
THE UNIVERSITY
of ADELAIDE
Annual Report
2018
AUSTRALIAN
INSTITUTE FOR
MACHINE LEARNING
adelaide.edu.au/aiml

OUR VISION

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

OUR MISSION

Transform the way we live for the better through research excellence in machine learning, artificial intelligence (AI) and computer vision.

ABOUT THE INSTITUTE

The Australian Institute for Machine Learning (AIML) is Australia's first university-based institute dedicated to research in machine learning. It launched in early 2018 with co-investment from the South Australian Government and the University of Adelaide.

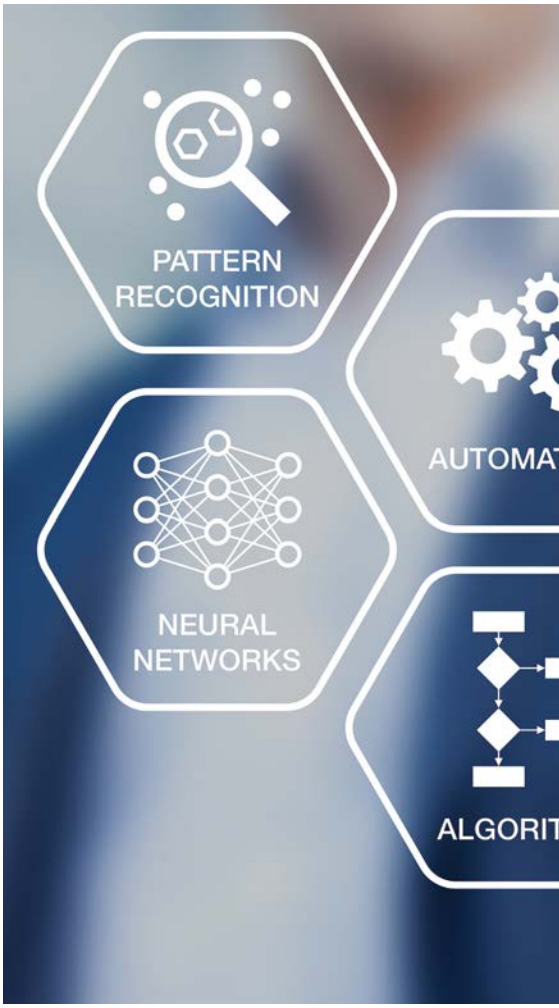

This annual report marks the completion of our first year of operation. We had a head start, however AIML was formed from the Australian Centre for Visual Technologies, which already had a long and successful history of high-impact research.

At the end of our first year, AIML already has over 100 researchers, making it the largest group of its kind in in Australia.

Our capabilities

- Machine learning
- Deep learning
- Image captioning
- Medical imaging
- Object recognition
- Parameter estimation
- Robust fitting
- Segmentation
- Video surveillance
- Augmented reality
- 3D modelling
- Tracking
- Visual question answering.

Our impacts

Machine learning is a powerful tool delivering significant benefits in almost every area of society. At AIML we are applying machine learning to a wide variety of applications with a diverse set of partners to achieve impact in a surprising range of fields.

Here are just some of the examples of how our research is making an impact.

Speeding disease diagnosis

The world's first AI microbiology screening technology for use in pathology labs.

Developed with LBT Innovations, the Automated Plate Assessment System went into production in 2017 and is attracting huge international interest. The system dramatically accelerates patient diagnosis and treatment, helping doctors to quickly address infectious diseases. A device commercialising this technology was the first intelligent medical device in its class to receive FDA 510(k) approval.

QUICK FACTS

Global spending on AI will increase from $12 billion in 2017 to $56 billion in 2021 (Deloittes, 2018).

AI has the potential to increase Australia's productivity by ~40 per cent (PWC, 2017).

Two thirds of Australian organisations are having difficulty finding suitable staff to lead AI technology integration (Leadership in the Age of AI – Infosys, 2018).

Accelerating crop farmers' adaptation to climate change

Image-analysis technology to analyse the physical structure of plants.

The system uses multiple images taken from numerous angles to construct computerised 3D models of the plants for analysis. It enables biologists to rapidly identify robust varieties able to thrive in harsh conditions, thus helping feed the world past peak human.

A step closer to reliable visual tracking

A unique algorithm for exceptional visual tracking.

Dragonflies catch their prey with more than 95 per cent accuracy. We teamed with engineers and neuroscientists to adapt the insect's neuronal processes to enable machines to visually track in the same manner. Fast, accurate visual tracking has applications in driverless cars, retail, and sports, amongst many other sectors.

Monitoring traffic incidents and road congestion

Analytics for incident detection, vehicle counts, and queue management.

