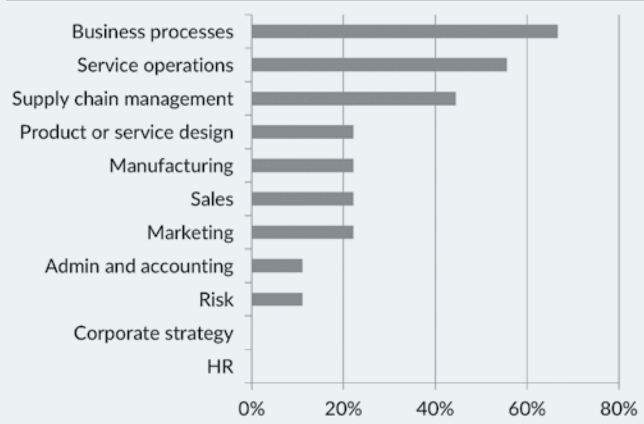

Source: Forsyth Barr analysis

Figure 47. Does your senior management use AI as a support for decision making?

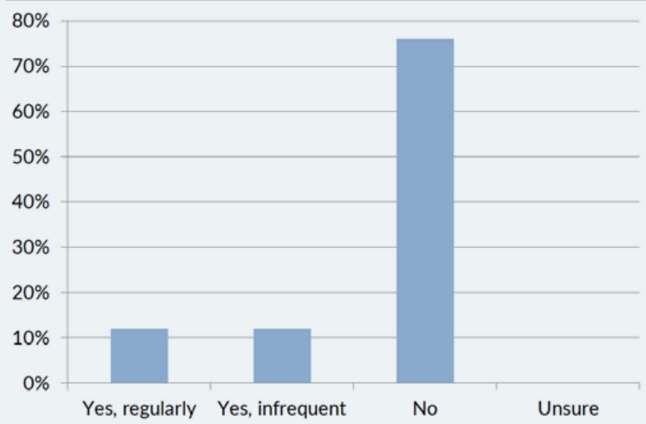

Source: Forsyth Barr analysis

Figure 48. What problems has your organisation encountered developing or utilising AI?

Source: Forsyth Barr analysis

Figure 49. Are you concerned about bias in your organisation's machine learning data?

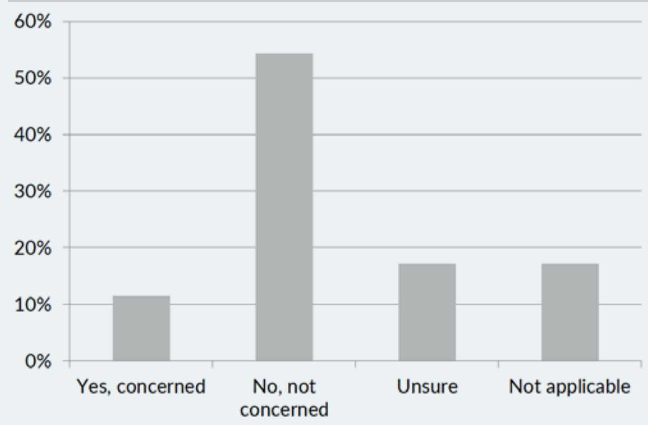

Source: Forsyth Barr analysis

Opportunities and economic risks across New Zealand corporates

AI provides both a huge opportunity and a meaningful risk to New Zealand Inc, in our view. NZ has a lot of long term structural advantages for capitalising on the potential of AI. It is an attractive place to live for global AI talent, it has connectivity with Europe, US and Asia, it is English speaking, and it already has numerous innovative small tech companies. NZ is a small country with a reputation for minimum start up time and costs. It has a well-developed tech infrastructure in the form of fibre to the premises and soon 5G will be rolled out across most of the country. It lends itself to being a test bed for innovation.

The ability to take advantage of AI ranges from the potential to be totally transformative in areas such as logistics and healthcare to more modest improvements in efficiency and customer care in sectors such as utilities. In sectors such as REITs and agriculture, the impact may not be obvious at first, but we believe it could be surprisingly transformational longer term. What will the optimal built environment look like if you don't need your own car to get around? What if you can watch a movie, sleep, or work in a hired car-pod that takes you door to door? How will agriculture change when AI and robotics managed vertical farms deliver between 240 and 400 times more

crop yield per metre (depending on crop and methods) while using ~99% less land vs traditional agriculture? Vertical farming may also use ~95% less water and zero pesticides. These factors vary with local conditions, but significant gains are possible.

We believe many New Zealand companies are well positioned to take advantage of AI. The exponential increase in value of data within healthcare should benefit Infratil's (IFT) diagnostic imaging business, Pacific Edge's (PEB) early detection tests, and Fisher & Paykel's (FPH) ability to manage utilisation of its products. There are massive opportunities for Mainfreight (MFT) and Freightways (FRE) to take advantage of progress made towards AI powered autonomous vehicles. A company like Arvida (ARV) could see its homecare offering turn into a significant contributor to earnings. Spark (SPK) has built out a strong offering under its IT services umbrella, supporting New Zealand businesses to take advantage of AI, and "digitally native" companies like Plexure (PX1) and Serko (SKO) are starting to build moats around their AI capabilities.

The risks to New Zealand's ability to take advantage of AI, however, are also significant. The most prominent, in our view, are a lack of urgency and a lack of local AI talent. A common thread through our proprietary survey is that companies appear to see potential in AI, but often do not have any imminent or urgent plans to do anything about it. Additionally, the few who do try to build out their AI capacity are, unfortunately, facing an acute shortage of qualified people to hire. An additional risk facing some of NZ inc is the constant mission creep of the US$1trillion club of US tech companies. Who knows where they will go next? Communication, retail, transport, healthcare are all areas where mega tech is already active today. Any company that rest on its laurels is highly vulnerable.

Next, we present an analysis of specific NZ industry sectors and their use or potential use of AI, followed by a brief analysis of selected NZ large cap, small cap, and unlisted companies.

Transport

Among our transport listed companies under coverage, 50% are currently engaged with AI applications, of which two thirds have been employed beyond a pilot stage. However, all the transport companies surveyed see AI as a means of enhancing business operations in the future. More than half also see AI as a means of reducing costs, enhancing customer experience, and improving decision-making.

To date, AI has had only a minor beneficial impact on covered transport companies that have employed applications beyond a pilot stage with only modest levels of investment. However, we expect AI to increasingly add value within the sector. For example, Air New Zealand (AIR) has recently employed FLYR labs' Cirrus Revenue Operating System to operate its revenue management function. FLYR leverages AI that captures and analyses massive amounts of data, enabling AIR to more accurately forecast demand and automatically set revenue-optimal fares. It has delivered an overall +7% increase in revenue in trials with airlines while reducing forecast errors by 10x as compared to legacy vendors.

Many of the AI applications that will transform the transport sector (by lowering costs, reducing congestion, improving safety, and enhancing customer service and engagement), have been in development for some time. They are often not new ideas. For example, despite significant investment in autonomous vehicles since 2010 there are still no fully self-driving cars available for consumers.

Autonomous vehicles: An autonomous vehicle can sense its environment and operate to achieve its objectives with little or no human input. Key applications of this technology include personal self-driving vehicles, shared taxis, industrial fleets, delivery trucks, and platooning (driving a group of vehicles together). The technology is currently being tested globally by a large number of entities, but its implementation will likely be subject to regulation given the safety and accident liability issues at stake. The public transport and freight industry will also be key beneficiaries of autonomous vehicles given the forecast significant reduction

Figure 50. Autonomous trucking funding has climbed...

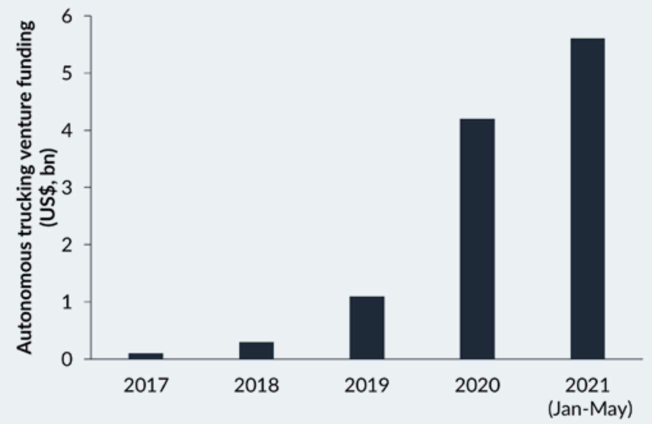

Source: Pitchbook, Forsyth Barr analysis

Figure 51. ...so too has investment in drone technology

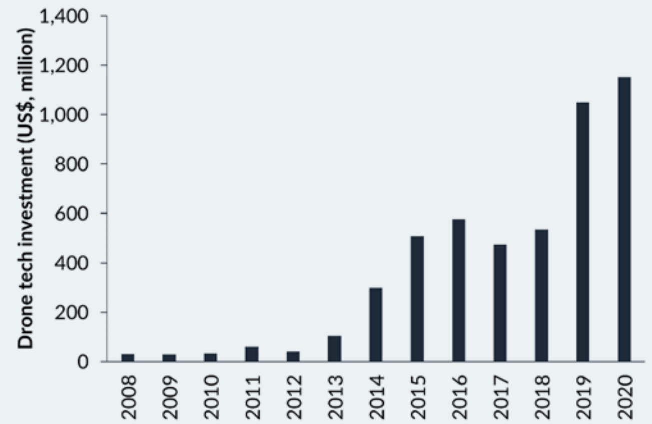

Source: DRONEII, Forsyth Barr analysis

Many of the AI applications that will transform the transport sector (by lowering costs, reducing congestion, improving safety, and

