Machine Learning in Australia

National Pulse Report 2021

DIUS

Contents

Introduction

Machine Learning (ML) has an important role to play. Here at DiUS, we believe it will define the future of technology in many ways. Accordingly, we have invested in a specialist ML practice to take us—and our clients—forward in this emerging area.

We've been talking with many companies about leveraging ML over the past five years and found most organisations want to adopt ML, yet this transformative technology is not being adopted at the rate it should be.

Indeed, our experience has been that a great proportion of organisations struggle to move beyond proof-of-concept or pilot stage. Picking the right problem, data challenges, model accuracy and application integration can be big blockers, either delaying or preventing ML project success.

However, we recognise that our observations may be incomplete, so we undertook a pulse of our clients and the broader Australian market.

The survey results from the 205 respondents confirmed what we are seeing: more organisations could be driving success with ML. There is a strong appetite for ML in the Australian market with 82% of organisations interested in ML, but only 21% in production. On the positive side, the majority of those that are in production are reporting positive business outcomes.

In this report, we dig into what's behind this gap; what's preventing more organisations converting an interest in ML into success? We also offer insights and tips that may help organisations wanting to progress their ML Journey.

Ultimately, the pulse survey confirms there is no one formula for success; there's no one application or area doing disproportionately better than others. There are some significant challenges, however they are not insurmountable with organisations still progressing despite these hurdles.

Here at DiUS, we remain ever optimistic, understanding that ML is still an emerging area and there is no universal 'best practice' yet. There's no set and forget solution for ML. The right approach, we've found, is to focus on the right problems, take an experimental approach, invest continuously in the latest advancements including cloud vendor tooling and open-source frameworks, build inhouse ML and supporting capabilities and ensure organisational buy-in.

Key findings

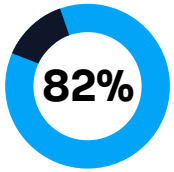

8 in 10 have started their ML journey (82%) and 54% have adopted this new technology by applying it to solve business problems with initiatives either in proof of concept or production.

ML projects are delivering business value, of those respondents with ML models in production 81% report successful business outcomes.

ML adoption is going to accelerate, 86% of respondents see ML as critical or one of several important technologies going forward, and 49% of those who have not yet started plan to do so in the next 12-24 months.

Top ML use cases are internally focused, the top two business areas are operational efficiency (48%) and business decision making (46%). Going forward, respondents plan a shift to both an internal and external focus: operational efficiency (57%) and customer experience (51%).

Succeed with ML by making it a priority, 79% of respondents achieving success with ML have an ML strategy, suggesting that focus and investment drive outcomes.

Invest in your data, data-related challenges are either the top or second most commonly reported challenges once the ML journey is started. The importance of data quality, data engineering and building appropriate data infrastructure and pipelines to enable ML initiatives cannot be overstated.

Even advanced ML adopters are challenged by skill gaps, only 69% of organisations with models in production report sufficient ML capability.

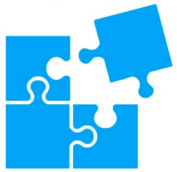

It's horses for courses in ML, there's so much diversity in where and how organisations are applying ML. Business areas, applications and types of ML all show a very broad distribution across multiple aspects of each dimension. It's important to find the right problem to solve and the right approach to solve it.

ML adoption

Interest vs adoption

ML is firmly on the agenda for Australian organisations with 82% of respondents having embarked on an ML journey within their business. We've defined the start of an ML journey as the consideration stage, where a business focuses on selecting and prioritising the most promising use cases for an ML initiative, and adoption is when organisations have initiatives either in production or at a proof-of-concept stage.

That eight out of ten Australia organisations have started an ML journey and over half (54%) are adopting ML is an encouraging result, given the challenges typically experienced with the adoption of emerging technology.

Our local results are consistent with other global surveys on AI/ML. Australia is amongst the leading countries adopting ML. Kaggle's worldwide survey in October 2020 of 20,036 data professionals includes Australia as an early adopter with 53% reporting ML adoption activities, whether that's in proof-of-concept or production. The country with the highest rate of adoption is Israel, with 63%. Globally, it's 45%. In a 2020 global study by McKinsey, the global adoption rate of ML was 50%.

Ultimately however, it remains that the survey findings show only 21% of organisations are putting models into production. And that's a big gap between the large proportion that are interested in ML and those that are using it to drive business outcomes.

Figure 1: How far has your organisation progressed with ML?

Stage of ML journey	%
Not started	18%
Considering	28%
Experimenting	33%
Applying	21%

N = 205, all respondents

The survey findings suggest key ML adoption is correlated with having an ML strategy and possessing sufficient ML capability. A not-unexpected result for the adoption of any new and emerging technology, however a key result to highlight for an organisation seeking success with ML because of the challenging and complex nature of ML. As this field is still so new, best practices and tooling are still solidifying and success requires focused investment coupled with specialist skills.