In collaboration with the SA Government, AIML are working to improve the analysis and prediction of traffic patterns. This allows the real-time identification of incidents like traffic accidents and obstructions in order to guide the traffic systems' response and reduce travel times. Future work includes identifying the highest value opportunities to improve the traffic systems, by adding bicycle lanes and altering public transport routes, for instance.

AIML's strategies

Recruit and retain high calibre people

AIML has built a reputation as a high-performing research group because we have talented staff and students. Our standards for recruitment are very high, and we are able to promote a fantastic work environment. This continues to strengthen the international reputation of the group.

High quality research

AIML is ranked in the top four in the world on many metrics because of our commitment to high quality research, and research translation. Visibly achieving high-quality research in an extremely competitive field is critical to our ability to attract the top staff, students, and commercial partners.

Research impact

AIML's history of high-impact pure and applied research is critical to our credibility, and thus to our continued success. Our reputation has been developed through decades of world-class research, and collaboration across a variety of fields.

AIML's key actions

- Collaborate with world-leading researchers and companies to develop high-impact products and solutions
- Identify and exploit new opportunities for high-impact research and research translation
- Raise the profile of South Australia and the University of Adelaide as a a focus for high-quality research
- Identify and attract high-quality staff and students to grow AIML's research capacity.

Professor Mike Brooks, AIML Advisory Board Chair, the University of Adelaide.

AIML ADVISORY BOARD CHAIR'S REPORT

The 21st Century is a time of considerable change. Artificial intelligence is beginning to leave a major impact on our lives and has huge growth potential.

That's why, in 2018, the University of Adelaide, with co-funding from the South Australian Government, established our sixth University Research Institute – the Australian Institute for Machine Learning (AIML).

AIML aligns extremely well with the University of Adelaide's Strategic Plan, Future Making, which is designed to help realise its purpose as a catalyst of knowledge creation and innovation, an engine of social advancement and an active participant in the local, national and global economies.

AIML is also strongly engaged with the five pillars of the Strategic Plan:

1. **Connected to the Global World of Ideas:** Some companies are ahead of others when it comes to thinking strategically about artificial intelligence. They are the kind of companies with which AIML is partnering.
2. **A Magnet for Talent**: AIML hosts top performing researchers from all over the world, helping to extend our global reach.
3. **Research that Shapes the Future**: Artificial intelligence is the key to the future, and we want Adelaide to be on the front foot. An institute dedicated to advancing machine learning ticks this box.
4. **A 21st Century Education for Growing Community of Learners**: AIML is providing our bright young minds with opportunities to stay in South Australia to study and pursue a career in emerging technologies.
5. **The Beating Heart of Adelaide**: AIML has already become a catalyst to bring companies and people into South Australia and to generate significant economic activity.

This first annual report provides an overview of AIML's highly-productive first year of operation. We hope you enjoy it.

DIRECTOR'S REPORT

The first year of AIML was incredibly busy, both with the process of setting up an Institute and dealing with the incredible enthusiasm that its commencement generated.

Over the course of the year we met with many hundreds of interested parties and secured a host of important relationships. It is only the beginning, but we've made incredible progress already.

AIML was built on the core of a strong existing research group, the Australian Centre for Visual Technologies (ACVT). The ACVT started with 5 people in 2007, and had grown to over 80 in 2018. Although AIML has grown to over 100 people, it is still fundamentally a single research group working in one core technology area. The event that made the transition from a Centre to an Institute possible was securing the support of the South Australian state government, that came in the form of additional funding, and a new building. The funding has been pivotal in allowing us to expand our engagement with SMEs, the government, and the community more broadly. Access to the building has been slowed by unavoidable delays in its preparation, but will enable AIML to grow larger and engage further. The future looks very bright.

Our focus in year one was to continue to establish AIML by filling key roles, increasing our visibility nationally, and securing additional research funding. This period saw a particular focus on the business development process of the Institute, and the development of infrastructure and business systems. We formalised the relationship with the State Government, particularly around the process of generating and approving projects, and developed relationships with the key thought leaders and decision makers within South Australia, and nationally.

The first year also saw the initiation of the four Research Programs that we developed with the South Australian Government:

- **Artificial Intelligence Skills Development.** We developed a new major in artificial intelligence for the university's Bachelor courses.
- **Defence Industry Engagement.** We signed a three-year collaborative agreement with Lockheed Martin.
- **Government Efficiency Engagement**. We commenced the AddInsight Traffic Management System, which will help South Australians timely information to navigate our roads.
- **SME Engagement and Global R&D**. We started planning how to integrate and adopt machine learning to benefit South Australian SMEs, many of whom have already approached us with project ideas.