Figure 52. Leading use cases for AI in retail and consumer goods

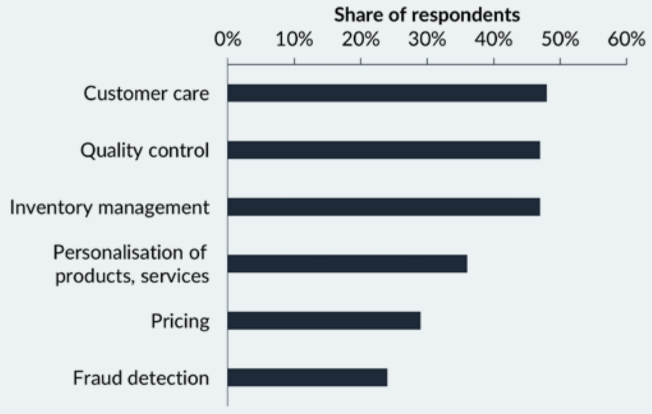

Source: MIT Technology Review Insights survey (2020), Forsyth Barr analysis

Figure 53. Organisations are already realising significant benefits from chatbots

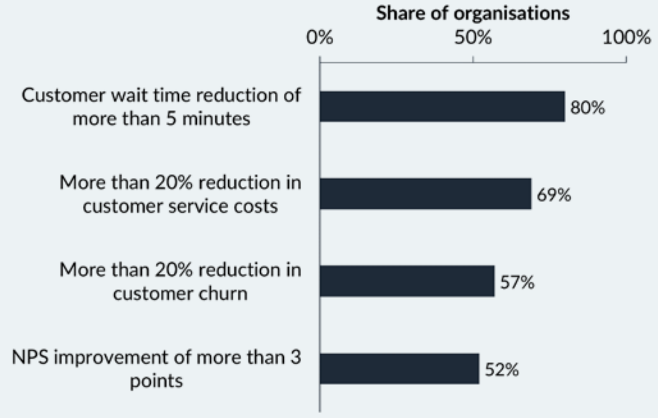

Source: Capgemini (2019), Forsyth Barr analysis

Figure 54. AI can increase the economic value of wind energy

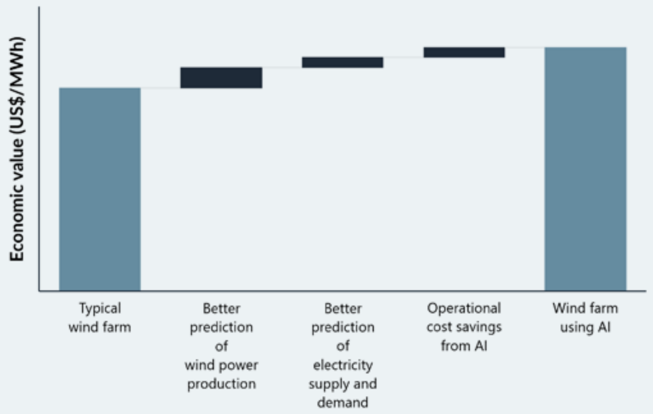

Source: DeepMind, Forsyth Barr analysis

Figure 55. Expected business efficiency impact of AI today and in the near future

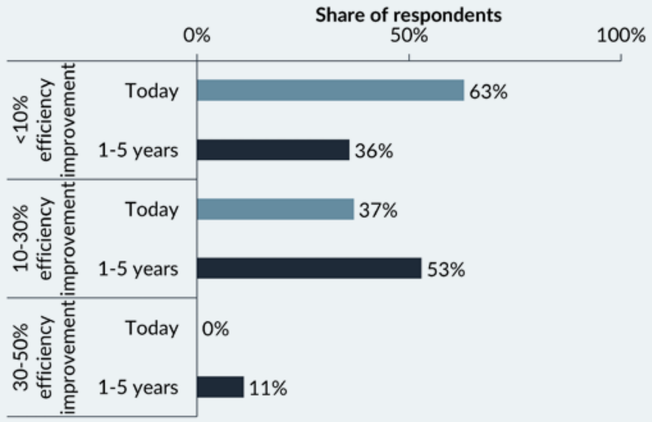

Source: Roland Berger, Forsyth Barr analysis in labour and maintenance. This is likely to benefit the industry more than individual players and could reduce barriers to entry significantly.

Drones: A drone taxi, or UAV (unmanned autonomous vehicle) will transport passengers short distances at high speeds bypassing congestion and the need for high cost land based infrastructure. Smaller drones already transport time-sensitive packages, but not yet on a mass scale given regulatory hurdles. We expect drones to transform the parcel industry and lower the last mile cost for the B2C channel.

Warehouse automation: Many AI applications can be applied in warehousing facilities to reduce labour costs, increase safety, and improve accuracy. These include automated picking robots, package shipping, product tracking, voice assistants for workers, and augmented reality wearables.

Airline delay prediction and avoidance: The air transport industry is synonymous with delays. Passengers are usually the last people to hear about delays. Delay prediction and avoidance based on weather, maintenance problems, or technical glitches could reduce passenger airport dwell times (potentially negatively impacting airport retail income) and improve customer experience.

Traffic management: Predictive route planning with real time changes could reduce congestion for both commuters and freight companies and improve safety. Sensors and cameras deployed along roads collect the data on traffic, which AI then uses for traffic management and route planning purposes.

Predictive route planning could reduce congestion and improve safety.

Retail

Heightened consumer expectations and the blending of physical and digital purchasing channels by retailers mean that AI already offers competitive advantages and will increasingly be a key competitive edge and business tool for most major retailers.

Among our retail listed companies under coverage two thirds have invested in AI applications. Of these virtually all have been adopted beyond a pilot stage. Retailers appear more concentrated on top-line AI opportunities with the most applications addressing marketing (i.e. personalisation through better understanding of customer behaviour), sales (i.e. pricing optimisation) and customer service. Ultimately, we expect AI to shift retail more towards the scientific end of the art-science continuum.

The retailers leading the charge are those with more developed ecommerce strategies, yet other applications where AI is being deployed currently include inventory management, invoice processing, and supplier reconciliations.

Over next few years, AI applications will impact the sector in the following ways:

Pricing flexibility: AI can help to optimise retailers pricing strategies. First, AI can automatically adjust prices by seasonal trends, competitive products, and consumer demand (for example in the US, Kroger Edge technology has introduced smart shelf tags instead of paper price tags). Second, AI can analyse the outcome of multiple pricing strategies based on historic and real time data to set the optimal price for a particular product.

Enhancing in-store experience: AI can automate in-store operations and reduce operational expenses in retail stores. It can replace sales staff to assist customers in the store, reduce queues through cashier free shopping (for example Amazon's Amazon Go and Just Walk Out shopping technology, which responds every time a customer picks up a product from the store shelf), automatically replenishes stock by real-time stock monitoring, enhances merchandising by understanding how customers navigate stores, digitises store displays, and virtualises fitting rooms.

