THE UNIVERSITY of ADELAIDE

# Annual report 2021

adelaide.edu.au/aiml

## Australian Institute for Machine Learning

make history.

# Contents


About the Institute / **2**

Welcome from the Deputy Vice-Chancellor (Research) / **3**

2021 overview / **4**

AIML capabilities / **7**

**Partnership with the Government of South Australia / 8**

- Partnership programs / **8**
- Key achievements against the partnership objectives / **8**

**Governance / 10**

- Advisory board / **10**
- Project review committee / **10**
- Early career research committee / **10**

**Our research / 11**

- Bridging the 2D / 3D domain gap with democratic AI / **12**
- Machine learning for diagnostic health care / **14**
- AI among the vines / **16**
- Safe, ethical and explainable medical AI / **18**
- Computer vision for space AI / **20**
- Say cheers with an AI beer / **22**
- Award-winning industry collaborations / **23**
- Blockbuster AI behind the perfect Hollywood face / **24**

Centre for Augmented Reasoning / **26**

Art Intelligence Agency / **28**

2021 AIML publications / **31**

# About the Institute

The Australian Institute for Machine Learning (AIML) is a global leader in fundamental and applied research.

Established as a partnership between the University of Adelaide and the Government of South Australia, AIML commenced operations in 2018 as the anchor tenant of the Lot Fourteen innovation precinct.

Now with more than 160 members, AIML has grown rapidly to become one of the largest university-based machine learning groups in Australia and one of the best in the world.

AIML makes an important national and international contribution to pushing the boundaries of what machine learning can do, and how it can be applied to almost every aspect of our lives.

Machine learning underpins the business models of the world's largest corporations, and has the potential to deliver great social, economic and environmental benefits.

At AIML, we collaborate with world-leading companies to develop high-tech products and solutions to everyday problems.

Our industry and business partners include experts in many fields including agriculture, space, medicine, transport, defence, cybersecurity, and the creative sector.

**Our vision**

To be global leaders in machine learning research, and high-impact research translation.

**Our mission**

Research excellence in machine learning (ML), artificial intelligence (AI) and computer vision (CV).



# Welcome

## Welcome from the Deputy Vice-Chancellor (Research)

The University of Adelaide's Australian Institute for Machine Learning (AIML) is one of the world's leading entities for machine learning research and translation.

Now with more than 160 members, AIML has doubled in size since its inception in 2018; this is testament to the rapid growth in the global machine learning sector and the University's strong commitment to South Australia's innovation ecosystem at Lot Fourteen.

Throughout 2021, AIML members delivered impressive achievements in both fundamental and applied machine learning science, with researchers and students producing numerous publications in leading journals and conference proceedings.

AIML works closely with government and industry clients and partners to help them realise the benefits that machine learning presents. Also, to push forward the science of machine learning. Artificial intelligence can solve challenging problems and help people in a variety of ways, and—as you'll read in the articles in this report—we're seeing this impact in a range of sectors such as agriculture, diagnostic healthcare, and even the global entertainment industry.

AIML's trajectory over the last four years has been remarkable, and I'm sure the Institute will continue to build leading science and technology capabilities for Australia into the future.

I hope you enjoy reading AIML's 2021 Annual Report.

**Professor Anton Middelberg**
Deputy Vice-Chancellor
and Vice-President (Research)

# 2021 overview

AIML is a model of the close partnership between the University of Adelaide and the Government of South Australia – it shows the power of universities in leading innovation and building high-tech sectors.

With AIML now in its fourth year, the University of Adelaide has matured to become Australia's largest machine learning research site. Our members work to secure funding from traditional grant agencies as well as industry collaborators, and deliver project outcomes for a diverse range of innovative and ambitious clients and partners.

## Locally based, globally connected

As Lot Fourteen's first and currently largest tenant, AIML is well positioned to attract the smartest people to Adelaide, from across Australia and the world. AIML is a key and active participant in the State Government's *Hi-Tech Sector Plan 2030* to increase innovation and high-tech employment in South Australia.

Despite the disruption of Covid-19 and the associated focus on remote work, the Lot Fourteen precinct is vibrant and alive with opportunities to network, collaborate and try new things. While our machine learning research has a global perspective, our Lot Fourteen 'home' provides the perfect incubation environment to train the next generation of expert AI talent; and support research, innovation and commercialisation endeavours.

The precinct is now home to more than 100 commercial enterprises—around half are established businesses, and half are startups—and this number increases almost every month. More than 1,300 people work or study here, and the seven-hectare project is still in its development phase, with construction of the 16-level, Entrepreneur and Innovation Centre imminent.

AIML's growing AI talent pool is proving a major attractor for multinational technology companies to open operations in Adelaide. In February, tech giant Amazon expanded its Lot Fourteen presence, choosing to establish a machine learning team here in order to access AIML's expertise.

In September, AIML partnered with Microsoft in a memorandum of understanding for *Project AI Off Earth* — a collaboration to conduct modelling and simulation of complex space systems and operations, build AI algorithms for on-board satellite data processing, and develop systems for remote satellite operation and optimisation. It's part of the rapidly growing space industry here in South Australia.

## Our leaders and high achievers

In May, Professor Ian Reid, Head of the School of Computer Science and AIML senior researcher, was named a Fellow of the Australian Academy of Science; a well-deserved recognition of his work over two decades pioneering new research in computer vision and robotics.

In June, Dr Qi Wu was recognised in *InDaily's* '40 under 40' awards, receiving the Adelaide BioMed City Discovery Award for his work as AIML's Director Vision-and-Language Methods.

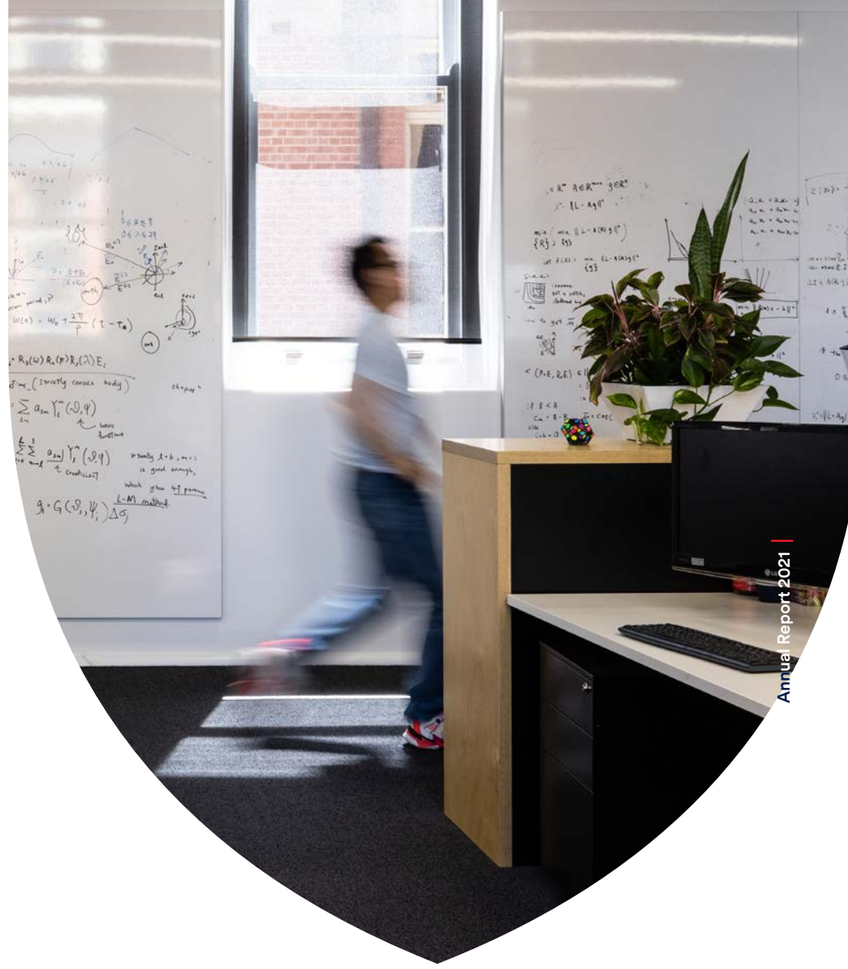



In November, AIML Director Professor Simon Lucey took a co-appointed role at Argo AI, a US-based company that builds software, hardware, maps, and infrastructure for self-driving vehicles. This dual appointment is exemplar of the new international model for the AI sector, where senior researchers hold academic and industry positions at the same time, helping foster collaboration opportunities for researchers and students.

Professor Ian Reid and Professor Anton van den Hengel were also honoured at the 2021 Australasian AI Awards for their significant contributions and service to the field. Professor Reid received the Distinguished Research Contribution award for his landmark research in deep learning for scene understanding, visual tracking, and visual simultaneous localisation and mapping (vSLAM). Professor van den Hengel received the Outstanding Service award for his work establishing AIML and the Centre for Augmented Reasoning, as well as his efforts to see AI shape the cultural landscape through the Art Intelligence Agency.

## Academic excellence

Our machine learning researchers have continued to excel academically in 2021. AIML's members produced approximately 100 publications throughout the year, contributing notable advances in fundamental machine learning research.

AIML researchers and students were authors of 28 papers presented at CVPR '21, the Conference on Computer Vision and Pattern Recognition. Held annually in North America, it's regarded as the world's flagship AI conference, producing the fourth-most impactful scientific papers (according to a ranking by Google Scholar Index). One AIML paper, by PhD student Wei Yin and collaborators at Adobe Research, was shortlisted as a candidate for CVPR's best paper — placing it in the top 2% of all papers accepted by CVPR.