As part of our role as a national leader in the field, AIML has made an ongoing contribution to the debate around research funding for AI, and developed a National Artificial Intelligence Research Paper that has been well received.

Importantly, while all of this work was being carried out, the research capability and reputation of AIML has continued to grow. We have moved into new application areas, and strengthened our engagement in existing fields. Recent appointments and announcements indicate that AIML growth in scale and reputation are destined to continue in future years.

The above inevitably represents a small fraction of the work carried out by an amazing team. It has been a privilege to be a part of it.

ADVISORY BOARD REPORTS

Professor Mike Brooks

Chair AIML Advisory Board, Provost of the University of Adelaide

After announcing that the Australian Institute for Machine Learning was to be established, we were inundated with enquiries from companies, government and research institutions about how to partner with us. Australia needed an institute for machine learning, and we gave it to them.

Professor Anton van den Hengel

Director AIML

The first year of AIML has seen incredible growth in our capability, research income, and visibility. It's particularly gratifying to see that so many of the existing and new staff have driven these results.

Professor Anton Middelberg

Executive Dean Faculty of Engineering Computer and Math Sciences (ECMS)

We have successfully launched new degrees and courses in Cybersecurity and Machine Learning. Teaching staff report that they need to do a lot of remedial maths. The Board is keen to do more with secondary and university schools to better prepare the next generation of AI researchers.

Dr Tony Lindsay

Director Lockheed Martin STELaRLab (Science, Technology, Engineering Leadership and Research Laboratory)

It's wonderful for a new institute to secure a major contract in its first year. AIML's three-year strategic partnership agreement with Lockheed Martin, with significant cash and in-kind investment from LM, is just that. We'll work together to advance machine learning techniques in intentional machines.

Ms Lusia Guthrie

Independent Business Consultant, Chair BioMelbourne Network, Chair Medicines Manufacturing Innovation Centre

One of the keys to AIML's success is going to be in how well it brings innovative ideas and products to global markets. The Board will assist researchers wherever possible, from forming strategic partnerships with important companies and governments, to concept and product development.

Mr Adam Reid

Executive Director Industry, Innovation, Science and Small Business, Department for Innovation and Skills

AIML was established with the help of a $7.1 million funding agreement with the State Government. AIML's focus on skills development, defence industry engagement, government efficiency, and SME and global R&D will drive improvements in productivity that underpin sustainable economic growth in South Australia.

Developing talent in rapidly advancing technology domains, including machine learning, is critical to ensure that South Australia is a destination for technology based businesses that can solve globally relevant problems and build successful businesses in South Australia.

AIML's collaborative research partnerships with defence companies will address the priority needs of the defence industry and contribute to Australia's necessary industrial capabilities. This will continue to build on South Australia's long history of success in defence research, underpinned by strong education and industry partnerships.

AIML delivers an innovative model for effective research collaboration between the University, industry, start-ups and government.

RESEARCH THEMES

The core of machine learning is the development of systems that are able to learn by example. This is important for a range of tasks, but particularly for those that humans have difficulty in specifying algorithmically.

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

Our research can be broken down into 6 major themes.

RESEARCH PROGRAM

MACHINE LEARNING THEORY

Advancing the mathematical fundamentals of the field.

UNDERSTANDING HUMAN BEHAVIOUR

Extracting information about human behaviour from large volumes of data.

ROBOTIC VISION

Enabling machines that see (note the strong collaboration with the Australian Centre for Robotic Vision in this endeavour).

VISUAL QUESTION ANSWERING (VQA)

Enabling natural language interactions with systems that exploit visual information.

TRUSTED AUTONOMOUS SYSTEMS

Machines that work cooperatively with humans.

PHOTOGRAMMETRY AND 3D MODELLING

Advancing the mathematical and statistical methods that recover geometric information from images.