Figure 2: ML Journey Stage, ML Strategy and ML Capability

N = 205, all respondents

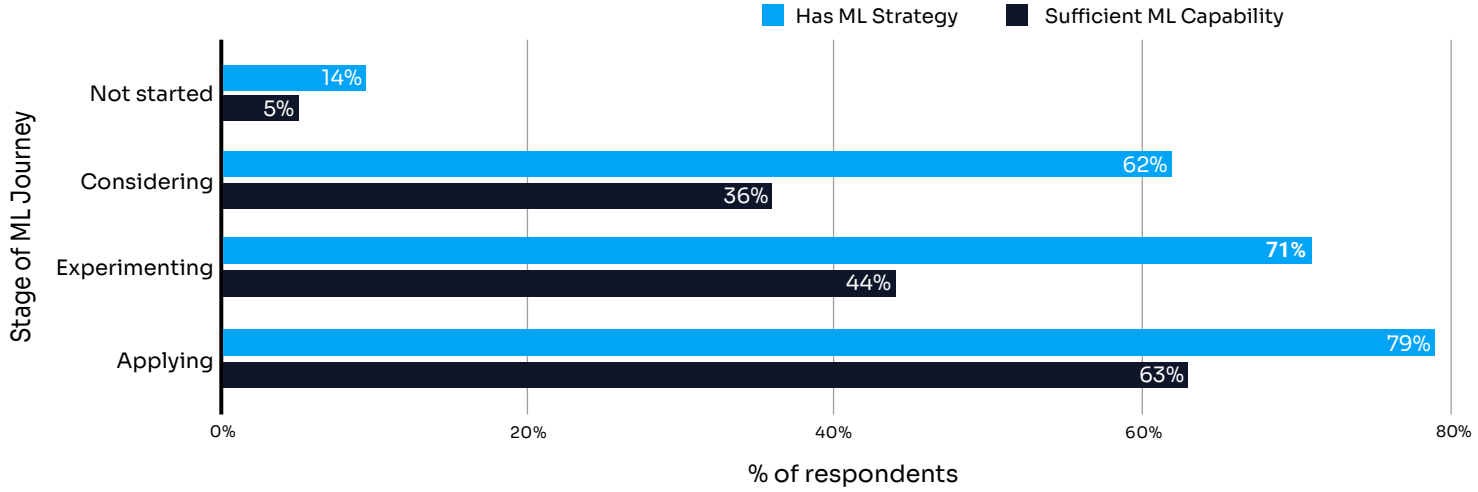

We expect ML interest in Australia to continue increasing, with:

- 49% of those respondents who have not yet started with ML, report organisational plans to start exploring ML in the next 12-24 months.
- 86% expect that ML is going to be a critical part of future organisational success or one of several important technologies contributing to success.

Given the promise of ML coupled with the strong interest demonstrated by the survey findings, we would expect more Australian organisations to be moving into the 'Applying' stage of the ML journey in the near future. However, the lag or gap between actual ML implementation and interest in ML means we can't expect with certainty that these figures will translate into future ML adoption within organisations. Our experience has been that the amount of effort and investment required to do ML well can be underestimated by organisations.

ML journey stage definitions

Not started	The organisation has not started exploring the application of ML other than possible formulation of an initial ML strategy.
Considering	The organisation has identified an appropriate opportunity or problem to solve with ML but has not yet begun.
Experimenting	The organisation is currently experimenting with a Proof of Concept or Pilot investigation.
Applying	The organisation has one (or more) ML models in production or being readied for production.

Business areas

Looking at the business areas where ML is being adopted—focusing on the group of organisations that are actively applying ML in a Proof of Concept (PoC) or in production—the most commonly reported business areas relate to operational efficiency, 48%, and business decision making, 46%, suggesting that internal processes may be the main focus of organisations' ML initiatives right now.

Our commercial experience supports an internal focus with the most common use cases we've seen for ML in areas that are labour intensive, repetitive and error prone as well as where predictive or prescriptive outputs can be derived to improve a process or support a decision.

There were no significant differences in the business areas where ML was being applied across the stages of the ML journey. However, as organisations progress through the ML journey, the number of business areas where ML is being applied increases. This confirms our commercial experience that organisations are starting with one business area, proving value and then expanding ML across other areas.

Figure 3: Business areas where ML is being adopted

N = 110, respondents in the experimenting and applying stages

When looking at future plans for applying ML in business areas, a similar broad distribution is found, albeit with slightly different proportions: operational efficiency remains the most frequent however the biggest increase was in customer experience. This suggests that internal focus will remain important, but ML's remit is moving into externally-focused areas for an organisation. Commercially at DiUS this is supported by a growing demand within the market for ML-powered digital products that create brand new customer experiences.

For example, DiUS recently worked with bolttech to use ML to quickly and easily onboard customers onto device protection plans. Read the case study.

It's clear ML is regarded as a valuable technology, suggesting that investment will continue to be made across multiple business areas and applications into the future.

Figure 4: Future plans for ML adoption in business areas

N = 205, all respondents

Business areas definitions

Business decision making	Business decision making can be supported by AI / ML in the form of expert assistance or automation.
Customer experience	Customer experience (CX) is a customer's perception of their experience interacting with a business. AI/ML can enhance this through personalisation or providing a new kind of experience. Examples are the integration of multi-channel communication, the appropriate automation of responses and providing personalised recommendations.
Operational efficiency	AI/ML can improve the efficiency of an organisation's day-to-day operations by supporting workers to perform tasks more quickly and consistently, or to perform them as if they had specialised expertise. For example, computer vision is used to screen microbial growth on sample cultures in pathology laboratories to allow microbiologists to focus on more complex cases.
Security	Various aspects of AI/ML can be applied to security. Computer vision can support facial recognition and automated analysis of CCTV footage. Structured / Tabular Data algorithms can be applied to financial crime detection. Anomaly detection algorithms can assist with cybersecurity intrusion detection.
Sales & marketing	The effectiveness and consumer experience of Sales and Marketing can be improved if promotions and offers are targeted at the consumers most likely to be interested. AI/ML has been used for many years to segment potential customers according to their interests and enable more targeted advertising. Most online stores provide some form of AI/ML-powered recommendation to consumers based on previous browsing and purchasing behaviour.

Organisation size

There weren't any strong differences regarding ML adoption within Australian industry sectors within our pulse survey. Based on our commercial experience, we would have expected fintech, industrial and other digital-based businesses that collect a large amount of data to present more strongly in comparison to other industries. The survey findings did however show that larger organisations have a higher ML adoption rate than smaller ones, a similar trend as reported in Kaggle's worldwide survey on Machine Learning.

We posited that larger organisations would be more likely to have a dedicated strategy and capability to drive ML success. It might also be that larger and more mature organisations have a history of data governance and collecting data, making ML a much easier next step in terms of leveraging data for insight and value.

Once organisation size reaches 250 employees the survey findings show a tipping point in terms of ML adoption. There is a jump in the share of organisations that are hands-on with ML—adopting ML via a PoC or through models in production—versus those that have not started or are focused on finding the most promising problems to solve using ML.