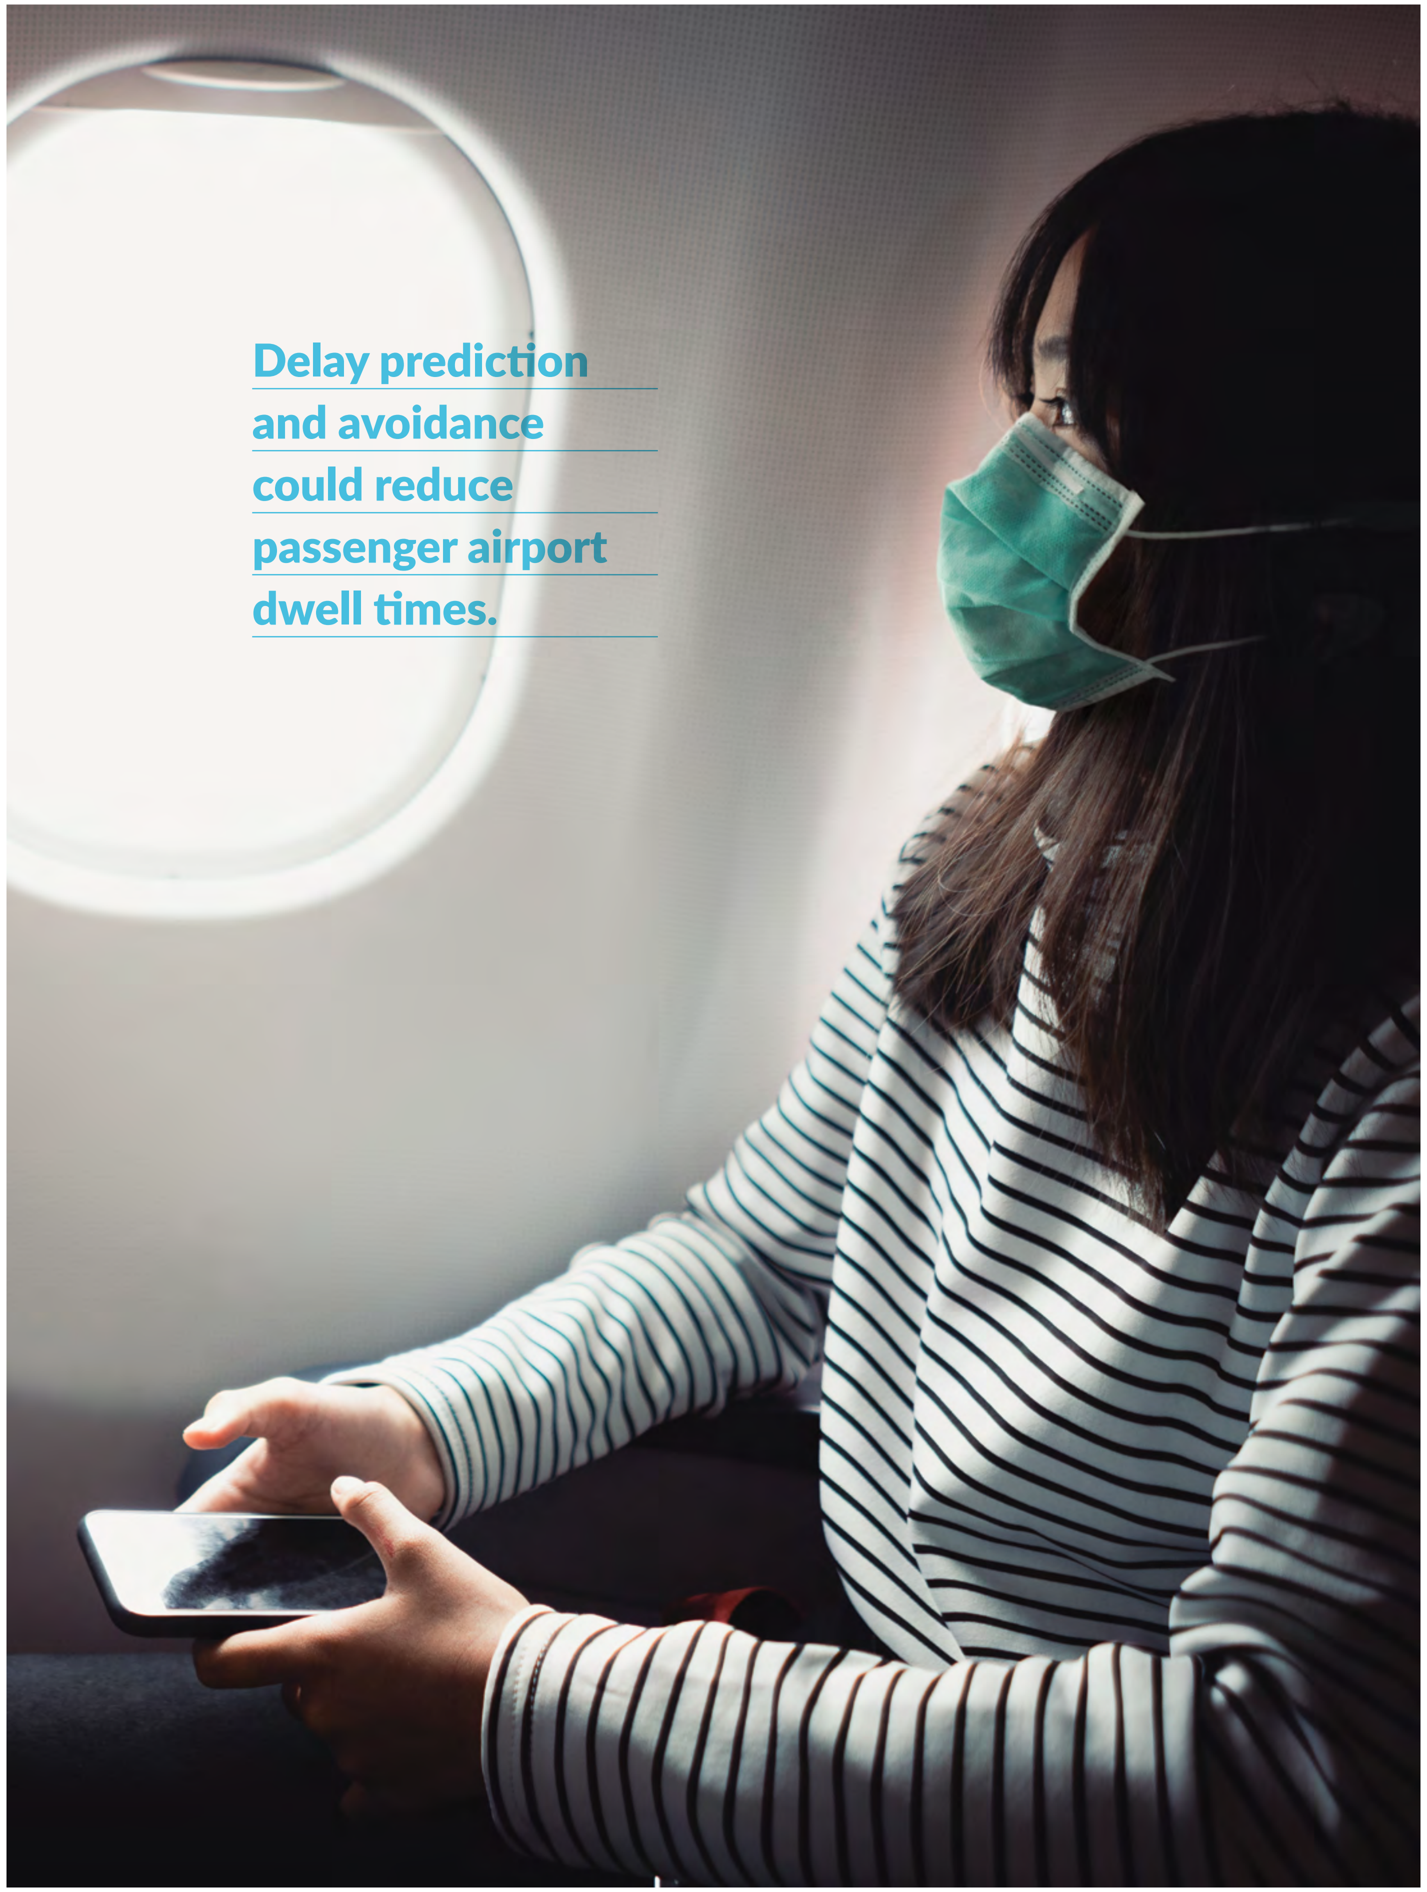

Chatbots assist with customer service: Omnichannel retail extends the reach of customer service beyond the retail store. AI can dramatically enhance customer engagement by powering interactive chat platforms. There are still many limitations of natural language understanding, problem interpretation, and solution generation. Chatbots can currently answer frequently asked customer questions, recommend products, address grievances, and collect valuable data before diverting the call to a human, if needed.

Inventory management: AI can execute real time purchasing reflecting fluid demand forecasting. These systems can take into account a history of sales, location, weather, trends, promotions, and other parameters. This can enhance product availability and working capital. For example, retailer Morrisons has improved stock forecasting and replenishment in 491 stores with the help of BlueYonder. This resulted in a ~30% reduction of in-store shelf gaps.

Retailer Morrisons improved stock forecasting in 491 stores resulting in a 30% reduction of in-store shelf gaps.

Utilities

AI has the potential to cut energy waste, lower energy costs, and facilitate and accelerate the use of renewable energy sources globally. AI can also improve the planning, operation, and control of power systems.

AI is forecast to address cost issues initially in the utilities sector, with all respondents to our survey expecting it to lower the industry's cost base and to enhance business operations and processes in the future. However, progress to-date has been slow with only half of the sector having deployed AI beyond a pilot stage. These AI applications were primarily applied to customer service and sales. Of the other organisations, only one is not expecting to invest in AI over the next two years.

AI applications that will impact the utilities sector over next few years include:

Predictive maintenance: Electricity generators can use AI to move beyond prevention-only maintenance to include predictive maintenance and fault predictions through equipment monitoring and better data analysis. AI paired with sensors can be used to monitor equipment and identify current operational problems that might have gone unnoticed by humans acting alone.

Load forecasting and optimisation: AI is especially critical for the energy industry to accurately forecast and predict in real-time how much power needs to be produced. More accurate power predictions result in lower production costs, reduced energy waste, retained revenue and more stable wholesale prices.