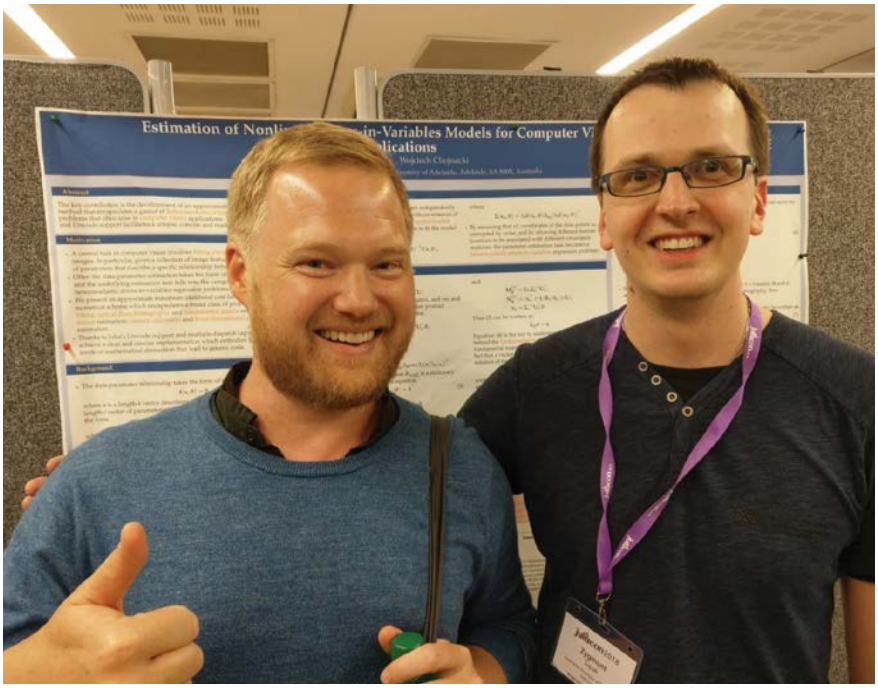

Research Theme Leaders

Research Theme Leaders will be appointed from time to time to build strength and critical mass in particular areas of strength. Theme Leaders will develop a strategy for increasing research income and activity in their area of responsibility. They can draw upon the support of the Institute (staff and seed funding) to develop and implement their strategy.

Theme Leaders may choose to be appointed for a defined period of time to enable a concentrated effort to progress a strategy in a particular domain.