Figure 5: Organisation size and ML journey stage

N = 205, all respondents

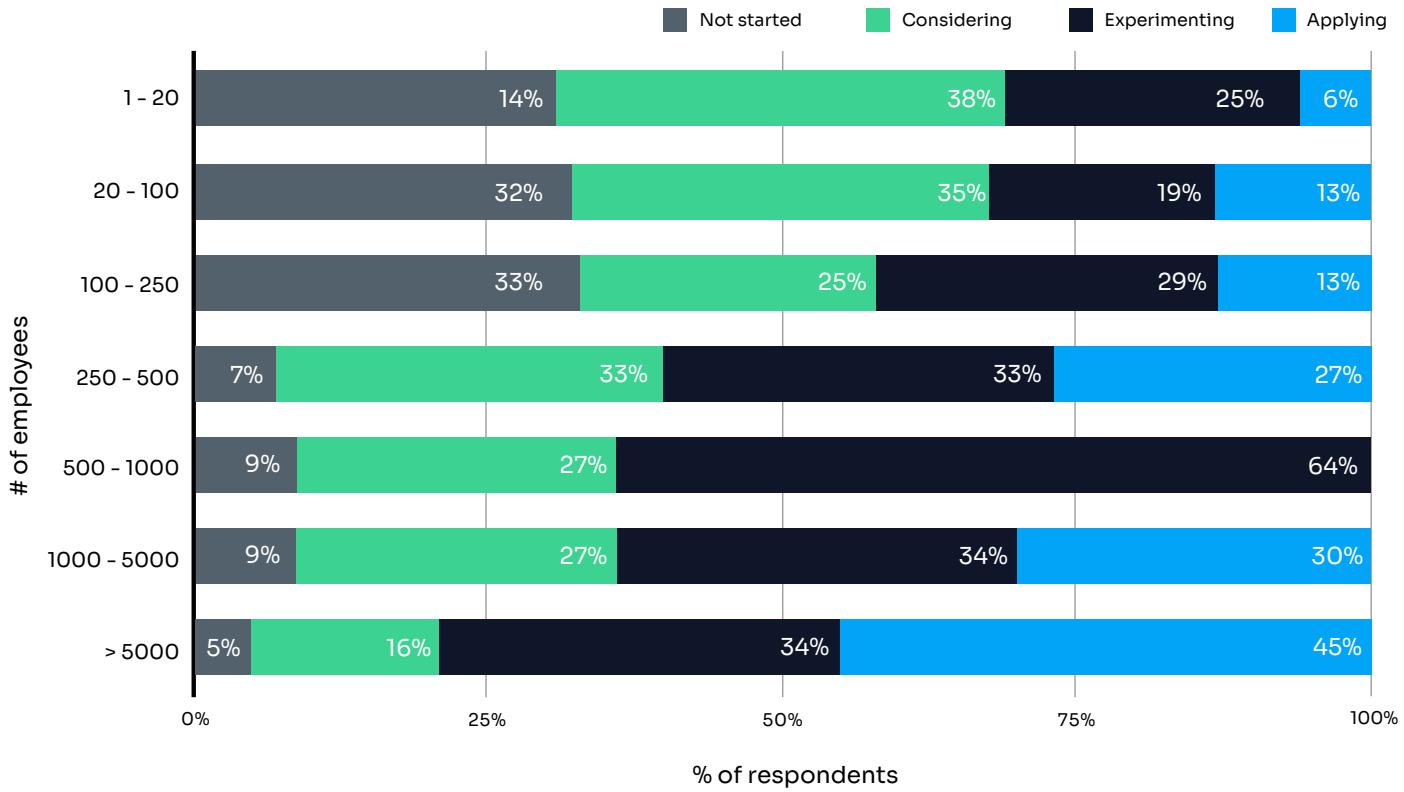

A notable outlier to the relationship between ML adoption and organisation size is the 500 - 1000 employees cohort. This group has the largest share of respondents that are experimenting in the entire sample but appear to be stuck at that stage of the ML Journey. That is, no ML models are being operationalised or productionsed.

The 500 -1000 employees cohort is slightly smaller than the others, so we dug deeper into the data to see if the outlier result was sample related. The only notable observation was despite organisations in this cohort experiencing similar challenges to others, the challenges are not slowing ML adoption, rather the challenges are more likely to halt their ML programs: 9% as compared to 5% across the entire sample.

We want to acknowledge that although our pulse report shows that larger organisations are doing better, our commercial experience is that smaller organisations, including startups, are in fact succeeding with ML. For startups and smaller organisations, we've observed that unless ML is part of a core business proposition, it's less likely to be considered until they achieve market fit and are ready to scale.

ML implementation

To gain insight into how ML is being implemented, we focused on the group of organisations that are actively applying ML in a PoC or in production—the 54% that are adopting ML.

ML application areas

The survey findings show a fairly even distribution across ML application areas*. Consistent with our commercial experience, there is no one application area significantly dominating the ML landscape. Rather, the survey findings show there's a lot of interest across the field of ML with organisations looking at how this emerging area can add business value in general, as well as seeking solutions for specific use cases.
Recommendation Systems and Tabular / Structured Data have emerged as the two top ML application areas right now. However, the distribution of ML application areas within business areas is quite equal, suggesting that the number of different use cases are increasing and requiring different applications of ML to provide solutions.

The popularity of Recommendation Systems likely comes in part from the fact this application area is well established, having been around for a long time, so it makes sense that it's being adopted. Tabular / Structured Data's popularity is also expected because the underlying algorithms are well-described and readily accessible.

* ML algorithms can be grouped according to the kind of input data they accept, which we refer to here as "application areas".

Figure 6: Current ML application areas

N = 110, respondents in the experimenting and applying stages

Looking at the survey findings on respondents' thoughts on ML future application areas, this broad distribution is not likely to change anytime soon, but the popularity of certain application areas will ebb and flow.