Optimising power output: AI has more application in renewable vs. traditional fossil fuel power generation operations given the variability of renewable solar and wind supply. For example, AI has been able to increase the value of wind energy generated by Google's wind farms by 20%. It has made wind farms more viable by using its DeepMind machine learning algorithms to predict wind energy output. The company can now schedule deliveries of energy output, which are more valuable to the grid than standard, nontime-based deliveries.

Electricity trading: AI energy trading platforms analyse thousands of variables to forecast market prices and anticipate grid conditions that are likely to influence prices. They determine optimal economic bids, which tell the power generation company the prices at which it is willing to continue generating.

Customer engagement: AI monitoring of customer energy usage can allow better tailored and more personalised services for customers, providing a key customer retention tool.

Food, beverages and agriculture

AI has the potential to enhance New Zealand's competitive advantage in growing and producing food and beverages. Yet, the listed agricultural sector has been slow in its uptake of AI. To date, only two of seven companies surveyed adopted AI applications beyond a pilot stage.

All agricultural companies surveyed see AI as a means of enhancing business operations and processes. Most see AI also as a means of enhancing the customer experience. Only one company has in-house AI capability, which was a single staff member.

Globally, AI has already penetrated many parts of food and beverages production. From supply chain management, waste management, predictive maintenance, scheduled ordering, weather predictions, and food safety compliance, to new product development, AI is already reshaping the sector.

In New Zealand, Fonterra is already a leading user and facilitator of AI in a variety of different areas. For example, it is using AI to detect improperly sealed or faulty bags of powdered milk in its factories. In addition, it has recently agreed to stock an AI based farming tool from Dutch company Connecterra in its Farm Source stores.

These key AI applications will impact the sector significantly over next five to ten years:

Smart farming: A variety of AI applications enhance crop yield and reduce growing and production costs, including crop yield predictions, soil monitoring, and robocroping (the use of harvesting robots).

Sorting fresh produce: Historically, food processors hired people to undertake the monotonous and routine identification and sorting actions linked to food selection. AI replaces these humans by easily recognising and sorting apples, grapes, honey, fish, mussels, and other food products.

Food safety: AI enhanced sensor and camera technology are already starting to be deployed and have the potential to significantly improve food safety and extend the shelf lives of products.

Minimising food waste: UN data suggests ~14% of food is lost between harvest and retail, with further losses occurring thereafter. Food service industry data suggests that 5%–15% of food purchased ends up in the bin. Cameras with AI algorithms can recognise foods that typically end up in the waste bin, allowing better understanding of food waste patterns, This improves food utilisation and enhanced sustainability.

Cleaning processing equipment: Processing equipment in the food manufacturing and food service industries requires time consuming and highly labour intensive cleaning operations. AI can reduce cleaning costs by better monitoring food residue and microbial debris within equipment and providing intelligent controls for cleaning robots and automation.

Fonterra uses AI to detect improperly sealed or faulty bags of powdered milk in its factories.

REITs

Among our REITs under coverage, the majority responded that they do not currently have an AI vision or strategy. Only about one third have invested in any AI application, and the ones that had invested in AI had done so on a very modest scale. Despite these modest investments and low uptake of AI, the majority of REITs expected AI to increase industry returns in the future and nearly half expected AI to reduce costs. Why the low uptake? Comments from the ones that had started to adopt AI suggest a lack of understanding of the opportunity as well as the potential for added complexity were problems that they had encountered.

Key AI applications that will impact the REITs sector over next few years

The REITs sector may face less dramatic immediate impact from AI. However in the long term we see considerable room to improve efficiencies and gain a cost advantage for those that adopt AI intelligently. Smart Cities and Smart Buildings have been discussed as a major investment theme for over a decade and they have so far failed to deliver the promised returns - but it is still early. However, several trends are slowly coming together which could reverberate through the sector. Specifically, distributed electricity generation, smart appliances, high bandwidth fibre to the premises, and the Internet of Things (IOT) have laid the hardware foundation for AI to enter the built environment. While the impact is likely to be long term, we believe tenant management and leasing, building systems, optimising of transactions and construction monitoring are areas where an astute operator can use AI to gain a meaningful competitive advantage.

Building systems: Already today there are highly advanced AI driven building systems optimising HVAC, lighting, and people movement within buildings. Operational and energy costs of a Smart Building compared to a less smart one is material, yet the vast majority of buildings in NZ are still less smart. A relatively simple upgrade of elevator software, for example, can increase capacity by 15–20%, effectively adding a whole elevator shaft in an averagely sized office tower.

Optimised transactions: An under appreciated area where we believe AI can impact the REITs relates to optimising transactions. Today, the vast majority of tenders focus almost entirely on price. As demonstrated by Sandholm's team, as discussed earlier, there are many variables in addition to price that are important to both the buyer and the seller. This includes factors such as time of settlement, financing status, delivery guarantees, contingent liabilities, and granular product features. By applying AI to optimise tenders the value for both buyers and sellers can improve.

Construction monitoring and optimising: Most REITs in NZ are also significant property developers. AI mature REITs have the potential to incorporate the potential of AI in the design phase of the building to be 'smart city ready'. Additionally, developers already use AI to improve progress monitoring and inspection, optimise workflow, and reduce complexity and downtime.

AI powered transport solutions. Self-driving, smart service delivery and smart sharing have the potential to fundamentally alter how space is utilised in high density environments.

For REITs that adopt AI intelligently, we see room for improved efficiencies and cost advantages.

Figure 56. AI and connectivity expected to unlock significant value in agriculture

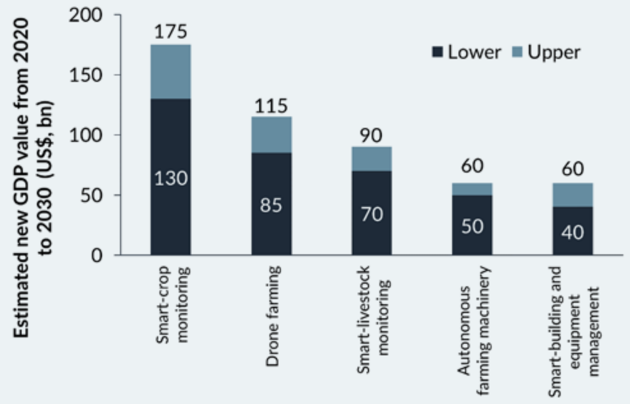

Source: McKinsey, Forsyth Barr analysis

Figure 57. Growing concerns about food waste supports productivity saving investment