In early December, AIML postdoctoral researcher Dr Michele Sasdelli was awarded the DST Best Contribution to Science Award at DICTA 2021 (the Digital Image Computing: Techniques and Applications conference).

The Institute also appointed a new Future Making Fellow; one of 11 across the University chosen from a field of more than 160 applicants.

## Centre for Augmented Reasoning

The opening in November of the University's new Centre for Augmented Reasoning represents a significant advancement in AIML's research and innovation capacity. The $20 million centre is an investment by the Australian Government in people and research to develop the core capability Australia needs in order to compete globally.

Building on AIML's existing research strengths, the Centre is training a new generation of AI experts working across four main research themes, and will further establish the University as a key strategic institution for machine learning in Australia. Research grants and scholarships supporting PhD students are attracting the best talent to South Australia, while also growing high-tech jobs here at the University and at Lot Fourteen.

Commencing in 2022, the Centre's $3.5 million innovation fund will support local collaboration opportunities, strategic development programs, and new business ventures.

# 2021 overview

## Working with business and industry

A core program of our partnership with the State Government, our engagement with South Australian small-to-medium enterprises has seen businesses develop new products, create new jobs and upskill their staff.

In November, AIML was recognised at the 2021 SA Science and Innovation Excellence Awards winning the category of Excellence in Science and Industry Collaboration. Since 2018 AIML has helped more than 25 organisations realise the benefits of machine learning technology for their product and business capability, with tangible results.

In the case of Adelaide-based visual effects company Rising Sun Pictures, their collaboration with AIML translated into $1 million in increased revenue with an additional $3 million anticipated in 2022.


## Interesting collaborations

AIML members continue to engage and collaborate internationally in new and unexpected projects outside machine learning fields. These initiatives help share the benefit of AIML's technological capability, while also exposing our scientists to new ideas and new ways of thinking.

The Art Intelligence Agency, our partnership with the Sia Furler Institute, has resulted in a significant cultural contribution to the art world. Laurie Anderson, the agency's inaugural artist-in-residence, is the subject of a major solo exhibition at the Smithsonian Institution's Hirshhorn Museum in Washington DC, and also delivered a poetry lecture series for Harvard University; both projects shaped by AIML's custom-built AI technology.

AMIL also worked with Barossa Valley Brewing to create Australia's first AI-designed beer. The two undergraduate computer science students who built the neural network even temporarily transformed a robotic arm into an automated 'bartender'. The story was popular and made its way across international news, stimulating interest from craft beer lovers. Behind this new product lies a very serious message — the AI revolution is coming, and even small enterprises like craft brewers need to develop their AI skills, or risk being left behind.

## 2022 and beyond

The machine learning research environment moves quickly, and so are we.  We're building our network of valued industry clients and partners, and we're working hard to ensure all South Australians can enjoy the benefits that machine learning technology offers.

We appreciate the enthusiastic and continued support of the Government of South Australia, particularly the Department of the Premier and Cabinet, the Department for Industry, Innovation and Science, and the Department for Trade and Investment; as well as our government agency project partners the Department for Education, the Department for Energy and Mining, the Department for Infrastructure and Transport, and the Department of Primary Industries and Regions (PIRSA).

We thank our Advisory Board members, academic leaders, Early and Mid-career Researchers' Committee, and all our hardworking members for their valuable contribution to machine learning science.

**Professor Anton Middelberg**

Deputy Vice-Chancellor
and Vice-President (Research)
Chair of the Advisory Board

**Professor Simon Lucey**

Director, Australian Institute
for Machine Learning



# AIML capabilities

Artificial intelligence (AI) has become increasingly part of our everyday lives.

AI is all around us, embedded in products we now take for granted, and is powering new developments in industry, agriculture, medicine, manufacturing and entertainment.

AIML members regularly win international competitions for excellence in tailoring AI to address specific tasks, often in completely new areas of application and with short lead times for team building and technology development. AI offers us great opportunities to apply new technologies and solve some of the world's intractable problems.

AIML's expertise lies across many areas of applied artificial intelligence:

**Computer vision / robotic vision** — for object counting; extracting information from images and videos, tracking and detecting anomalies; or identifying objects.

**Natural language processing / understanding language** — to extract useful information from documents and datasets; to sort, order and analyse information for patterns; or for chatbots.

**Visual question answering** — where computers are queried with open-ended questions about images. These questions require an understanding of vision, language and common-sense knowledge to answer.

**Advanced learning and reasoning** — to extract meaningful insights from big data, to make AI systems explainable and to push machine learning beyond today's capabilities.

**Signal processing** — to interpret complex signals such as audio and radar, and to review medical images and more.

**Reinforcement learning** — where a machine learning model is trained to make a sequence of decisions. The AI faces a game-like situation and learns through trial and error to find a solution to the problem. This is the type of software that was applied for a computer to beat a human at games such as checkers, go and more recently Fortnite.

**Deep learning** — imitates the working of the human brain in processing data and creating patterns for use in decision making. It is applied in virtual assistants, vision for driverless cars and in face recognition.

**Generative adversarial networks** — involve using machine learning to create a new instance of a particular input (images, audio, video). For example, based on a series of text the machine will generate an image or a piece of art.

**Meta learning and continual learning** — can learn models from multiple tasks and domains and help adapt to new tasks and domains. This is particularly useful when the data for a single task, or some tasks, is small, but data from all tasks as a whole is large. Meta learning can utilise data from all tasks to help learning in some or all tasks. Continual learning can utilise the past tasks (but without keeping all past tasks' data), to learn and adapt to new tasks, without catastrophic forgetting for the past tasks.

**Causality and probabilistic graphical models** — learn and model complex relationships among variables, and naturally offer explainability for the predictions. They can answer interventional questions such as how to change irrigation or pruning practices to improve profit in a farm. They can answer 'what-if' type of counterfactual questions, and generate infinite amounts of counterfactual data to circumvent the small data / sample problem, and improve generalisation performance.

# Partnership with the Government of South Australia

The Australian Institute for Machine Learning (AIML) is a partnership between the University of Adelaide and the Government of South Australia.

AIML was formed as a collaborative partnership between government and academia, for the purpose of driving the state's transition to a more modern and high tech economy. Just four years later, that partnership has paid off. AIML has grown into Australia's leading site for machine learning research and innovation.

AIML has played a key role in attracting global businesses to set up operations here at Lot Fourteen, and South Australia has now built a brand reputation as the place for AI capability.

## Partnership programs

Commencing in 2017, AIML's partnership with the South Australian government comprises four main programs designed to accelerate growth across the state's key strategic areas. AIML researchers are aligned with AIML's own engineering team, giving us an advantage in implementing new AI developments into software solutions for our clients and partners.

### 1) Artificial intelligence skills development

South Australia needs a skilled workforce to support its growing AI and high tech sectors. This program supports undergraduate students make the transition to postgraduate studies, and helps attract the world's best AI research talent to Adelaide.

### 2) Defence industry engagement

AIML collaborates with various defence industry partners. Research and development hours allocated under this program support South Australia's defence workforce, and contribute to Australia's sovereign security capabilities.

### 3) Government efficiency engagement

AIML works with government agencies to identify areas where our AI technology could help improve government operations. Researchers and engineers work on projects that provide solutions for government, and improve service delivery for South Australians.

### 4) SME engagement and global research and development

Since the partnership commenced, AIML's engineering team has worked with 26 local SMEs to help them access AI and machine learning capability to enhance or innovate their product and service offerings.





# Key achievements against the partnership objectives

**Leverage investment of at least $10m from Commonwealth and industry sources**

**$45M**

**In Commonwealth and industry funding, 2018-2021**

**An additional 20 FTE in staff and students within AIML, over the term of the agreement**

**+80**

**Additional AIML members since opening**

**Collaborate with South Australian SME's to develop at least five new products or services**

**13**

**New products and services completed or in development**

**Attraction of at least two international research and development projects to AIML**

**9**

**International R&D projects completed and in progress**

**Delivery of at least three South Australian government AI related projects**

**11**

**Government agencies have worked with AIML to improve their efficiency / service delivery**

**Adoption of AIML technology into at least five other industry projects, products or services**

**26**

**SME projects completed and in progress**

# Governance

## Advisory board

**Professor Anton Middelberg**

Board chair
Deputy Vice-Chancellor and Vice-President (Research)

**Professor Katrina Falkner**

Executive Dean, Faculty of Sciences, Engineering and Technology

**Ms Lusia Guthrie**

Independent business consultant; Chair, Bio Melbourne Network; Chair, Medicines Manufacturing Innovation Centre

**Mr Adam Reid**

Chief Executive Officer, Department for Industry, Innovation and Science

**Dr Tony Lindsay**

Director, STELaRLab, Lockheed Martin Australia

## State government project review committee

**Dr Kathy Nicholson - Chair (May - Dec)**
Operations Manager, AIML

**Dr Paul Dalby - Chair (Jan - April)**
Business Development Manager, AIML

**Dr Andrew Dunbar**
Executive Director, Innovation and Science, Department for Industry, Innovation and Science

**Mr Kim Scott**
Managing Director, TAO Consulting
Chair, Australian Cyber Collaboration Centre

## Early career research committee

**Mr Matthew Howe - Chair (Sept-Dec)**

**Mr Anthony Manchin - Chair (Jan-Aug)**

Dr Rafael Felix Alves

Mr James Paul Bockman

Dr Tong He

Mr Michael Mogford

Mr Mahdi Kazemi Moghaddam

Mr Adrian Orenstein

Ms Violetta Shevchenko

Mr Thomas Martin Walker

Mr Yang Zao

Dr Huangying Zhan

## Academic leadership

Dr Ehsan Abbasnejad

Professor Gustavo Carneiro

Professor Tat-Jun Chin

Dr Feras Dayoub

Associate Professor Anthony Dick

Professor Mark Jenkinson

Dr Lingqiao Liu

Dr Yifan Liu

Professor Lyle Palmer

Professor Ian Reid

Dr Jamie Sherrah

Professor Javen Shi

Dr Jack Valmadre

Professor Anton van den Hengel

Dr Johan Verjans

Dr Qi Wu



# Our research



A type of artificial intelligence, machine learning is the development of computer systems that can learn by example from data.