Figure 7: Future ML application areas

N = 205, all respondents

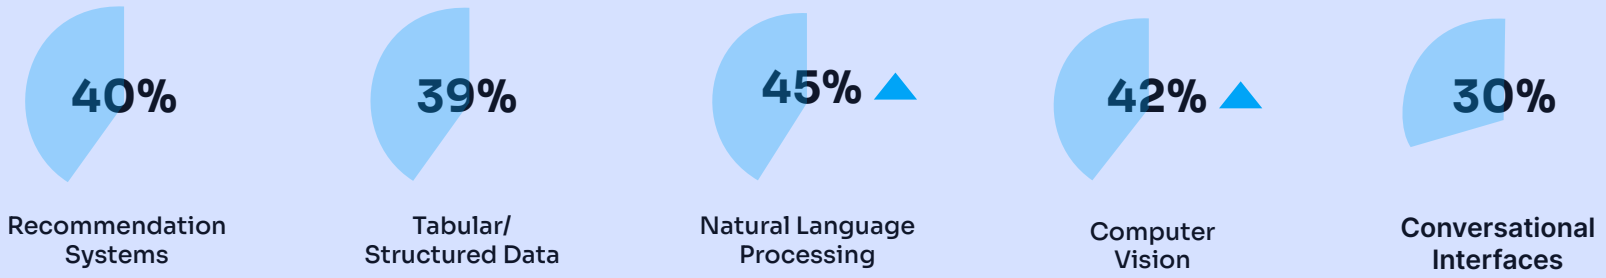

Natural Language Processing (NLP) overtakes recommendation systems as the top area organisations are looking to in the future and this is consistent with the strong interest in NLP's potential from our commercial conversations. NLP's ability to unlock the value in unstructured text makes it possible to power many new ML applications in industries such as customer service, human resources and healthcare.

DiUS and ML application areas

Here is an overview of what each ML application area means and some real-life examples from DiUS' work.

Computer vision

Computer vision enables computers to interpret and understand the visual world through digital images and videos. Using deep learning models, machines can learn to accurately identify and classify objects — and then take action on what they "see".

Examples of how computer vision is being used today include optical character recognition, machine inspection, automated checkouts in retail, 3D model building, medical imaging, automotive safety, surveillance, and fingerprint recognition and biometrics.

DiUS CLIENT STORIES

Datarock: ML-powered drill core analysis software

Launched in 2019, **Datarock** is a joint venture between DiUS and Solve Geosolutions. It is an ML-powered SaaS solution for mining and exploration. Datarock analyses digital photos of rock samples in drill core trays using computer vision, enabling important geological information about the site to be delivered swiftly and help efficiently direct mining operations.

For example, it may have previously taken a company up to six weeks to obtain geo-chemical assay results from a mining site using geologists. Using Datarock, companies can get proxies within hours. This near real-time power enables more intelligent decision-making regarding the viability of a particular mining site.

Read our **case study** to find out more. This solution was a finalist in both the CRN Impact Awards and the ARN Innovation Awards in 2020.

Conversational interfaces

A conversational interface is a user interface for computers that emulates a conversation with a real human either via text or voice. They're designed to foster a personal connection between organisations and their customers. They can also facilitate better user experiences by having a built-in customer service representative for a digital product.

Voice interfaces include virtual assistants like Siri, Cortana, Alexa, or Google Assistant in mobile devices. Text interfaces are usually used for websites and apps, and use an interface called a chatbot. Chatbots can find you makeup, help you order flowers, book you a cab, or help you buy a plane ticket.

DiUS CLIENT STORIES

nib: chatbot uses ML to assist with customer enquiries

DiUS developed a conversational interface to help nib to build nibby - a chatbot that helps its members with their everyday enquiries about health insurance and reduces pressure on frontline employees.

nib was looking for a solution to proactively and quickly direct enquiries to the correct service area without the need for the customer to go via a live chat agent. nibby meant nib was the first Australian health insurer to introduce artificial intelligence (AI) services to assist customers with their health insurance enquiries.

In its first three months, nibby handled more than 3,000 customer interactions with success in routing the queries to the appropriate specialist channel. In the first 12 months, nibby saved 535 hours of customer handling time.

Read our case study to find out more. DiUS was also awarded the CRN Impact Award 2018 for Customer Experience for the nibby project.