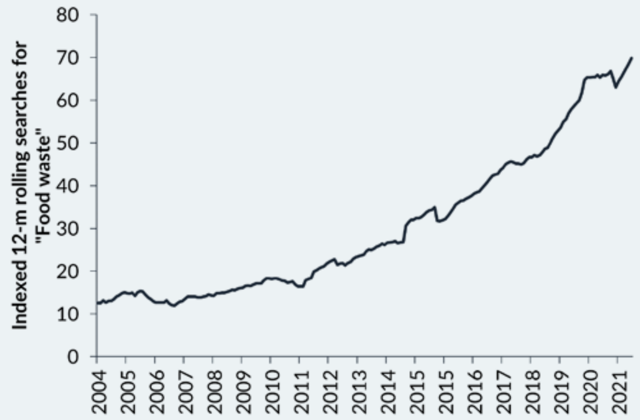

Source: Google trends, Forsyth Barr analysis

Figure 58. Energy savings from smart technologies

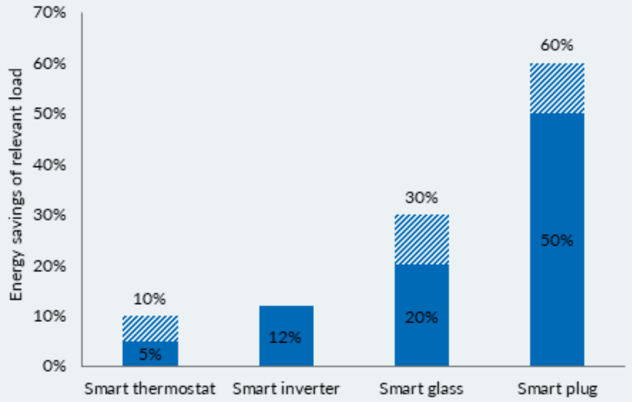

Source: Forsyth Barr analysis, ACEEE

Figure 59. IoT connections in Europe smart buildings

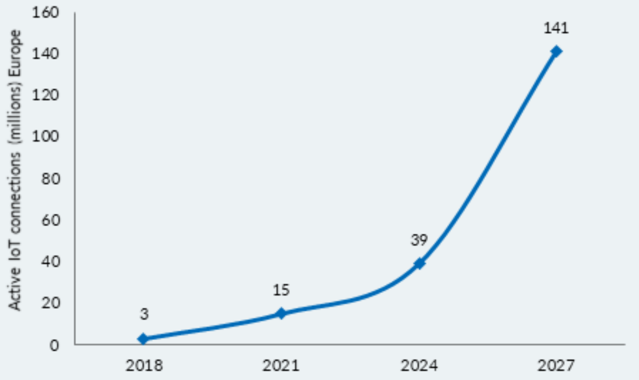

Source: Forsyth Barr analysis, ETNO

Figure 60. Cost of genome sequencing

Source: Forsyth Barr analysis, NIH

Figure 61. Value of wearable data

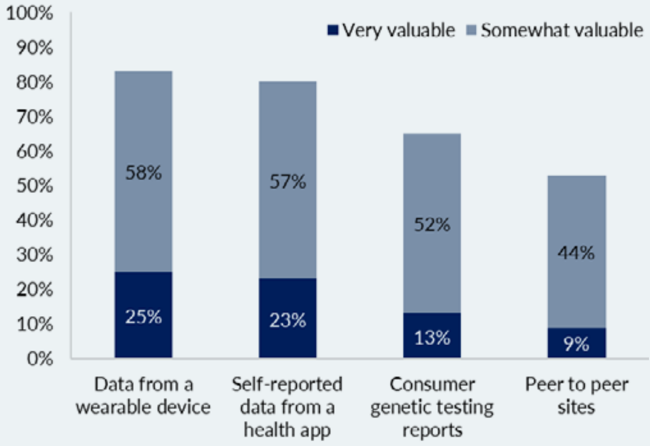

Source: Forsyth Barr analysis, Stanford Medicine Report

Healthcare

Among our healthcare listed companies under coverage, the majority that responded are currently engaged with AI applications and half are beyond a pilot stage. AI is primarily seen as a means to improve medical decision making and enhance the customer experience but also as a means to increase revenues, reduce costs, and enhance processes.

While the engagement with AI appears relatively high and the impact is expected to be positive within numerous areas, the actual impact to date is yet to be determined. We expect the healthcare sector to be fully transformed by AI. It has the potential to impact the core health delivery business as well as business administrative processes and customer experience.

Three broad sector trends that we believe will determine the long term winners and losers within the healthcare sector:

- The hunger for health data will increase in importance and value: Gathering and categorising data will be a key point of differentiation. We see Pacific Edge (PEB) as well positioned for this trend as well as Infratil's (IFT) Qscan and Pacific Radiology. Improving the quality of data (and access to that data) should also give Fisher and Paykel (FPH) a better gauge of utilisation and could support it with cost optimisation of its sales force and potential TAM expansion.
- Pre-emptive and predictive care vs treatments of symptoms: We expect healthcare spending in predictive care to increase as a share of total healthcare spending. We see both Infratil's diagnostic imaging and PEB as well positioned to take advantage of this trend. Increased early diagnostics could also support growth for FPH's homecare division.
- Software over hardware: The (so far) modest investments by NZ's healthcare companies into AI according to our survey puts them at risk of ending up down the value chain. Digitally native and well-resourced companies could take an ever increasing share of the pie. Google is investing heavily in Google health, to give one emblematic example.

Key AI applications that will impact the healthcare sector

The healthcare sector is one of the sectors outside of tech per se where AI will have dramatic long term implications and where commercialisation is already taking place. Advancements in monitoring, diagnostics, smart therapy planning, and the exponential increase in personalised health data will transition the healthcare industry from sick care to precision medicine and real heath care. Three trends which we believe will have major impacts on the healthcare sector over the next few years are:

DNA analytics and predictive care: It took 13 years and an estimated US$300m to complete the draft sequencing of the 3.2bn base pairs of the human genome, and total costs to complete the project has been estimated at well over US$1bn in today's money. The project was finished two years ahead of time in 2003. Today, it costs less than US$600 and takes less than a day to sequence human DNA. We are still only scraping the surface of what this knowledge and technology can achieve, but it is expected that with the assistance of AI, it will influence most areas of healthcare. This is particularly so when it comes to predictive and pre-emptive care.

Diagnosis and imaging: Among the first applications of AI in healthcare to be approved by the FDA in the US were the use of AI to analyse medical imaging such as Magnetic Resonance Imaging (MRI) and Positron Emission Tomography (PET). The list of illnesses where AI plays an important part in assisting physicians to diagnose and treat is constantly growing. Five years ago, AI powered algorithms were shown to improve accuracy of breast cancer diagnostics, saving time for physicians as well as reducing unnecessary surgical interventions. When IFT bought Qscan and Pacific radiology, the synergies and scales benefits from AI was given as one of the key drivers of the long term economics of the deal.

Well-being monitoring and the democratisation of health care and diagnostics. Medical devices for personal use already includes ECG monitoring, oxygen levels in babies and infants, temperature and heart rate monitoring, skin temperature as an indication of inflammation and ovulation and much more. AI, in combination with the advancement of wearable sensor kits, has the potential to transform many aspects of healthcare. Individuals increasingly have access to diagnostic tools that rival those used by physicians today.

Aged care

Among our aged care listed companies under coverage all that responded are currently engaged with AI applications. As in the healthcare sector, as a whole, AI is seen as a means to improve decision making and enhance the customer experience as well as increase revenues, improve decision making and enhance business operations.

We expect AI to increasingly add value within the sector. The large listed aged care operators that can afford to invest for the long term will increasingly get an advantage over smaller private operators.

Key AI applications that will impact aged care sector

The aged care sector may at first glance appear to be largely unaffected by AI, but long term we see several vectors through which the sector will be impacted. While the impact may be subtle at first we believe AI driven operational efficiencies, actuarial risk assessment AI assisted home care and personalised wellbeing can impact the sector. The biggest threat as well as opportunity relates to the ability to deliver personalised care remotely into the home. Another potential source of uncertainty relates to AI enhanced diagnosis and drug discovery in the treatment and prevention of dementia.

Patient surveillance and safety: The fastest area of growth within aged care is dementia. All the listed aged care operators now include dementia care in new villages. We expect the demand for dementia care in NZ to increase significantly over the next decade. Caring for dementia patients is very expensive and the built environment needs to be significantly altered relative to traditional care to allow for the "escape" risk. AI has the potential to change this meaningfully in several ways. One such example would be facial recognition technology allowing for dementia patients to move more freely within the age care villages. Longer term, AI may aid in drug discovery to slow or effectively treat dementia in its early stages.

Patient monitoring and tracking: There are many AI assisted patient monitoring and tracking systems in active use today, including systems from HeadsUp, Honeywell, Vocera, Oura, and others. AI has the potential to completely transform the way health care is delivered. Specifically, we believe that an increased focus on pre-emptive care and wellness over just delivering healthcare to sick people has the potential to reduce the cost of care while, at the same time, increasing the wellbeing of the aged care residents.

Home care vs. retirement village living: AI has the potential to extend the age of vital health (healthspan), and to dramatically improve the ability to care for older people in their own homes. Wearable monitoring devices and remote diagnostics in combination with alert systems and virtual home visits could result in people being able to stay healthy at home substantially longer. This has the potential to be a threat to the aged care operators. However, if they operate proactively, we think the opposite is more likely to be true. The number one reason why people move into a retirement village is for social life and to be close to a spouse that may already or soon need care. Intelligently applied, future AI systems have the potential to create safer, health monitored environments, and a sales funnel for innovative new services with lower total costs for aged care operators.

Technology — enablers and adopters

The NZ listed technology sector consists of an eclectic mix of primarily small cap stocks. The potential for AI to transform and disrupt within the technology sector is by default orders of magnitude higher than within more traditional industries. The likely impact will be all-encompassing. Some of NZ's tech companies can be classed as early adopters (such as PX1), while others are yet to determine a fully comprehensive AI vision and strategy. Some of the AI adoption hesitancy can be explained by concerns about bias in machine learning data, potential for applications to be hacked or have unanticipated consequences, limited team experience on heterogeneous projects, and a general unfamiliarity by many business leaders about the competitive advantages AI can enable for their organisation. All of the companies surveyed below have stated that recruiting AI talent is the current major hurdle to development and deployment of AI applications.

Plexure (PX1): Within the wider marketing sector, PX1 provides a number of personalised mobile marketing services to the QSR and grocery operations leveraging its machine learning and AI. An early adopter of AI, the company's objective remains to drive commercial uplift for customers by forecasting consumer behavioural patterns. This allows PX1 customers to better process large datasets of accumulated information and subsequently allow marketers to provide insightful and personalised marketing content to their customers. PX1 business customers have reported that using PX1's AI and machine learning services has resulted in an increase in the average basket size. PX1's success in the mobile marketing and AI sector has been driven by the global shift within marketing away from traditional print and towards personalisation and digital channels.

For PX1, this has led to increased revenues and the ability to enhance its customers experience. The company continues to build out its AI functionality, adding additional features and offer new modules and add-ons. PX1 aims to provide tools for companies to gather and consolidate customer data. This includes recent transactions, in-app behaviours, demographics, and offline purchases connected through physical loyalty cards. This kind of data allows PX1 companies to make estimates and predictions around future behaviours while improving customer experience on an individual level. This in turn drives greater commitment from PX1's customers, boosts customer satisfaction and retention and grows the customer brand.