Machine learning offers great opportunities for analysis, mapping, prediction and understanding in fields where vast data is available — these include vision, language, finance, medicine, science, space, agriculture and more.

Our focus at AIML is to build fundamental new methods and technologies in machine learning.

> “People want the data they’re collecting to be realistic, they don’t want a white background.”
>
> **Professor Simon Lucey**



# Bridging the 2D/3D domain gap with democratic AI

## How AI can understand our big 3D world

**Professor Simon Lucey**

Humans move effortlessly around our rich and detailed three-dimensional world without much second thought. But, like most mammals, our eyes actually sense the world two-dimensionally — it’s our brains that take those 2D images and interpret them into a 3D understanding of reality.

Even without the stereo visual input from our two eyes, we’re experts at looking at a flat 2D image and instantly ‘lifting’ it back to its 3D origins, we do it every time we watch TV or look at photos on our phone.

But computers and robots have a much harder time doing this ‘lifting’. It’s a problem that AI researchers are working to fix.

Making computers able to understand 3D space from only 2D input is considered such an important capability—with diverse applications ranging from mobile phones to driverless vehicles—that AIML Director Professor Simon Lucey is working to build a geometric reasoning system that can exhibit human-like performance.

“When cameras try to sense the world, like humans do, what’s coming into the robot is still just 2D. It’s missing that component that we have in our brains that can lift it out to 3D, that’s what we’re trying to give it,” Lucey says.

If 3D understanding from normal cameras is so difficult, why not instead equip computer vision systems with proper 3D sensors like LiDAR, a sensing method that uses lasers? It’s not that easy. Building new hardware is slow and expensive, and often out of reach for smaller tech startups seeking to innovate AI research commercially.

“You could take ten years and billions of dollars and it would still be very, very risky to generate...but when you’re doing something in software, you can deploy it straight away, and you can continually update and make it better,” Lucey explains.

Building computer vision systems that can understand the real world typically requires vast troves of labelled training data using something called supervised machine learning. Millions of images, each labelled ‘dog’, ‘strawberry’ or ‘President Obama’; or thousands of hours of driving footage where coloured boxes are drawn to mark each pedestrian, stop sign and traffic light.

AI researchers are using the vast collections of labelled 2D training data, and working out how to apply it so AI systems can develop a 3D geometric understanding similar to that of humans.

“How can I take 2D supervision that humans can easily provide,” asks Professor Lucey, “and, using some elegant math, allow it to act as 3D supervision for modern AI systems?”

One application of this kind of computer vision is something called 3D motion capture, where earlier advances brought us Gollum in The Lord of the Rings movies. It’s still a popular technique and one that’s widely used in film visual effects, video game production and even medicine and sports science. But even today it uses a number of expensive and finely calibrated cameras, and sometimes still requires people to wear special reflective dots on their body and perform in front of a greenscreen, and that’s a problem.

“People want the data they’re collecting to be realistic, they don’t want a white background. They don’t want a greenscreen. They would love to be out in the field, or in areas that are highly unconstrained. And the sheer cost of this limits the application of this technology at the moment,” says Professor Lucey.

But in a recent project that saw Professor Lucey work with researchers from Apple and Carnegie Mellon University, the team was able to demonstrate a new AI method for 3D motion capture that is sure to make the technology far more accessible and affordable.

“The work we’ve done on this paper has tried to ask the question: how few cameras could we get away with if we were willing to use AI to do this 3D lifting trick?”

The team used something called a neural prior — a mathematical way of giving an AI system an initial set of beliefs in terms of probability distribution, before any real data is provided.

As a result, the new method can perform 3D motion capture from normal video footage (no greenscreens or special reflective dots required) using only two or three uncalibrated camera views. It delivers similar 3D reconstruction accuracy that would otherwise require as many as 40-60 cameras using earlier methods.

Professor Lucey highlights the importance of AI research that focuses on finding efficiencies and significant cost breakthroughs as a way of bringing technology to those who’d otherwise not be able to afford it.

“You could be a small startup and you could use this, whereas with other methods you’d need to be very well resourced financially,” he says.

“It’s democratic AI.”

**Story:** Eddie Major.

# Machine learning for diagnostic health care

Dr Zhibin Liao

Have you ever wondered what happens after you get a pathology test? If you're in South Australia, there's a good chance your sample will end up with SA Pathology, where teams of pathologists and scientists work around the clock, every day of the year, to provide comprehensive pathology services for the state's public health sector.

AIML is working with SA Pathology to research ways of using artificial intelligence (AI) in diagnostic health care. While still in its early stages, SA Pathology hopes the adoption of machine learning can help pathologists by speeding up some of their more time-consuming tasks, and potentially even help doctors improve outcomes for their patients and prevent avoidable hospital admissions.

Machine learning researcher Dr Zhibin Liao works in medical image analysis. In one SA Pathology project, Dr Liao and colleagues are working on a computer vision method to detect and count cells and give more information to pathologists.

"Cell counting is a pathology method to measure the number of cells of interest in a sample, but it's very time consuming," Dr Liao says. "So, if a machine learning algorithm could accurately count them faster than a human, that would save a lot of time."

Ki-67 is a protein that is found in the nucleus of cells that are actively growing and dividing, but not in resting cells, making it an excellent marker to determine active cell proliferation. For some cancers, Ki-67 is an important indicator of disease prognosis and how well it will respond to chemotherapy treatment.

Pathologists are able to identify the number of cells with Ki-67 in a given sample using immunostaining, a method of identifying specific proteins using chemical dyes. When viewed under a microscope, cells with a high level of Ki-67 appear dark brown, and those without appear blue.

While automated cell counters have long been available for many pathology tests, the Ki-67 test requires pathologists to count each cell manually; a time-consuming process.

Dr Liao is developing an AI system that can detect and count Ki-67 positive and negative cells from microscopy images using *unsupervised learning* — a type of algorithm that learns patterns from untagged data. The system detects and outlines each type of cell and annotates the images for a pathologist to then review and correct any errors, saving the pathologist time in annotating a large number of cells. The corrected images are then used to train machine learning models and improve the performance.

Professor Gelareh Farshid is a senior consultant pathologist at the Royal Adelaide Hospital and specialises in breast cancer and soft tissue pathology.

"The capacity of AI algorithms to provide objective, reproducible, accurate and fast Ki-67 data has enormous clinical appeal. The Ki-67 proliferative index is one of the factors that assists in treatment planning of some patients with breast cancer, including whether chemotherapy should be considered and the type of medications that may be effective," Professor Farshid says.

The Ki-67 cell counting project is at an early research stage, and is not being used clinically for patient tests, but SA Pathology is embracing the benefits that machine learning automation will bring.

SA Pathology Executive Director Lucas Semmler, says AI is going to "become part of the everyday, normal workload and pathology work practices in the future".

"There will always be a scientist or a pathologist involved in every pathology test. AI will provide significant support to a certain level, but there are definite benefits due to the significant volumes of data generated within pathology," Semmler says.