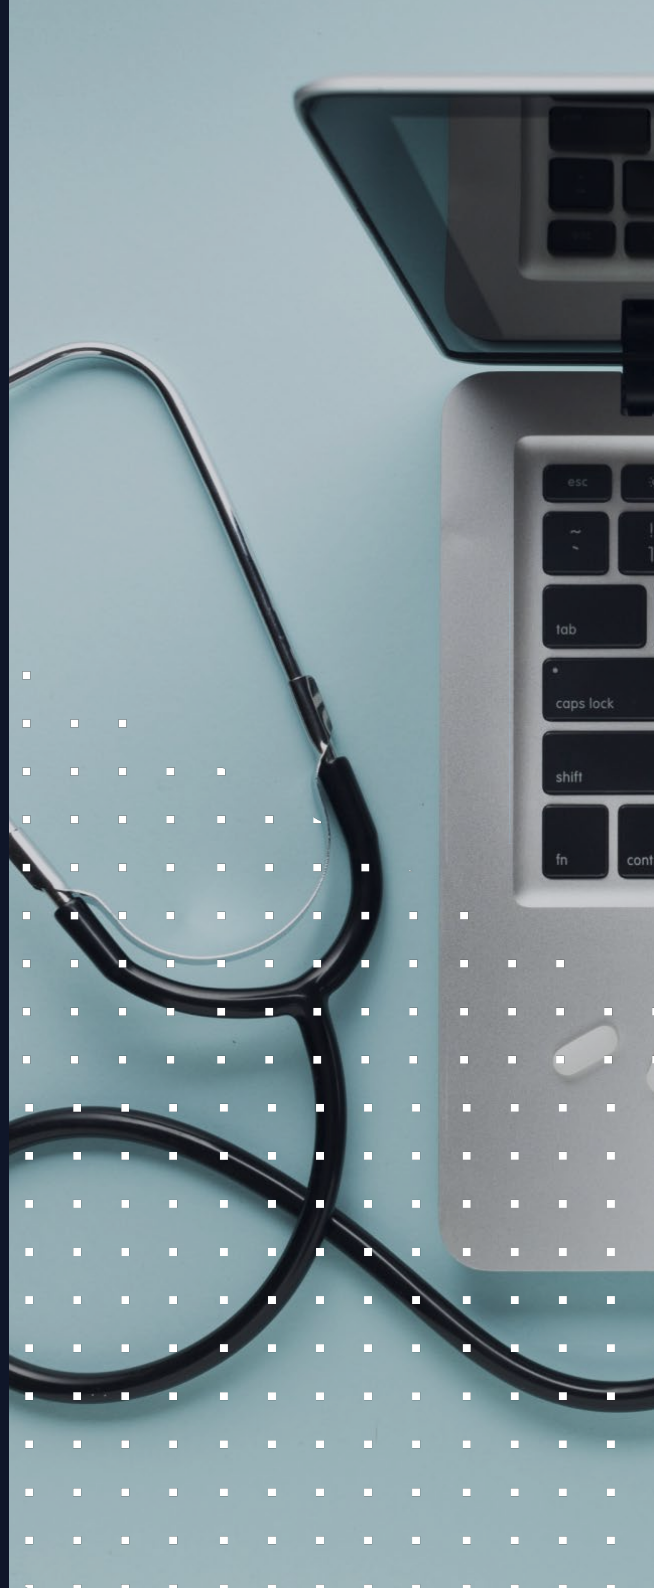

Natural Language Processing

Natural Language Processing (NLP) describes the interpretation of human language, either speech or text, by computers.

In NLP, human language is separated into fragments so that the grammatical structure of sentences and the meaning of words can be analysed and understood in context. This helps computers read and understand spoken or written text in the same way as humans.

A few simple examples of NLP that people use every day are spell and grammar check, autocomplete texting, voice to text messaging and related words and topics on search engines.

DiUS CLIENT STORIES

Using ML to detect brain bleeds in CT scans

DiUS proved the value of using ML to automate detection of brain slices in a CT scan, and detection of brain bleeds, to help prioritise these cases in a radiologist's work queue.

Radiology Services handles 50,000 radiology studies each month so it's important that time-critical cases like brain bleeds are put to the top of the queue. Two hours can be the difference between the life and death of a patient.

An NLP Classifier was built to perform textual analysis of metadata and content of radiologist reports to identify previous cases of brain bleeds so those images could be included in the training data set. Then computer vision was used to develop an image segmentation model to identify brain bleed evidence on CT scans.

Recommendation systems

A recommendation engine is a system that suggests products, services and information to users based on analysis of data from various sources, such as the history of the user and the behaviour of similar users.

Recommendation systems can increase engagement, conversion and purchases on sites where users (identified or not) are searching and browsing a wide range of content, products or services. High-performance recommendation engines employ a technique known as collaborative filtering, which finds similarities between products based on similar interactions of users. A simple example is how Netflix makes movie recommendations based on movies users have watched in the past.

DiUS CLIENT STORIES

An accommodation booking system that makes recommendations with minimal data

DiUS built a custom recommendation engine for an accommodation booking site to lift the conversion rate from searches into bookings, all without that essential user information that's traditionally used to power sophisticated recommendations.

The problem was solved by building proxy clusters for users and their interactions. The property clusters were built from query data (user searches), based on features such as the number of adults, whether or not there were infants, the number of stay days, a request for a pool and so on. Using only this new clustering approach produced an uplift of 48X the existing model, which had originally been created using a manual decision tree. The work was successful in demonstrating that a recommendation engine could be developed without personal user data and without being able to identify who a user is. Read our blog to find out more.

Tabular / structured data

Tabular / structured data refers to applications where the relevant input data fit well with the regular layout of a relational database or spreadsheet.

The most common forms of structured data are numbers and text, for example names, dates, addresses, credit card numbers, stock information and geolocation. The stores of these types of data are all likely to be labelled and quality checked, creating high quality sources of data for machine learning.

DiUS CLIENT STORIES

ML-powered sales forecasting

Identifying successful and repeatable sales models was key to growing a major retailer's brand overseas. DiUS helped identify products with similar sales patterns and used machine learning to generate more accurate sales forecasting.

Around two billion transactions were analysed using cluster analysis and time-series forecasting techniques using FBProphet—an open source forecasting algorithm from Facebook.

MLaaS vs custom models

The largest share of respondents that are adopting ML, 52%, are leveraging both Machine Learning as a Service (MLaaS) and custom ML models in their organisation, suggesting growing maturity and understanding of when to best leverage either option or both.

A driving factor is likely the acceleration of ML tooling from cloud-based data science and machine learning vendors, as well as the adoption of open source machine learning models and libraries.

Certainly this is supported by the observation that a majority of organisations employ a combination of both MLaaS and custom models, regardless of ML application area or business area.

This is supported by our commercial experience. An organisation may start with a custom model or use an MLaaS approach, but the convenience of MLaaS for some applications or a growing confidence with custom models leads to a hybrid approach.

Figure 8: ML model implementation

N = 110, respondents in the experimenting and applying stages

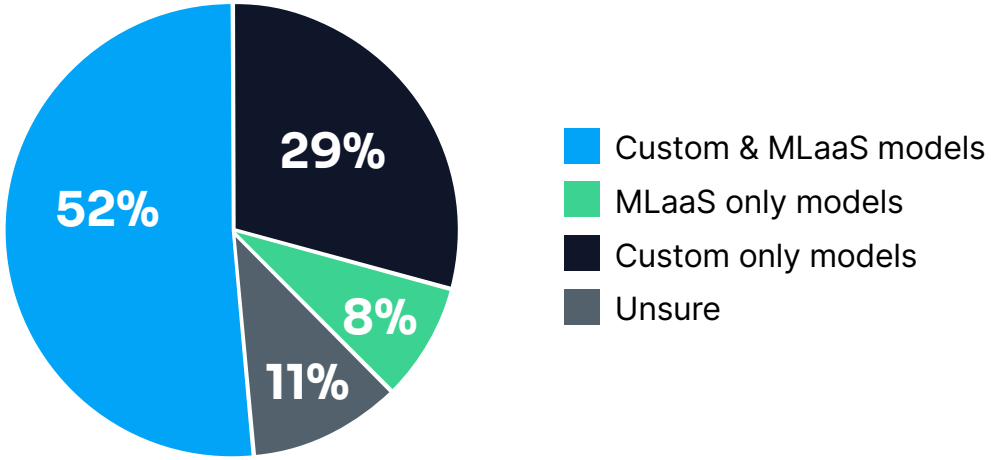

Some ML definitions

MLaaS	ML algorithms provided as a service by a public cloud provider. Some can be used with minimal understanding of ML algorithms and require no model training.
Custom models	ML algorithms provided by a library, or composed of components provided by a library, and then trained with data curated by the developer.
Custom & MLaaS models	The use of ML algorithms of both categories either in combination to solve a problem, or to solve different problems.