Pushpay (PPH): PPH is yet to develop a fully-fledged AI strategy, but we believe an opportunity to develop AI functionality in the church technology market exists. For a number of years, data collection within churches has grown in both quantity and significance. With megachurches in particular using attendance metrics as internal KPIs regarding the 'health' of a church, technology has been deployed to collect information around growth, age, donation history and other key demographic information. Up to now, church management systems (ChMS) have collected data (such as PPH's Church Community Builder) and provided pastors with updated metrics around current attendance and donations vs. history.

This has been met with popular demand as churches look to create a more personal feel with congregation members and also to follow up with those members needing assistance, prayer, emotional support, or counselling for financial distress. For church fundraising projects, the use of analytics data has also been a key driver, with pledges drawn up by donors throughout the year and pastors alerted to months where donations have lagged behind budget.

Content generation is also a key AI application opportunity. Pastors are competing for congregation members attention on an ongoing basis. Analytics about what data and messages are well received allows churches to communicate more effectively. This aligns with the goal of moving the church's mission forward, supporting people for their relevant needs, bringing more people into the community, and enabling members to participate more actively. Advances in technology are nothing new for the church. The first printing press technology enabled the ability to print the bible and was deemed by some as the catalyst to Europe's reformation.

Evangelists such as Nicky Gumble and Billy Graham were famed for radio communication over the past century. Hillsong and other modern megachurches have utilised advancements in television and media to spread their messages internationally. Today's Internet and smartphone apps have shown massive adoption in church technology, with over 5,000 apps in the app store related to 'church'. As a result, we see rapid adoption of AI in the church technology market, but it has not yet begun to reach its potential.

Serko (SKO): SKO uses AI within its flagship product Zeno in order to predict the different elements of a customer's journey from home to hotel via airport and hotel. Its AI powered recommendations use previous behaviours to push suggested locations, times and type of accommodation in an effort to provide the 'connected trip' for its customers. As one of the leading providers of AI within the booking and travel space, this was strongly validated as the company won a substantial contract to provide the backend AI and data functionality for Booking.com's Booking for Business platform. With SKO employed to provide its connected trip offering at scale, it will look to leverage its AI capabilities in order to grow future total transacted volumes (TTV) for Booking.com.

Going forward, we expect SKO to continue to innovate its leading AI technology for use within the travel and booking sector. We believe a range of untapped opportunities exist around the connected trip — both in the corporate and leisure booking sectors. SKO is well positioned to leverage its current market position and capabilities.

Spark (SPK): SPK has spent the last decade accumulating substantial AI technology know how under its fast growing IT-Services umbrella. If these capabilities were separately listed they would probably be NZ's largest technology company. Maybe one could describe Spark's technology ambitions as bringing the best of global technology to NZ inc. We believe this is a realistic and admirable ambition. Spark may not be able to compete with the largest global tech companies in the AI arena. However, it does have the potential to extract meaningful value by being the "last mile" partner for the likes of Amazon Web Services and Microsoft — in particular to small and medium sized companies that may not be sufficiently profitable for the trillion dollar club to target. Spark's IT services division now operates numerous smaller technology companies across the tech value chain from traditional areas such data. Centers, procurement and connectivity, to advanced cloud businesses, and AI consultancy. Spark's significant inroads into various areas within advanced technology including AI has the added benefit of educating the organisation in general and the leadership in particular about the potential in adopting leading technologies.

The most AI mature of Spark's tech companies is probably the fully owned Qrious which is one of the five companies that form the "Spark Business Group". Qrious is an AI powered "big data" analytics company with a mission statement to "unleash data-driven innovation that revolutionises the way organisations do business". Qrious works directly with AI, and one of its offerings is Enterprise AI which offers AI powered risk projections, customer centre enhancement, and data entry optimisation. Qrious collaborates with several of the global leaders with AI capabilities at scale such as Microsoft, Amazon Web Services, and Cloudera. Qrious also has multiple exclusive re-seller agreements with some of the world leaders within AI.

Listed aged care operators investing in AI will increasingly get an advantage over smaller private operators.

Figure 62. Dementia prevalence in NZ

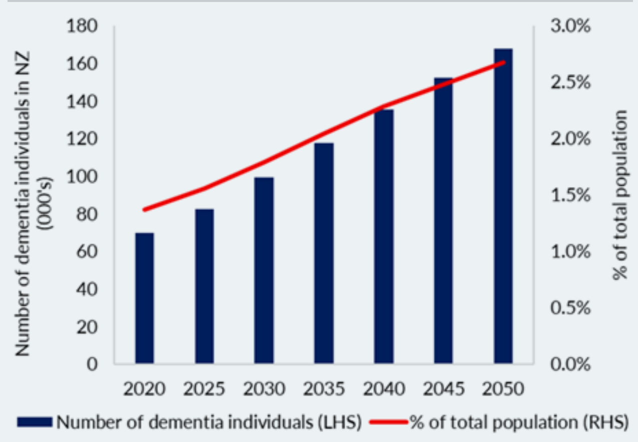

Source: Forsyth Barr analysis, Alzeihmers NZ

Figure 63. Australia homecare revenues

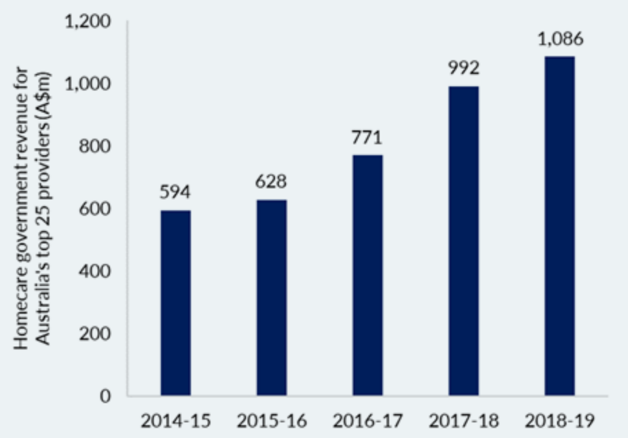

Source: Forsyth Barr analysis, KPMG

Appendix 1:
AI company investments and applications

Investing in AI company applications is unlike investing in most business opportunities. There is often more risk and more upside opportunity. Beyond the typical technology company start-up risk, additional risk comes from a potential global "winner take all" Internet product competition. Some applications may benefit from being localised, but others may develop global application strengths that blow past the nice to have local specialisations.

The upside opportunity is the flip side of the "winner take all" risk. A company in a small country such as New Zealand could develop a competitive application with industrial scale infrastructure layers such as Amazon's AWS, which provides global reach.

In the context of national opportunities on 15 November 2021, Google announced a US$1bn Google Research Laboratory investment in Australia (Sam Shed, 2021). AI is an explicit part of the new lab's research agenda. Google has also made investments in New Zealand over the past 15 years. New Zealand could benefit significantly by additional partnering with international companies on large AI research laboratories and AI applications investments.

Many New Zealand companies haven't yet fully digitised and taken advantage of machine learning utilities. This presents a major set of opportunities for local service providers to help NZ companies specialise global AI application components and solutions to their specific needs.

"AI Inside" vs. AI doing the heavy lifting

The giant semiconductor company Intel insists that personal computer makers include the label "Intel Inside" on the outside of their computers. AI has become so popular with investors and consumers that even the aspirational claim of "AI Inside" in product specifications has become common. Unfortunately, the claim of "AI Inside" may be true in the sense that an AI algorithm is included in the source code library, but that fact is irrelevant if the code is never executed as a meaningful part of the application solution.

Sometimes it doesn't matter much if a customer is buying a rice cooker. However, if investors need to know what is under the covers of a large investment in a company's application solution, then careful due diligence is required. It is common for reviewers of AI technical application papers to discover that a claim that AI is doing the heavy lifting in an application solution is exaggerated at best, and the system is really doing conventional mathematical optimisation.

AI due diligence checklist

Analysis of AI claims by startups or corporate application teams does not typically require deep technical knowledge, although that doesn't hurt. What is critical is asking a few key questions:

1) What kind of AI is used, and what role does this AI specifically play in the proposed solution?

2) How narrow or general is the result, or what is the actual range of the application?

3) Is there a demo available, and if so, is the demo engineered beyond typical performance in the real world? Can the demo examples be changed by us without severe performance degradation?

4) How carefully qualified are the application's performance claims — better than what "world class" standard?

5) How much training and tuning does the system require, and is there a shortcut or "cheat" in the success of the system?

For example, the IBM Watson AI took advantage of the fact that 90% of Jeopardy answers were Wikipedia headers. That made the system look much smarter than it actually was. Overall, due diligence should be penetrating but balanced in its assessments.

AI tools and applications

Most AI applications have common components which can be considered utilities. These components include: a "data wrangling tool" that enables the applications team to marshal and clean the data from noise, a data object storage and retrieval system, one or more pattern recognition or problem solving algorithms such as deep learning, reinforcement learning, or a mathematical optimisation module, a user management and interface tool, a customer service component, and a system test harness.