**Story:** Kurtis Eichler and Eddie Major.

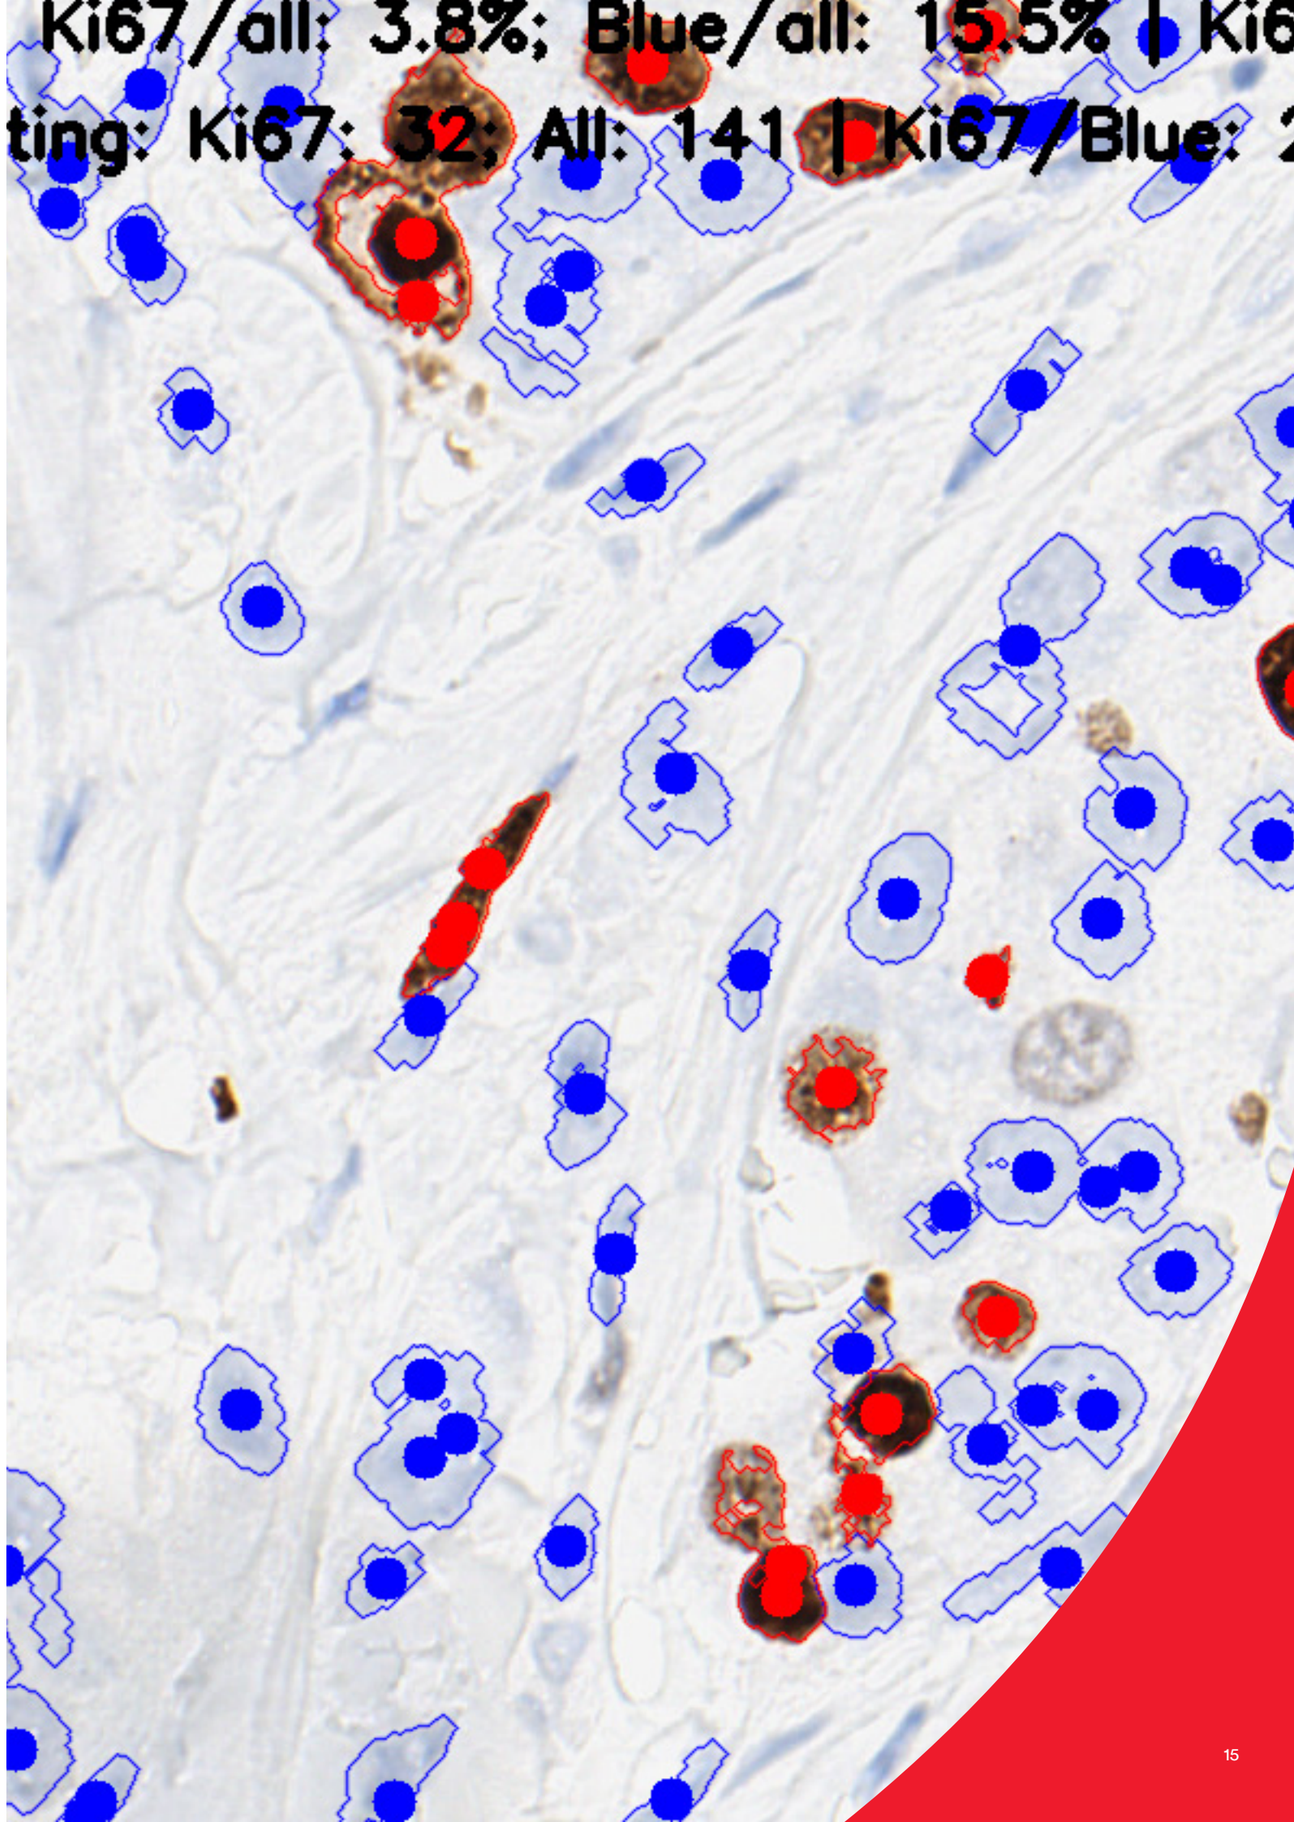

> **“The cameras are mimicking human eyes. Our AI behind it, is mimicking the viticulturalist’s expert brain.”**
>
> **Professor Javen Shi**



# AI among the vines

**Professor Javen Shi**

South Australia’s picturesque Riverland is well known as the home of Australia’s largest wine region, and it’s also now the location of a high-tech project that’s set to see the country’s $40 billion viticulture industry enjoy the benefits of artificial intelligence (AI).

Throughout a typical season, wine growers are faced with a number of decisions in managing their vineyards; such as when to prune the vine canopies, apply fertiliser and how best to control irrigation. To do this, viticulturalists are increasingly choosing high tech solutions — using sensors, probes and cameras to keep an eye on their grapes.

While camera systems for agriculture aren’t new, a team of University of Adelaide researchers—working in collaboration with Wine Australia and Riverland Wine—has developed an advanced agricultural computer vision system tailor-made for Australian grape growers.

*VitiBox* is a small unit that’s mounted on a quad bike or tractor, and driven around the vineyard while digital cameras continually take images of the vines. Once captured, the images are then fed into the *VitiVisor* system: an information, prediction and advisory platform for Australian winegrowers.

*Vitivisor’s* powerful algorithms then convert the image data into a range of key indices for managing the vines and predicting the crop’s outcomes.

This valuable insight was achieved by training the AI algorithms on a library of 175,000 images of several grape varieties grown in South Australia’s Riverland wine region. Experienced viticulturalists annotated the images to provide the AI system with a ground truth on which to make its predictions.

AIML’s Director of Advanced Reasoning and Learning, Professor Qinfeng (Javen) Shi, who is co-leading the project, calls it a machine brain.

“The cameras are mimicking human eyes. Our AI behind it, is mimicking the viticulturalist’s expert brain,” Professor Shi says.

“The artificial intelligence behind this new system can detect more than just canopy development. It can estimate many other key indices as well, such as pruning weight and the number of buds or shoots on the vine.”

From the camera vision, the system can detect inflorescences (grape flower clusters) and give a prediction of the harvest yield.

Like all aspects of the $3 million *VitiVisor* project, the *VitiBox* was designed with Australia’s wine growers in mind. The system is easily configurable, with options for multiple cameras, thermal imaging and even LiDAR sensors.

The ground-based vision system will be available from the open source platforms Hackster and Github in 2022.

**Story:** Kurtis Eichler and Eddie Major.

> “A big problem in medical practice is that these biases, and these disparities, are invisible.”
>
> **Dr Lauren Oakden-Rayner**

# Safe, ethical and explainable medical AI





**Dr Lauren Oakden-Rayner**

Millions of Australians undergo medical imaging every year to diagnose a range of potential health problems. Radiologists and other doctors use this technology to obtain detailed information and help plot their patient’s treatment and provide a better level of care.

Medical imaging systems—such as those used in x-rays, MRIs, CT scans, and sonography—are increasingly incorporated with artificial intelligence (AI) technology to assist doctors in making an accurate and timely diagnosis.

Supplied with large training datasets of medical images, AI software models are able to learn how to analyse medical scans for signs of disease or injury, with similar or greater accuracy than that of human doctors.

The technology is well established in clinical practice, with almost 350 AI systems already FDA-cleared for use in American hospitals and radiology practices to do things like detect bone fractures, measure heart blood flow, and even diagnose strokes and certain types of cancer.

However, when using AI in a high-stakes field such as medicine, doctors have to be able to understand the intricacies of the software’s decision-making process when considering a diagnosis.

There is a lot that medical AI researchers still don’t know, and they’re working to ensure that AI tools are safe, ethical and explainable. One of those researchers is Dr Lauren Oakden-Rayner, a practising clinical radiologist and Senior Research Fellow in Medical AI at AIML.

“I think there’s a lot of evidence that our current assessment methods don’t give us a good indication of how AI imaging systems will perform when they get out into clinics,” Dr Oakden-Rayner says.

Dr Oakden-Rayner, alongside a world-leading team of researchers from MIT, The University of Toronto, Imperial College London, and the University of Birmingham, authored a viewpoint article for Lancet Digital Health. The researchers advocate for a rigorous internal and external validation of AI models in the form of a medical algorithmic audit process to help doctors better understand how automated systems make decisions, and more importantly, how harmful they can be if they make mistakes.

This framework would help users, developers and regulators consider these potential errors in the diagnostic systems and anticipate consequences for patients.