Business impact

According to Gartner's predictions, "Through 2020, 80% of AI projects will remain alchemy, run by wizards whose talents will not scale in the organization". While we acknowledge that ML at scale is hard to achieve and that getting models to production is challenging, the survey results do not support a high rate of failure once organisations get to production.

Despite only 21% of respondents having one or models in production, Australian businesses are seeing business results from an investment in ML. The majority, 81%, of respondents with one or more models in production have experienced successful business outcomes. Further, only 14% of this cohort reported failing to achieve the business outcome / ROI as a challenge being faced.

This is a highly encouraging result, supporting the case for further investment in ML and also acts as an encouragement for organisations that are yet to dip a toe in the water.

Challenges to ML adoption

Stage of ML journey

Respondents were asked different questions about the challenges they were facing depending on their stage of the ML journey.

As expected, challenges being faced differ according to where organisations are on their ML journey. The top challenge when getting started or prioritising a business problem is making ML a priority (41%), as compared to a lack of ML skills (49%) when experimenting and data related (56%) followed by application / integration complexity (51%) when putting models into production.

Figure 9: Top three challenges faced by organisations at the different ML Journey stages

Not started

Have more important priorities	**41%**
Haven't identified an opportunity or problem to solve	**35%**
No business case or funding dedicated to ML	**24%**

N = 37, all respondents in the not started stage

Considering

Have more important priorities	**40%**
Data related	**34%**
Application / integration complexity	**33%**

N = 58, all respondents in the considering stage

Experimenting

Lack of ML knowledge and skills in-house	**49%**
Data related	**40%**
PoC objectives not well defined or measurable	**31%**

N = 67, all respondents in the experimenting stage

Applying

Data related	**56%**
Application / integration complexity	**51%**
Scaling and performance	**49%**

N = 45, all respondents in the applying stage

Data, Data, Data

One of the hardest parts of ML initiatives is data availability; it's vital to have data collected in a structured, clean and centralised way in order to train ML models. This is very often not the case. And indeed the survey findings show data-related challenges are in the top two challenges once the ML journey is started. All respondents that indicated a data-related challenge were also asked the underlying reasons. While there are differences across the ML journey stages, it's clear that again, there's no single reason behind the data-related challenges being experienced and more important, multiple reasons were common.

At the considering stage, where a business focuses on selecting and prioritising the most promising use cases for an ML initiative, lack of suitable / labelled data and a lack of data strategy had no responses. These challenges become apparent once an organisation starts experimenting with and training models.

At the applying stage the top challenge was inadequate data infrastructure and access (58%) consistent with a need for the appropriate data pipelines required when putting models into production. That a lack of suitable / labelled data and a lack of data strategy have the smallest share of respondents at this stage is consistent, as these are likely to have been addressed if an organisation has been training models for a while.

Notable was the emergence of data privacy as a key concern; it was the top reason behind data-related challenges in the considering (50%) and experimenting (51%) stages, and the second top reason (42%) in the applying stage. While data protection laws such as GDPR regulate the collection and use of personal data, much remains to be done around regulating fairness and explainability at a global level.

Figure 10: Stage of ML Journey & reasons behind data-related challenges

Considering

Regulatory or data privacy concerns	**50%**
Issues relating to dataset management	**50%**
Inadequate data infrastructure and access	**45%**
Lack of suitable / labelled data	**0%**
Lack of data strategy	**0%**

N = 20, all respondents in the considering stage that selected 'data related' as a challenge

Experimenting

Regulatory or data privacy concerns	**52%**
Lack of data strategy	**44%**
Lack of suitable / labelled data	**41%**
Inadequate data infrastructure and access	**41%**
Issues relating to dataset management	**41%**

N = 27, all respondents in the experimenting stage that selected 'data related' as a challenge

Applying

Inadequate data infrastructure and access	**58%**
Regulatory or data privacy concerns	**42%**
Issues relating to dataset management	**38%**
Lack of suitable / labelled data	**29%**
Lack of data strategy	**29%**

N = 24, all respondents in the applying tage that selected 'data related' as a challenge

ML capability

The survey findings show that most organisations, 82%, have some ML capability, however only 39% have what they consider to be sufficient.

And when we look across the ML journey, while the share of respondents claiming sufficient ML capability increases from 5% in the 'not started' category to 63% in the 'applying' category—the survey results do indicate a lack of ML skills in Australian organisations. A result consistent with the large amount of commentary seen in the market about the lack of skilled ML resources.

Figure 11: Amount of Internal ML capability

ML capability	% of respondents
None	9%
Limited	43%
Sufficient	39%
Unsure	9%

N = 205, all respondents

Impact on ML adoption

Despite the many challenges being faced, it's encouraging to see that ML adoption is being slowed, rather than prevented. The majority of respondents, 86%, report that challenges are reducing the speed of ML adoption as compared to only 5% reporting a halt altogether.