Each of these components may be bought as a commercial off the shelf generic product or specialised from open systems development modules found on GitHub and similar services. Commercial products cost more initially but the cost of upgrades and maintenance are borne by the providers.

Open system modules must be tested carefully for compliance to specifications, and in addition to necessary company specific specialisations, the costs of upgrades and maintenance must be borne by the applications company.

There is no right answer to the tradeoffs associated with using a commercial tool or specialising an open system component. However, for the list of common utility components discussed above there must be a very compelling reason to start building from scratch using a programming language such as Python.

Domain and task specific applications

Application system utilities or components may be specific for a particular generic task such as diagnosis, planning, or customer service. However, in order to work properly in a specific application area or domain, the generic task system often needs to be specialised for a specific domain, such as automotive vs. medical diagnosis. The data used in pattern recognition systems may include both the task and the application domain.

AI maturity models

Organisations differ significantly in their maturity using AI and machine learning. There are dozens of characterisations, each with some merit. One of the simplest and least technological models is from the Gartner Group. The primary merit of this model is that it characterises an organisation's functional relationship with AI. The focus is on what the organisation is actually doing with AI. The goal is to move through levels 1–5 until AI is used pervasively throughout the organisation, and it evolves to become an integrated part of the organisation's dynamic business DNA.

Figure 64. Some action steps for applying AI effectively

Some Action Steps for Applying AI Effectively

Start with high value problems, not AI technology
Disrupt your industry with AI/ML – or others will
Identify & track technical and business metrics
Utilize AI applications frameworks and best practices
Adopt a cloud-based AI platform, replace modules
Develop internal/external team & champions
Design for security, ethics, and future generations.

©Jacobstein, 2021

Source: Forsyth Barr analysis, Jacobstein

Figure 65. AI in New Zealand

Source: Forsyth Barr analysis, AI Forum NZ

Figure 66. AI maturity model

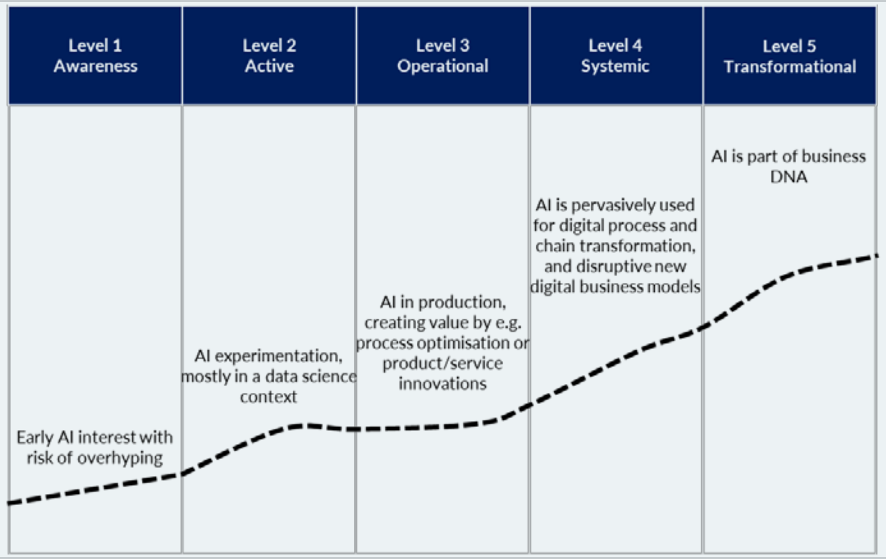

Source: Forsyth Barr analysis, Gartner

Appendix 2:
NZ Small cap and unlisted companies

As discussed above, NZ has many small companies which utilises Artificial Intelligence to deliver products and services. The list below is by no means exhaustive and ranges from companies that have built its entire business around AI, and larger listed companies like Vista (VGL) and Gentrack (GTK) which incorporates AI into its offering. We have not analysed these companies in detail but they caught our attention and we believe AI focussed investors in NZ will benefit from getting to know them better; odds are that one day a few of them will be large cap stocks.

Listed companies

Eroad (ERD): Eroad designs and manufactures in-vehicle hardware, operates secure payment and merchant gateways and offers webbased value-added services. The business aims to modernise road charging and compliance for road transport by replacing paper based systems with easy-to-use electronic systems. It is currently the largest provider of road user charges (RUC) compliance in New Zealand, and a leading provider of health and safety compliance, as well as fleet management solutions. We understand Eroad has limited use of AI currently.

Gentrack (GTK): Gentrack provides revenue, operations and customer engagement software platforms for a range of companies and sectors which includes a number of global utilities and airports customers. Gentrack utilises AI powered advanced data analytics to help companies organise their data from a range of sources and make it easily accessible.

ikeGPS (IKE): IKE seeks to be the standard for collecting, managing, and utilising AI and statistics to analyse pole and overhead asset information for electric utilities, communications companies, and their engineering service providers.

Livestock Improvement (LIC): Livestock Improvement Corp is an agri-tech and herd improvement co-operative. The company's AI powered MINDA herd management solution allows farmers to make better and faster decisions about their animals.

Rakon (RAK): Rakon is one of the world's largest manufacturers of frequency control products and timing solutions. Its products can be found in a wide range applications, from 5G networks to satellites and from emergency beacons to autonomous vehicles. At present the company does not explicitly claim to use AI functionality, but indicated that they would be open to deploying AI should an appropriate opportunity exist.

Scott Technology (SCT): Scott provides market-leading technology in order to transform industries using dangerous, dirty and repetitive manual processes with cost-effective, productive and safe automation. It focusses on five key industries; meat processing, mining laboratory automation, appliance manufacturing, materials handling and industrial automation. Their solutions are redefining end-to-end automation and cover all stages from design through to commissioning and service, and everything in between. Scott has used AI capabilities for a number of years — predominantly in its meat processing division in order to automate decisions (such as best cut lines when lamb boning). The company is expecting its other divisions to also use some AI applications in the future.

Vista Group International (VGL): Vista is the world leader in cinema management software. While Vista doesn't claim to use AI currently for internal purposes, it does provide a cutting-edge AI offering through its Movio business. The product leverages AI-driven algorithms to market specific movies based on consumers behaviour, preferences, past history and demographic.

Wellington Drive (WDT): Wellington Drive Technologies supplies electricity-saving, electronically commutated motors and fans worldwide. Their focus is on advanced motors, electronics and software that save power. The company makes motor housing from industrial plastics, rather than from stamped metal parts.

Unlisted companies

Actassa: Actassa provides AI-enabled conversational (chatbot) software solutions for law firms that have the potential to stimulate client engagement and increase firms' productivity.

Ambit: Ambit's AI-powered chatbots are designed to deliver automated customer service at scale to help businesses maximise revenue, reduce costs and better engage customers.

Amy: Amy is an AI-based private tutor for math that is intended to make learning math easier for everyone. It does this by giving students feedback and automatically addressing their knowledge gaps as they learn.

ARDA: ARDA is a software development kit that allows AI coaching features to be implemented in fitness hardware or applications.

Arria: Arria NLG is a form of artificial intelligence that transforms structured data into natural language. Arria dynamically turns data into written or spoken narrative at machine speed and on a massive scale. The vision is to give data the power of language.

BVT: BVT has an intimate understanding of the seismic requirements of building in New Zealand and abroad. BVT uses Technology Enabled Engineering, like 3D virtual reality capture, artificial intelligence and high-value engineering tools to empower its clients to solve problems with cutting-edge technology.

CBA Strategic IT: CBA Strategic IT is focussed on research, development and sales of Cognitive Solutions (AI) for the global market.

Dexibit: Dexibit enables leaders and teams to democratise data and switch from gut feel to insight informed with the latest in big data analytics and artificial intelligence. Dexibit's insights inspire decisions to get more visitors through the doors, engaged in rich experiences and returning loyally.

Halter: The Halter app uses a combination of the Internet of Things and artificial intelligence to help farmers guide and manage their dairy cows; maximising production and significantly reducing on-farm workload.

Imagr: Imagr has built a vision-only, white-label autonomous checkout solution that retailers can own, operate and scale themselves. Its proprietary AI doesn't need barcodes. It just recognises what is being put in the cart or being taken out.

Iris Data Science: Iris Data Science is an artificial intelligence and machine learning company. Its latest system OmniEye, is an innovative automated on-farm livestock monitoring system that provides real time locomotion scores for a herd and allows powerful early detection of lameness or other animal traits and welfare issues.

Molemap: MoleMap is New Zealand's leading pattern recognition powered detection, diagnosis, and surveillance service for melanoma and other skin cancers. Its range of comprehensive skin check and mole mapping services are designed to catch skin cancer early — when it's most treatable.

Ohmio: Developed by HMI technologies Ohmio's AI powered self-driving vehicle system allows vehicles to be deployed quickly. Its mapping capability provides the means for the vehicle to learn its route and improve its performance using machine learning in the process of repeating the route.