The issue of ethics and justice in medical AI made international news in August, when Dr Oakden-Rayner and colleagues demonstrated an AI model that can accurately detect a patient’s race from x-rays, even when the images are blurred or distorted — something even experienced doctors can’t do.

But exactly *how* the software is doing it, has left researchers stumped.

While a patient’s race or ethnicity is not considered clinically important when radiologists review x-rays, the study highlights the danger that medical AI systems may reinforce, or even amplify, existing racial biases and disparities in public health care.

“We understand that medical practice is biased and produces health disparities across racial groups; and we understand that machine learning models will reproduce biases that they’re trained on,” she says.

This, Dr Oakden-Rayner says, offers an opportunity for the medical community to intervene and help address some of those disparities.

“A big problem in medical practice is that these biases, and these disparities, are invisible.”

“But when you develop a machine learning model, you have to make some explicit decisions about how you deal with patients, and in this case, patients from different racial groups.”

“When you do that, you have the potential to intervene in a positive way”

Dr Oakden-Rayner says that, if implemented safely, AI systems offer many potential medical benefits to diverse communities.

“In a general sense, we have the promise of better quality medical care, we have the promise of better diagnostics, and more accurate decision-making. With that, automation means we can reduce costs as well, which allows us to expand access and coverage.”

**Story:** Kurtis Eichler and Eddie Major.

> “Ideally we want to design unmanned spacecraft to refuel satellites when they run out of power.”
>
> **Sofia McLeod**

# Computer vision for space AI

**Sofia McLeod**

Sofia McLeod spends a great deal of time thinking about exploring space. The AIML PhD candidate is researching ways to build an AI system that can safely land an autonomous spacecraft on a distant planetary or asteroid surface guided by visual input from a single event camera.

Inspired by how the human eye works, an event camera is a dynamic vision sensor where each pixel works independently to report changes in brightness as they occur. Whereas a typical camera sensor—like the one in your smartphone—records numerous whole image frames every second, even if there’s nothing new happening in view.

Because it only sends new data when conditions change, an event camera system is more data efficient, lightweight, and may use less power — all things that are vital in successful space missions, where resources are precious and limited.

Space exploration is an area currently seeing rapid development in autonomous technology. In April 2021, NASA completed the first powered controlled extraterrestrial aircraft flight as part of its Mars 2020 rover mission. The long distance from Earth meant the craft had to operate with a high degree of autonomy.

“If you’re looking at Ingenuity, which is the drone that’s now on Mars, the delay is approximately 20 minutes for human interaction,” McLeod says, referring to the maximum return radio signal transmission time due to the 50 - 200 million kilometre distance between Earth and Mars.

“We need both the rover and Ingenuity to be able to do this navigation by themselves, so they can avoid obstacles on their own. You have to remember that if a robot gets stuck or breaks down, we can’t go to Mars to repair it.”

Event-based computer vision technology in space isn’t just limited to autonomous landing, but has a range of potential applications.

“Ideally we want to design unmanned spacecraft to refuel satellites when they run out of power. To do this you’ll need computer vision to know that you’re aligned perfectly with the object you’re trying to dock with,” she says.

And like a future autonomous spacecraft’s camera-guided journey, McLeod’s pathway to computer science and machine learning at AIML was a visual one; she initially considered a career in design or visual effects for the entertainment industry.

“I’ve always been a visual person,” she says. “I was thinking of doing VFX or graphic design... but I was definitely better at algorithms.”

“I was just really fascinated by the concept of getting computers to see,” she says. “It’s just so intuitive.”

**Story:** Kurtis Eichler and Eddie Major.

# Say cheers with an AI beer

**SME Engagement**

For craft brewer Denham D'Silva, creating the perfect beer comes down a lot to old-fashioned trial and error. So the idea of introducing artificial intelligence into the time-worn production didn't fill the Barossa Valley Brewing founder with hope.

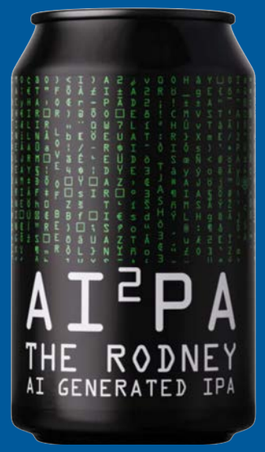


"The willingness to experiment and create interesting and premium beers has been a foundation of the brewery for 16 years," he says.

"So, to largely place this process in the hands of AI, was in a word, terrifying."

The unique idea came about while D'Silva and AIML founder Professor Anton van den Hengel were sharing a few cold ones.

The duo brought their two passions—technology and beer—together to create an India pale ale designed entirely by AI.

To achieve this, two University of Adelaide computer science interns, Christopher Fusco and Jay Vira, were set to work building a dataset of 260,000 existing beer recipes. From this, the students developed a neural network that learned how to brew beer by examining the dataset.

"We had to come up with our own mathematical formula using statistics from those original recipes," Fusco explains.

"By getting statistics on these variables, we were able to judge the significance of each variable, this also helped us deal with any possible biases that could have occurred in the data," Vira adds.

The project didn't remove the role of the master brewer entirely. The AI designed around 30 potential recipe candidates, but Barossa Valley Brewing made a final (human) decision on which would be brewed.

The beer was named The Rodney AI²PA, in homage to Australian roboticist Professor Rodney Brooks, inventor of the Roomba vacuum cleaner.

The beer is available in pubs and independent bottle shops across South Australia. It's also going global, with second-generation AI-enabled recipes now being brewed in the US.

As for D'Silva, he now sees AI as being vital in refining flavours and developing new products.

"Beer is traditionally a very hands-on process, and even more so for a small craft brewery like Barossa Valley," D'Silva says. "When you're a smaller craft brewery you can't compete on scale, so you have to be different and clever."

"Most people think of AI and machine learning as something that only the huge tech companies can do. This project has shown us that AI can take our artisanal skills and augment them to allow us to compete."

**Story:** Kurtis Eichler and Eddie Major.



# Award winning industry collaborations

**SME Engagement**

AIML Director Professor Simon Lucey thinks South Australia is batting above its average when it comes to innovation in artificial intelligence, and he wants to help businesses realise the technology's potential.

"Our mission here at AIML is to show that these sorts of innovations can be applied not just for international businesses, but for local businesses and even government; and to really showcase the excellence we have here in South Australia," Professor Lucey says.

"AI is really a productivity tool. If you want to get your business to be more productive and compete with people internationally, you can really leverage AI.

"What AI offers is a chance to complexify our economy and really give our kids the opportunity to have a really diverse set of jobs."

In December, AIML was recognised at the 2021 SA Science and Innovation Excellence Awards, winning the category of Excellence in Science and Industry Collaboration.

The award is the result of four years of hard work by machine learning experts and business development specialists.

AIML's industry solutions team worked with small-to-medium enterprises to help them get the most out of their business data. Since 2018, AIML has helped 21 businesses weave machine learning into their practices and supported them in the development of 13 new AI-enabled products.

Ultimately, this has created new jobs and provided existing staff with new skills.

One of those businesses is Pickstar, an Adelaide-based startup that allows customers to book sportspeople and athletes as brand ambassadors and guest speakers for events.

Pickstar CEO and founder James Begley recognises the importance of analytics in delivering successful bookings, with the company expanding internationally over the duration of its partnership with AIML.

"We've grown from a staff of nine to 35+ across AU, UK and USA. We are now exporting our technology to clients like the National Football League, and English Premier League clubs," Begley says.

"All our technology development resides here in Adelaide. We will be building our in-house data science team as we mature this new AI and machine learning capability."

**Story:** Kurtis Eichler and Eddie Major.

# Blockbuster AI behind the perfect Hollywood face

## SME Engagement

Last year, millions of people across the globe experienced the very latest in South Australian AI technology in an industry worth more than $100 billion. And if it all went according to plan, they probably didn't even notice it at all.

AIML researchers teamed up with Adelaide-based visual effects company Rising Sun Pictures (RSP) to use AI to create effects for some of Hollywood's biggest blockbusters, including Marvel Studios' *Shang-Chi and the Legend of the Ten Rings*.

This involved creating an AI method of replacing the faces of stunt performers in combat scenes with those of the lead actors.

Computer vision researchers Dr Ben Ward and Dr John Bastian worked on the film with RSP, before joining the VFX studio as full time employees in October.

"Rather than the traditional 2D and 3D face replacements typically used in high-intensity action scenes, the team used an AI deepfake method," Dr Ward says.

Deepfakes are synthetic media where a person in an existing image is convincingly replaced (or faked) with the likeness of someone else, using deep learning — a type of AI that learns from data and uses multiple software layers inspired by our brain's own network of neurons.

For *Shang-Chi*, this involved around 30,000 facial images across five characters, training five machine models in more than four million training iterations. The models were used for around 50 face replacements in six key action scenes.

"AI can help artists accomplish incredible artistic effects without the tedium of executing every frame in a long sequence," Dr Bastian says.

Dr Ward and Dr Bastian are currently using their facial replacement method on three new film projects.

Tony Clark, RSP's managing director, says AI can speed up delivery of visual effects and relieve artists of tedious, time consuming work.

"AI holds great promise for visual effects applications, especially in terms of accelerating labour-intensive tasks and augmenting human creativity," Clark says.

Our early work with AI has produced spectacular results and we are eager to push development further."

While AI might help bring the creative magic of Hollywood to life, the reality is show business is exactly that, business. For RSP, Clark says the collaboration with AIML helped the company generate $1 million in increased revenue with an additional $3 million forecast in 2022.

"Collaborating with AIML has enabled us to continue to deliver amazing images, and reinforces to global studio executives that RSP is among the best in the world at embracing and implementing advancements in technology such as AI," Clark says.