Figure 12: Impact of challenges on ML adoption

N = 168, respodents in the considering, experimenting and applying stages

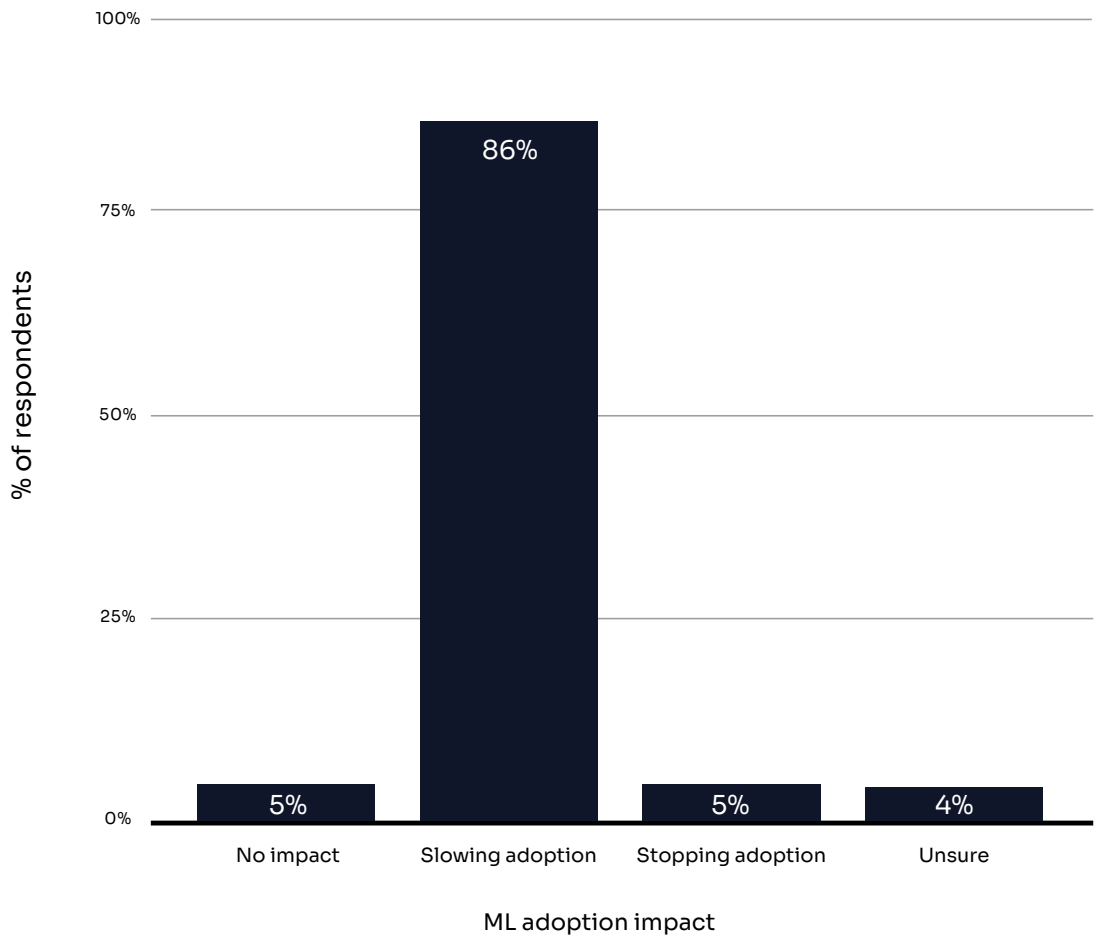

What’s next: how to progress your ML journey

ML can be applied in different ways to solve many business problems, therefore each organisation’s journey is unique. Overall, our recommendation for organisations that want to succeed with ML is to continuously invest from a technological, people and process perspective.

However, to help you move forward, we’ve looked at the survey results at each stage of the ML journey and developed some key considerations for organisations to look at.

For organisations that haven't started

Start by increasing your organisation's awareness of ML and how other organisations have successfully applied it.

Consider how ML might deliver business value for you, what internal resources your organisation is able to leverage and identify where more preparation might be needed.

Steps to get started:

- Educate yourself and the organisation about what ML is and what it is capable of
- Focus on understanding your key pain points and opportunities
- Identify which pain points / opportunities are potential candidates for ML
- Understand your organisation's internal data and ML capability
- Secure internal executive buy in and sponsorship for the adoption of ML in your organisation

Tips

Explore cloud services
They've done a lot of the heavy lifting in this area through purpose-built ML services that solve many general use cases. They are easy to understand and start with.

Internal efficiencies can provide quick wins
Look for manual tasks that can be repetitive, error prone or require subject matter expertise as these could be good candidates for task augmentation using ML.

Look within to understand internal capability
Success with ML requires more than just specialist data scientists; roles like data and cloud engineers also have important parts to play. It's ideal to identify these candidates early on to assist your internal capability building efforts.

For organisations that either want to start experimenting, or experiment well

When it comes to ML, it's not just performing one experiment, proving value and putting it into production and you are done. Like many other technologies, especially emerging ones, success requires continual investment and experimentation.

Steps to get started:

- Pick the right problems to solve with ML using the tips here
- Prioritise problems based on business value vs complexity
- Develop a roadmap of initiatives to inform a strategy or create a business case for investment

Then:

- Define the scope and goals of the first experiment
- Execute the experiment, assess against success criteria and capture learnings
- Determine next steps, whether it's moving the experiment into production or pivoting.

Tips

Don't start with the hardest or most complex problem to solve
Pick a problem that delivers business value, but is also solvable. Demonstrating progress and success early on will deliver learnings and the momentum to continue.

Use design thinking to select and prioritise the most promising ML ideas
The sweet spot for innovation is at the intersection of desirability, feasibility and viability. Applying the same lenses to ML ideas can help assess ones with the most potential.

You can start with a relatively small dataset
It's much better to start with what you have and explore ways of expanding your dataset through data augmentation techniques and using external data sources.

Look at data challenges as an opportunity
By running an ML experiment, you can gain insights into the quality of your data and its challenges and this will help inform or reinforce your data strategy for ML, which may include a refined data collection approach.

To experiment well, think about more than just the tech
Beyond proving the technology, experiments can be designed to validate assumptions that relate to desirability. These can be run early on in parallel to your tech activities.

Showcase and showcase often
Regularly showcasing your progress and learnings can improve visibility, buy-in and support for your initiative. It can also be a good source of feedback.

For organisations that are productionising and scaling models

When moving from PoC to production, there may be many other considerations such as:

- Does the model need to be re-written in the appropriate language and with robustness and design for production and scale?
- How will your model be integrated into an application or feature?
- Does your solution require a human-in-the-loop?
- What is the user experience associated with the application or feature? Is it usable or appealing?
- What automated infrastructure and data pipelines are required to support your model / solution?
- Do you have a framework to monitor, retrain and improve your model to ensure that it continues to perform in a real-world context?

Tips

Form cross-functional teams
We'd suggest a combination of technical, business and experience design roles to maintain a design thinking focus as you build and evolve the product. From the technical side, beyond modelling you may require other supporting roles such as data / cloud / software / Ops engineers.