Opum: OPUM's intelligent platform is an advanced orthopaedic data analytics solution, powered by AI. Opum uses AI to provide actionable data and insights to help manage patients in recovery, and for long-term knee health.

Orbica: Orbica uses geospatial, data, and geo-AI to add value to businesses. Geo-AI is where geospatial technology, imagery and artificial intelligence converge. Orbica has developed algorithms that can detect and classify features with a high degree of accuracy and speed, revolutionising traditional methods.

Orion Health: Previously listed Orion Health is leading ground-breaking research in machine learning, exploring meaningful ways to minimise waste, reduce operating costs, and help clinicians make more accurate decisions at the point of care. Significant amounts of data exist that will support better decision making, drawing on information from entire populations to treat and manage a person's health.

Quantiful: Quantiful's QU is an AI powered SaaS tool that generates highly accurate product forecasts and predicts what will drive demand in future.

Robotics Plus: Robotics plus is an agricultural robotics company with a number of products that combine robotics and AI to solve real world agricultural problems, including robotic apple packagers and a pioneering robotic log scaling machine.

Rush digital: Rush is an integrated design and technology studio whose engineers and UX designers use a range of technologies, like Machine Learning, Computer Vision, Cloud, Internet of Things and Natural Language Processing. These technologies are deployed to bring ideas to life across different customer touch-points.

ScaleXT: ScaleXT is using an AI powered precision hiring intelligence platform. This platform includes humans at the helm and is designed to make hiring faster and more effective.

Soul Machines: Soul Machines' patented Digital Brain allows their animated chatbot agents to talk, analyse data, and interact. The agents are AI plus digital animations generated dynamically, and Soul Machines refer to their chatbots as "Digital People". The intention is to create agents that achieve zero touch, strong user engagement, and unparalleled brand experiences from retail to healthcare.

Spacetime: Spacetime leverages unstructured data and artificial intelligence tools to build solutions that automate processes and transform user experiences including smart document analysis, natural language search, chatbot and virtual agents, and classification and routing.

Springload: Empathy-emulating human-centred design is how Springload approaches artificial intelligence. By making machine learning, natural language processing, chatbots, and computer vision more people-focussed Springload believes it can create more authentic experiences.

Touchpoint: Touchpoint's Ipiphany is an AI Text Analytics technology that uses natural language processing to read and categorise text feedback. Ipiphany is used by customer-centric enterprises to drive business understanding and deliver relevant customer insights at the speed and scale they require to make a difference.

Trademe: Previously listed TradeMe uses AI to predict the items people want to buy based on previous searches and other users' similar searches. The technology knows not only what we tend to buy, but what people like us tend to buy, and it remembers that data to rank ad listings in an order that makes sense to individual browsers.

Wine searcher: Wine-Searcher hosts a large and accurate collection of drinks-related price and location data available. Its technology includes a bespoke label matching AI system.

Xtracta: The company boasts a world-class Artificial Intelligence (AI) research team that brings together considerable academic knowledge. Several team members hold PhDs in AI and have commercial applications experience from around the world.

Appendix 3: References

Adam Stookes, et. al. 2021. "Open-ended Learning Leads to Generally Capable Agents". arXiv:2107.12808.

AI Forum New Zealand, Artificial Intelligence: Shaping a Future New Zealand. 2018.

AI Forum New Zealand, Trustworthy AI in Aotearoa: AI Principles. 2020.

Ali Alvi and Paresh Kharya, "Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, the World's Largest and Most Powerful Generative Language Model". Microsoft Research Blog, October 11, 2021.

Andrew Ng: "Why AI Is the New Electricity", Insights by Stanford Business (March 11, 2017).

Beth Kindig, "Palantir IPO: Deep Dive Analysis" ["The truest, closest competitor to Palantir is Semantic AI"...], Forbes, September 29, 2020.

Dan Jackson, The Netflix Prize: How a $1 Million Contest Changed Binge-Watching Forever, Thrillist Entertainment, 8/7/2017.

Daniel Zhang, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara Grosz, Terah Lyons, James Manyika, Juan Carlos Niebles, Michael Sellitto, Yoav Shoham, Jack Clark, and Raymond Perrault, "The AI Index 2021 Annual Report," AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA, March 2021.

Enrique Dans, "Stop Experimenting With Machine Learning And Start Actually Using It", Forbes, Jul 21, 2019.

David Bollier (Rapporteur), "Artificial Intelligence and the Good Society: The Search for New Metrics, Governance, and Philosophical Perspective". A Report on the 3rd Annual Aspen Institute Roundtable on Artificial Intelligence. 2019.

David De Cremer and Garry Kasparov, "AI Should Augment Human Intelligence, Not Replace It." in Harvard Business Review Special Issue, How AI is Changing Work, February 8, 2022.

Gartner Group, AI Maturity Model, February 2019.

Gary Marcus, "Deep Learning: A Critical Appraisal". arXiv, 2018.

Gary Markus and Ernest Davis, Rebooting AI: Building Artificial Intelligence We Can Trust. Pantheon, New York, September 2019.

Henry Kissinger, Eric Schmidt, Daniel Huttenlocher, The Age of AI: And Our Human Future. 2021. Little, Brown, and Company, New York.

Hitzler, Pascal, "A Review of the Semantic Web Field", Communications of the ACM, February 2021, Vol. 64 No. 2, Pages 76-83.

Investment Banking Council of America, "AI in Investment banking - The New Frontier", September 11, 2019.

John Riedl, et. al., Word of Mouse: The Marketing Power of Collaborative Filtering, Business Plus, 2002.

Jonathan Stokes, et. al, "A Deep Learning Approach to Antibiotic Discovery" Cell Volume 180, Issue 4, 20 February 2020, 688-702.

Kai-Fu Lee and Chen Qiufan, AI 2041: Ten Visions for Our Future. Currency, New York, 2021.

Kai-Fu Lee, AI Super Powers: China, Silicon Valley, and the New World Order, Houghton Mifflin Harcourt, New York, 2018.

Kevin "Rotterdam" van der Kooi, Waypoint, Feb 7, 2019

Markham, "EVs approaching 'iPhone moment' and that changes everything – Tony Seba", EnergiMedia, June 1, 2017.

Michael Lewis, Flash Boys: A Wall Street Revolt, W. W. Norton & Company, 2014.

National Security Commission on Artificial Intelligence, Final Report, 2021.

Neil Jacobstein, "Autonomous vehicles: An imperfect path to saving millions of lives", Science Robotics, Vol 4, Issue 28, March 27, 2019.

OpenAI, OpenAI.com/blog/better-language-models, February 14, 2019.

Paul Cohen, DARPA's Big Mechanism Program, 2015 Phys. Biol. 12 045008.

Petar Veličković and Charles Blundell, Neural Algorithmic Reasoning, Patterns Cell Press

Reid Smith & J. Eckroth. "Building AI Applications: Yesterday, Today, and Tomorrow", AI Magazine, 38(1), 2017, 6-22.

Ryan Poplin, et. al.," Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning", Nature Biomedical Engineering Volume 2, p158–164 (2018).

Sam Shead, "Google's Australia Investment Could Be a Big Boost for the Nation's A.I. Scene." CNBC, November 19 2021.

Steven Pinker, "We're told to fear robots. But why do we think they will turn on us?" Popular Science, February 13, 2018.

Stuart Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, 4th Edition, Pearson, 2020.

Tim Berners-Lee, James Hendler, Ora Lassila, "The Semantic Web: A New Form of Web Content That is Meaningful to Computers Will Unleash a Revolution of New Possibilities". Scientific American, Vol. 284, No. 5, May 2001.

Tuomas Sandholm and Anton Likhodedov, "Automated Design of Revenue-Maximizing Combinatorial Auctions", Operations Research, 2015, vol. 63, issue 5, 1000-1025.

Tuomas Sandholm, et. al., "Expressive commerce and its application to sourcing", IAAI06, Proceedings of the 18th Conference on Innovative Applications of Artificial Intelligence, Vol2, July 2006, p 1736-1743.

Vikram Mahidhar and Thomas H. Davenport, "Why Companies That Wait to Adopt AI May Never Catch Up." in Harvard Business Review Special Issue, How AI is Changing Work, February 8, 2022.

Wall Street Journal, "How Artificial Intelligence Will Change Everything: Baidu's Andrew Ng and Singularity's Neil Jacobstein say this time, the hype about artificial intelligence is real." March 6, 2017.

Will Knight, "The Pentagon Inches Toward Letting AI Control Weapons: Drills involving swarms of drones raise questions about whether machines could outperform a human operator in complex scenarios.", Wired, May 18, 2021.

William Fedus, Barrett Zoph, and Noam Shazeer. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity", ArXiv: 2101.03961v1, csLG, 11 Jan 2021.

FORSYTH BARR

This report is printed on environmentally responsible paper stocks manufactured under the environmental management system standard AS/NZS ISO 14001:2004 using Elemental Chlorine Free (ECF) pulp sourced from sustainable well-managed forests.

Processes for manufacture include use of water-based sealants, with disposal and/or recycling of waste materials according to best business practices.