**Story:** Kurtis Eichler and Eddie Major

# Centre for Augmented Reasoning

Artificial intelligence is transforming every industry and sector of the economy. We're building the expertise Australia needs in order to transition to the coming AI-enabled global economy.

Established in November 2021, the new Centre for Augmented Reasoning (CAR) is a $20 million investment by the Department of Education, Skills and Employment in people and research. Sitting within AIML and building on the institute's existing research strengths, CAR is training the next generation of machine learning experts and helping AIML be a leading voice in Australia's AI landscape.

## What is augmented reasoning?

Augmented reasoning is a new and emerging field of AI that combines an advanced ability for computers to recognise patterns using traditional machine learning, with the ability for them to apply reason and learn from prior information and by interacting with humans.

Rather than teaching humans to learn how to 'talk' in computer language, augmented reasoning can help us make computers better at understanding people and our needs, through more natural conversation and interaction. Augmented reasoning gives AI an ability to solve some of the frustrations and problems that we all experience with current computers and technology today.

## Investing in people

CAR has created new high-tech jobs at the University and in the heart of South Australia's growing innovation precinct, Lot Fourteen.

New researchers started coming onboard at the end of 2021, and the Centre will soon be home to:

- 22 postdoctoral researchers working across four machine learning research themes
- 8 scholarship-supported PhD students

## Investing in innovation

Starting in 2022, a $3.5 million innovation fund for machine learning investment will support local commercialisation activities and new business ventures. The fund aims to deliver at least 10 new commercial opportunities coming from the centre's research over the life of the program, these opportunities include:

- seed funding to launch new start-ups involving the University's researchers and graduates
- collaborations with external companies to co-develop new AI-enabled products and capabilities
- projects to extend the impact and reach of our research.

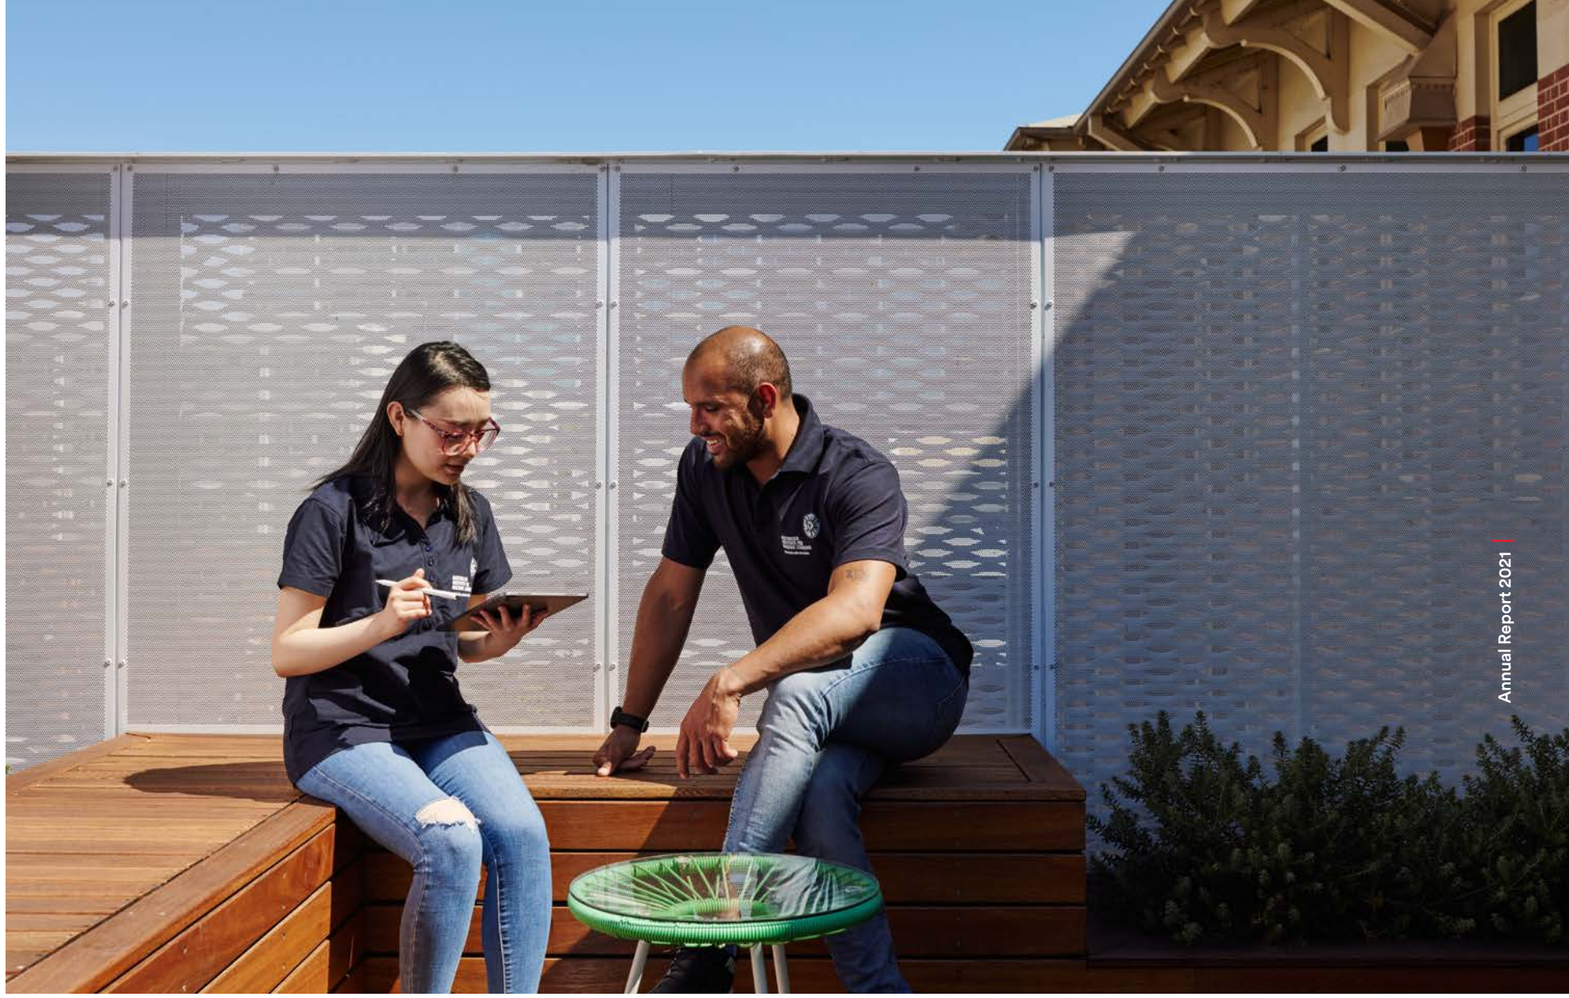



## AI and the Australian community

Thanks to the Centre's community engagement program, AIML is now even more engaged with a broad range of stakeholders across state and federal government, industry, business and the broader community. We're working with journalists, writers, artists and musicians, to explore what AI can mean for all Australians and how it will impact and improve our lives.

Right: Former Senator Rex Patrick; Centre for Augmented Reasoning Director, Professor Anton van den Hengel; Member for Dunstan, the Hon. Steven Marshall; University of Adelaide Vice-Chancellor and President, Professor Peter Høj.

# Art Intelligence

An ongoing creative collaboration between AIML and the Sia Furler Institute that combines the latest AI technology with leading artists from around the world.

The Art Intelligence Agency proves that when you combine the world's most creatively daring artists with AIML's machine learning science, fascinating things can happen. We're bringing together the best people in art, tech and business to collaborate on the frontier of artificial intelligence (AI).

## Artist-in-residence

This year saw the delivery and exhibition of various projects created by the agency's inaugural artist-in-residence, Laurie Anderson.

Throughout 2021 Anderson presented a series of live streamed lectures as the recipient of the 2021 Charles Eliot Norton Professorship in Poetry at Harvard University's Mahindra Center for the Humanities. Unlike previous poetry readings at Harvard, Anderson's six part series, titled *Spending the War Without You*, was an avant-garde performance made for the Covid-19 Zoom video conferencing times; a floating head against a virtual background, with visual effects subtly augmented by AIML's custom computer vision technology.

In September, the Smithsonian Institution's Hirshhorn Museum presented Anderson's largest ever US exhibition. *Laurie Anderson: The Weather* showcases key artworks from across her half-century career, and reveals new pieces created during her agency residency.

One such artwork, Scroll, 2021, pictured left, is an AI-generated Bible written in Anderson's 'voice', produced as the result of AIML's deep learning neural network being supplied with English translations of religious texts as well as Anderson's own extensive written archive.

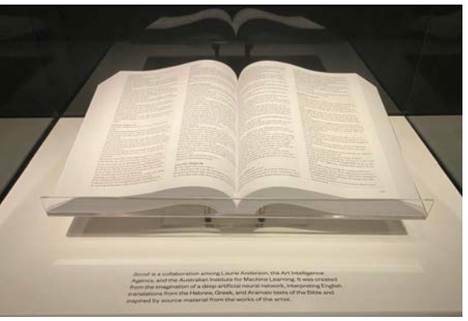


Given a question, AIML's AI text engine can also write in imitation of Anderson's idiolectic style, as well as that of her late husband, musician Lou Reed. Speaking to The New York Times, Anderson admitted the experimental AI text generator isn't perfect—at least two thirds of the output is either nonsense or just boring—but is sometimes deeply moving.