Don't overlook user research and experience design
Experience design is essential to the success of digital solutions and even more so for ML-powered applications. Introducing ML can radically change the user interaction and therefore customer perceptions and needs should be at the forefront of design.

Experimentation doesn't end once you've launched a product
Continue the experimental mindset to improve your ML models and the overall product experience. This could mean replacing your model or adding new models to improve your results. experience.

Feasibility, viability, and desirability lenses

Desirability - Human lens
Does it address a need?

- Is the pain point or opportunity a real problem? Is it a high-priority problem for the business?
- Is the solution something that will be appealing or trusted by the user?
- Are there any privacy or ethical concerns with the solution?

Feasibility - Technical lens
Can the problem be solved using ML models?

- Do ML models exist that can solve your problem?
- Do you have the appropriate data to train the models and is it accessible?
- Can the solution be built within any time constraints?
- Are there any barriers to integrating or operationalising the model within your systems or applications?

Viability - Business lens
Can it create the desired business impact?

- Can the solution be achieved and maintained within budget constraints?
- Will the solution achieve the desired ROI? Will it deliver the intended cost savings, efficiency gains, increased revenue or market share?

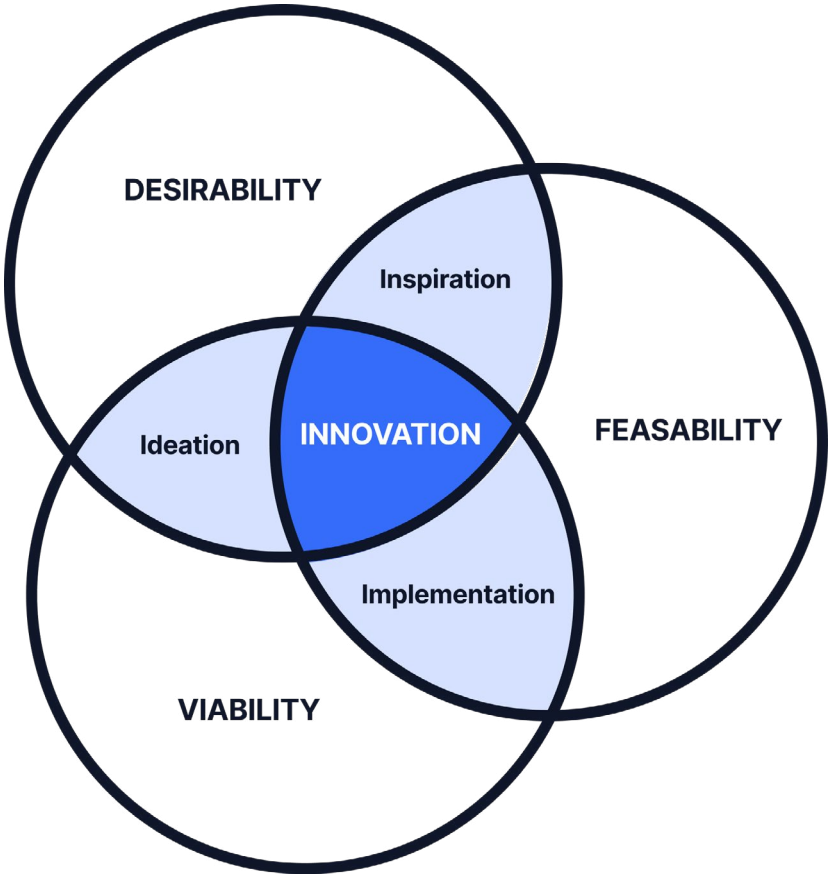

Survey demographics

DiUS conducted a pulse survey to capture insights into how Australian organisations are adopting, using and driving success with Machine Learning. The survey, fielded from December 2020 through February 2021, helped identify the key challenges and priorities for ML projects. The 205 respondents were predominately senior executives and technology practitioners across 18 industries. Notably, 70% of respondents were directly involved in their organisation's ML initiatives.

Industries

Aerospace	3%
Automotive	2%
Communications	3%
Computers & Electronics	5%
Financial Services	13%
Health	10%
Leisure	4%
Life Sciences	2%
Logistics	3%
Manufacturing	4%
Media & Advertising	2%
Mining, Oil & Gas	3%
Professional Services	12%
Public Service	10%
Retail	5%
Software and Internet	13%
Utilities	3%
Other	3%

Number of employees

1 - 20 employees	16%
20 - 100 employees	15%
100 - 250 employees	12%
250 - 500 employees	7%
500 - 1000 employees	11%
1000 - 5,000 employees	21%
5,000 employees or more	18%

Job role

Academic / Researcher	2%
Advisor / Consultant	9%
Business Executive	13%
Developer / Engineer	10%
Entrepreneur (Founder / Co-Founder)	3%
IT Executive	22%
IT Professional or Technical Manager	29%
Sales / Marketing	5%
Solution or Systems Architect	1%
System Administrator	6%

Direct involvement in ML Initiatives

Yes	70%
No	22%

Approach to adopting emerging technology

Innovator	26%
Early adopter	32%
Early majority	25%
Late majority	14%
Laggard	1%
Unsure	1%

DiUS is an Australian technology consultancy that specialises in using emerging technology to solve difficult problems, get new ideas to market or disrupt traditional business models.

With a cross-functional team of 150+ people across Sydney and Melbourne, DiUS provides game-changing approaches to cloud enablement and product development— coupled with expertise in Internet of Things (IoT), big data, artificial intelligence and machine learning (ML).

In 2019 DiUS was the first to attain the AWS ML Competency in ANZ off the back of a world-class delivery process and several successful AI/ML client case studies. At DiUS we build, productionise and scale ML-powered applications that can be deployed to the web/mobile or IoT devices such as wearables, drones and monitoring equipment. We drive innovation and impact through leveraging AWS ML-as-a-Service as well as building custom state-of-the-art ML models.

https://dius.com.au/
info@dius.com.au

Melbourne
Level 3, 31 Queen Street
Melbourne 3000 VIC Australia
P +61 3 9008 5400

Sydney
Level 2, 50 York Street
Sydney 2000 NSW Australia
P +61 2 8014 6640

Scan the QR code to find out more about how DiUS can help you with ML.