*'She did a few more Lou poems, including one in which he spoke in the most intimate terms: "my eyes are thin and dry, my heart is beating very fast."*

*"Wonderful," she said. "Just great. He's talking to me from somewhere else. I definitely do feel that. The line is pretty thin for me."*

— Laurie Anderson Has a Message for Us Humans, The New York Times Magazine, 6 October 2021.

*Laurie Anderson: The Weather* runs from 24 September 2021 to 7 August 2022 at the Hirshhorn Museum and Sculpture Gallery, in Washington D.C. *Spending the War Without You* is available to watch on YouTube.

## Hackathon 2021

In July, Art Intelligence Agency's second Hackathon saw AIML's machine learning experts explore new creative frontiers alongside contemporary artists working in a variety of visual, digital, and performance art realms.

The goal of the annual day-long event is to spark new startup ideas that can germinate and thrive in South Australia's creative economy. Hackathon is all about high-speed collaboration and collision: participating artists, engineers, sociologists, and scientists are all propelled into new and unfamiliar territory.

Hackathon 2021 resulted in several high-level concepts, with at least one major project now confirmed for delivery in 2022.

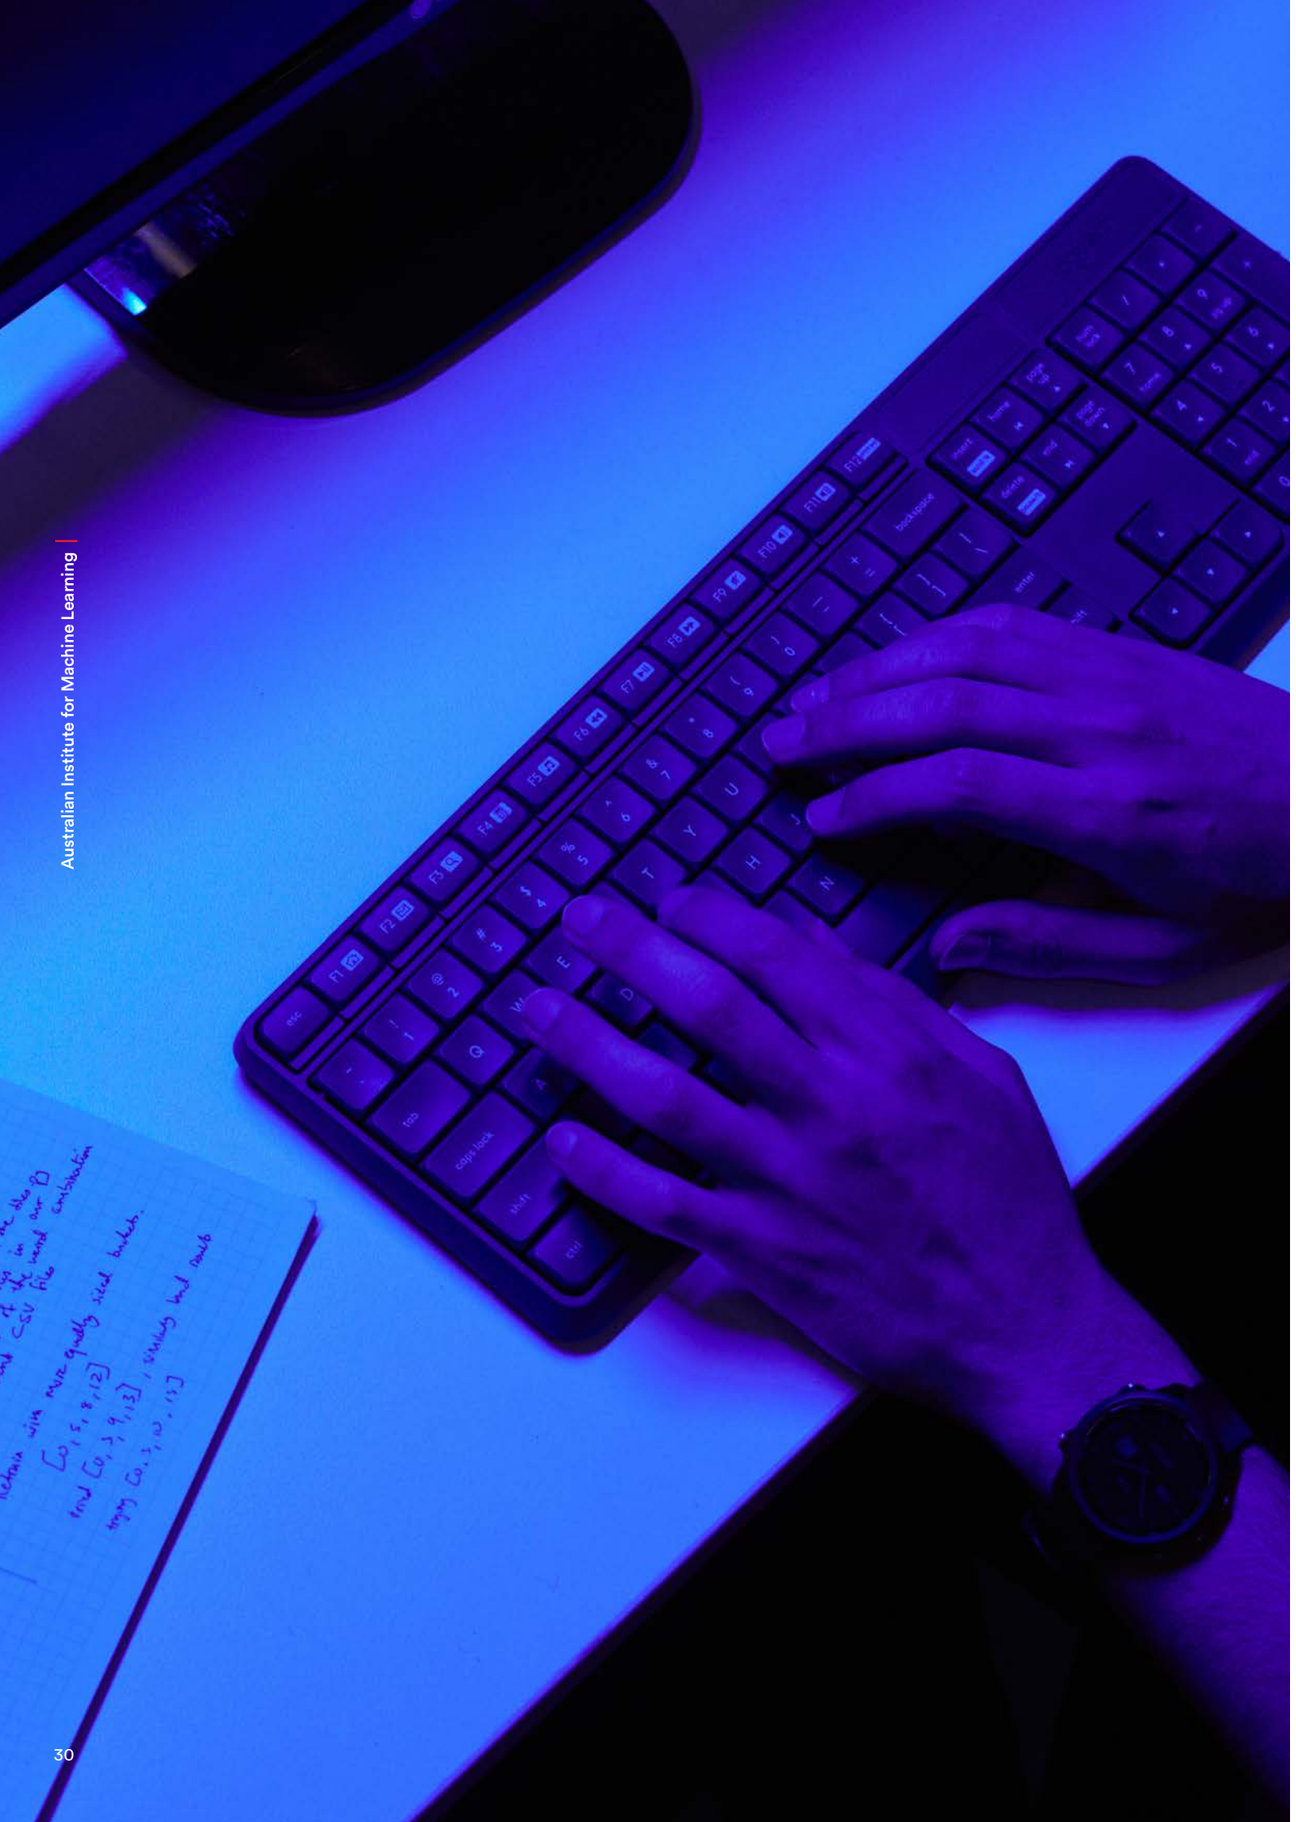

# 2021 AIML publications

Throughout 2021 AIML members authored more than 100 papers in leading international journals and conference proceedings, contributing notable advances in fundamental machine learning research.

A full list is published on the AIML website: adelaide.edu.au/aiml/publications

AIML researchers and students were the authors of 28 papers accepted at CVPR 2021 — the Conference on Computer Vision and Pattern Recognition. Held annually in North America, CVPR is regarded as the world's flagship AI conference, producing the fourth-most impactful scientific papers; according to a ranking by Google Scholar Index.

CVPR 2021 received more than 7,000 submissions and accepted around 1,660 — an acceptance rate just under 24%. One AIML paper, by PhD student Wei Yin and collaborators at Adobe Research, was among 32 shortlisted as candidates for CVPR's best paper, placing it in the top 2% of all accepted papers. Dr Qi Wu's Vision and Language Methods research group presented six CVPR papers.

## CVPR 2021 publications

1. A´lvaro Parra, Shin-Fang Chng, Tat-Jun Chin, Anders Eriksson, Ian Reid. Rotation Coordinate Descent for Fast Globally Optimal Rotation Averaging: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
2. Chaorui Deng, Shizhe Chen, Da Chen, Yuan He, Qi Wu. Sketch, Ground, and Refine: Top-Down Dense Video Captioning: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
3. Chaoyang Wang, Simon Lucey. PAUL: Procrustean Autoencoder for Unsupervised Lifting: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
4. Chen Gao, Jinyu Chen, Si Liu, Luting Wang, Qiong Zhang, Qi Wu. Room-and-Object Aware Knowledge Reasoning for Remote Embodied Referring Expression: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
5. Daqi Liu, Alvaro Parra, Tat-Jun Chin. Spatiotemporal Registration for Event-based Visual Odometry: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
6. Delian Ruan, Yan Yan, Shenqi Lai, Zhenhua Chai, Chunhua Shen, Hanzi Wang. Feature Decomposition and Reconstruction Learning for Effective Facial Expression Recognition: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*

# 2021 AIML publications


7. Dongyan Guo, Yanyan Shao, Ying Cui, Zhenhua Wang, Liyan Zhang, Chunhua Shen. Graph Attention Tracking: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
8. George Cazenavette, Calvin Murdock, Simon Lucey. Architectural Adversarial Robustness: The Case for Deep Pursuit: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
9. Haoyang Zhang, Ying Wang, Feras Dayoub, Niko S. VarifocalNet: An IoU-aware Dense Object Detector: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
10. Islam Nassar, Samitha Herath, Ehsan Abbasnejad, Wray Buntine, Gholamreza Haffari. All Labels Are Not Created Equal: Enhancing Semi-supervision via Label Grouping and Co-training: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
11. Guanghui Xu, Shuaicheng Niu, Mingkui Tan, Yucheng Luo, Qing Du, Qi Wu. Towards Accurate Text-based Image Captioning with Content Diversity Exploration: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
12. Jianpeng Zhang, Yutong Xie, Yong Xia, and Chunhua Shen. DoDNet: Learning to Segment Multi-Organ and Tumors from Multiple Partially Labeled Datasets. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
13. Peng Chen, Jing Liu, Bohan Zhuang, Mingkui Tan, Chunhua Shen. AQD: Towards Accurate Quantized Object Detection: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
14. Ruibo Li, Guosheng Lin, Tong He, Fayao Liu, Chunhua Shen. HCRF-Flow: Scene Flow from Point Clouds with Continuous High-order CRFs and Position-aware Flow Embedding: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
15. Ruwan Tennakoon, David Suter, Erchuan Zhang, Tat-Jun Chin and Alireza Bab-Hadiashar. Consensus Maximisation Using Influences of Monotone Boolean Functions: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
16. Tong He, Chunhua Shen, Anton van den Hengel. DyCo3D: Robust Instance Segmentation of 3D Point Clouds through Dynamic Convolution: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
17. Weian Mao, Zhi Tian, Xinlong Wang, Chunhua Shen. FCPose: Fully Convolutional Multi-Person Pose Estimation with Dynamic Instance-Aware Convolutions: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
18. Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, Chunhua Shen. Learning to Recover 3D Scene Shape from a Single Image: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
19. Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, Lei Li. Dense Contrastive Learning for Self-Supervised Visual Pre-Training: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
20. Xueqian Li, Jhony Kaesemodel Pontes, Simon Lucey. PointNetLK Revisited: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
21. Yao Lu, S¨oren Pirk, Jan Dlabal, Anthony Brohan, Ankita Pasad, Zhao Chen, Vincent Casser, Anelia Angelova, Ariel Gordon. Taskology: Utilizing Task Relations at Scale: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
22. Yazhou Yao, Zeren Sun, Chuanyi Zhang, Fumin Shen, Qi Wu, Jian Zhang, Zhenmin Tang Jo-SRC. A Contrastive Approach for Combating Noisy Labels: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
23. Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez-Opazo, Stephen Gould. VLNoeBERT: A Recurrent Vision-and-Language BERT for Navigation: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
24. Yifan Liu, Hao Chen, Yu Chen, Wei Yin, Chunhua Shen. Generic Perceptual Loss for Modeling Structured Output Dependencies: I*EEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*


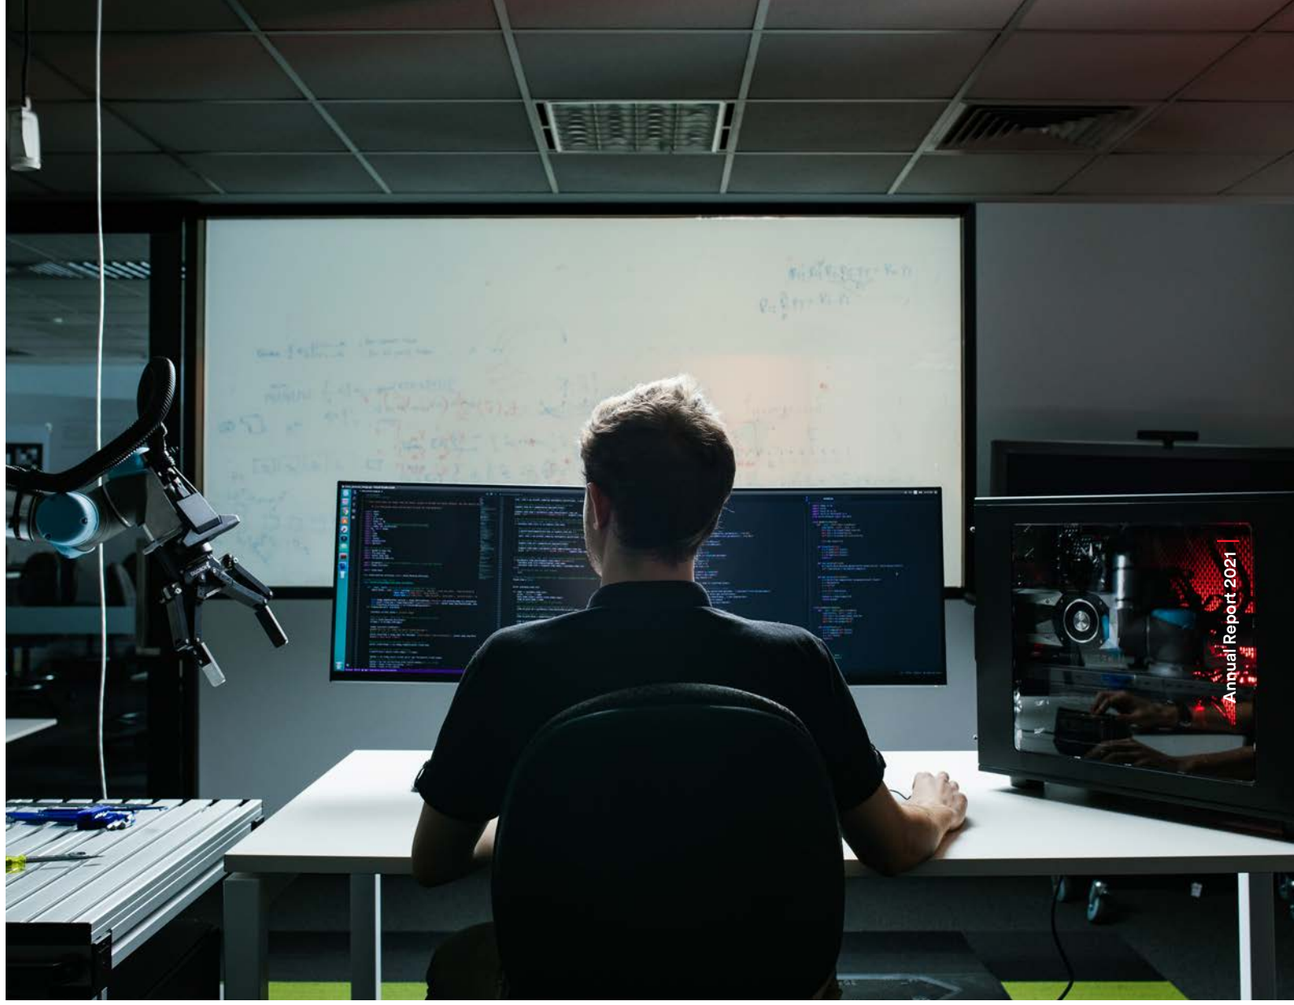




25. Ying Shu, Yan Yan, Si Chen, Jing-Hao Xue, Chunhua Shen, Hanzi Wang. Learning Spatial-Semantic Relationship for Facial Attribute Recognition with Limited Labeled Data: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
26. Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, Huaxia Xia. Meituan End-to-End Video Instance Segmentation with Transformers: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
27. Yutong Dai, Hao Lu, Chunhua Shen. Learning Affinity-Aware Upsampling for Deep Image Matting: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*
28. Zhi Tian, Chunhua Shen, Xinlong Wang, Hao Chen BoxInst: High-Performance Instance Segmentation with Box Annotations: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2021.*

## Further enquiries

**Australian Institute for Machine Learning**
The University of Adelaide SA 5005 Australia
**enquiries** aiml@adelaide.edu.au
**web** adelaide.edu.au/aiml
**twitter** twitter.com/TheAIML


September 2022. Job no. UA30491 - CD
CRICOS 00123M

**Kaurna acknowledgement**
We acknowledge and pay our respects to the Kaurna people, the original custodians of the Adelaide Plains and the land on which the University of Adelaide's campuses at North Terrace, Waite, and Roseworthy are built. We acknowledge the deep feelings of attachment and relationship of the Kaurna people to country and we respect and value their past, present and ongoing connection to the land and cultural beliefs. The University continues to develop respectful and reciprocal relationships with all Indigenous peoples in Australia, and with other Indigenous peoples throughout the world.