Machine Learning Theory
Professor Chunhua Shen

Understanding Human Behaviour
Associate Professor Javen Shi

Robotic Vision
Professor Ian Reid

Visual Question Answering
Dr Qi Wu

Trusted Autonomous Systems
Professor Anton van den Hengel

Photogrammetry and 3D Modelling
Associate Professor Tat-Jun Chin

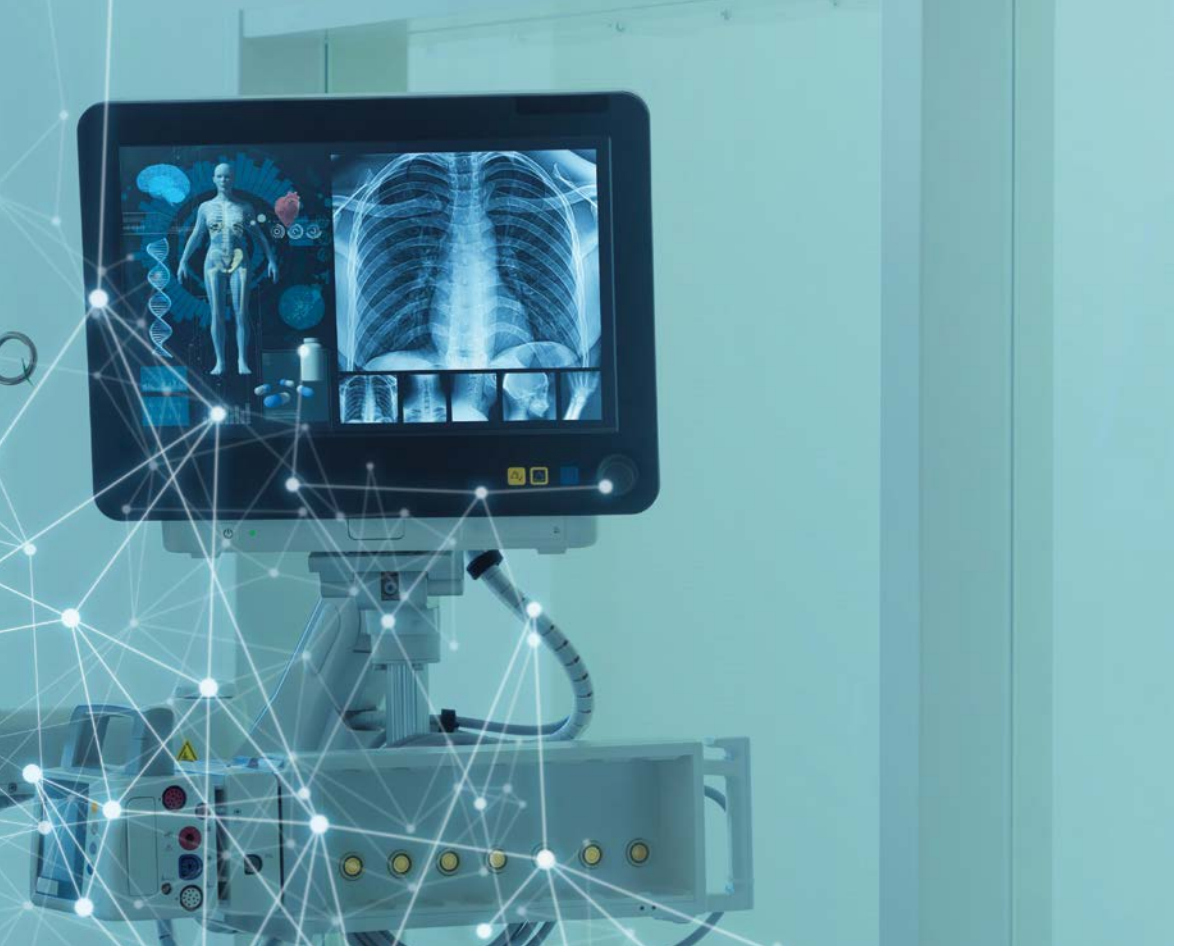

MACHINE LEARNING THEORY

Machine learning continues to be a fascinating field of research. At its most basic, it is the science of getting computers to act and learn over time without being explicitly programmed. But its areas of research and application are exploding.

What can machines learn?

After teaching computers what to do with examples of data and information, machines can learn to:

- detect disease and other medical problems
- operate machinery and drive cars
- recognise faces and groups of people
- predict retail buying trends and socio-economic patterns.

And much more!

Research strengths at AIML

A particular strength of researchers at AIML is our ability to teach computers using only weakly labelled data. Our researchers are making major contributions in four key aspects of machine learning theory:

- **Deep learning**, by advancing the mathematical tools that underpin the training of computers to perform human-like tasks
- **Systems optimisation**, by developing the theory, algorithms and tools that can predict environmental factors such as electricity prices and the weather, which are constantly changing
- **Robust statistics**, by developing procedures to analyse data to make sure that information from machines remains informative and efficient
- **Probabilistic graphic models**, by improving the way that machines model complex relationships among variables, to improve the reliability of the inferences they make.

ROBOTIC VISION

When machines have the ability to see, their capabilities grow substantially. Using one or more video cameras, they are able to collect visual data to understand the physical world.

Why pursue vision for machines?

As humans, we rely heavily on our vision to perform all sorts of tasks, including to see where we're going, who is in a room, and the emotions of faces. If we could enable machines to see as well – or better than – we do, this would open up amazing possibilities.

Research strengths at AIML

Here are three key areas of robotic vision we focus on:

- **Visual simultaneous localisation and mapping (SLAM)**. This technology uses a 3D vision camera to determine the position and orientation of the machine, while mapping the unknown surrounding environment. It enables field robots, drones and autonomous vehicles to navigate independently. The technology is not reliant on satellite information, but is able to accurately measure the physical world.
- **Semantic vision**. It's one thing to collect vision data, but how can we be sure that machines can make sense of it? Semantic vision is the field of processing the often huge amount of data in a way that produces meaningful, understandable information.
- **Vision and language**

Australian Centre for Robotic Vision

AIML is a founding partner of the ACRV, which carries out breakthrough science and technology under four research objectives: robust vision, vision and action, semantic vision, and algorithms and architecture.

AIML's Professor Ian Reid is Deputy Director of the centre, and leads the Scene Understanding Project. AIML's Director Anton van den Hengel leads the Vision and Language Project. Many of AIML's researchers and students play a role in the centre's projects.

TRUSTED AUTONOMOUS SYSTEMS

Trusted autonomous systems are reliable and independent. They do not need to be operated by humans. But they work alongside humans, and may communicate, cooperate and negotiate with us or other autonomous systems to achieve goals.

The road to independence

In the beginning, machines were under full human control. Then we created supervised systems (for example, automated sewing machines), and then automatic systems (such as car assembly robots).

Autonomous systems are the final step. They can operate without any human intervention at all.

How can they be trusted?

The 'trusted' part is the most challenging. Trusted autonomous systems are often being applied in situations where there are big consequences for failing, such as driverless cars and robotic defence technologies. So we need to make sure they are fully reliable with very low probability for error.

Research strengths at AIML

Solving this is our key focus at AIML. We have:

- developed autonomous systems that are aware of the uncertain environment and can perform complex tasks successfully
- carried out theoretical and practical research to develop systems capable of making transparent and explainable decisions, asking questions when uncertain about decisions or the surrounding environment, and understanding their interactions with the world and applying reasoning to their surrounds
- made significant progress towards intelligently controlled autonomous vehicles.

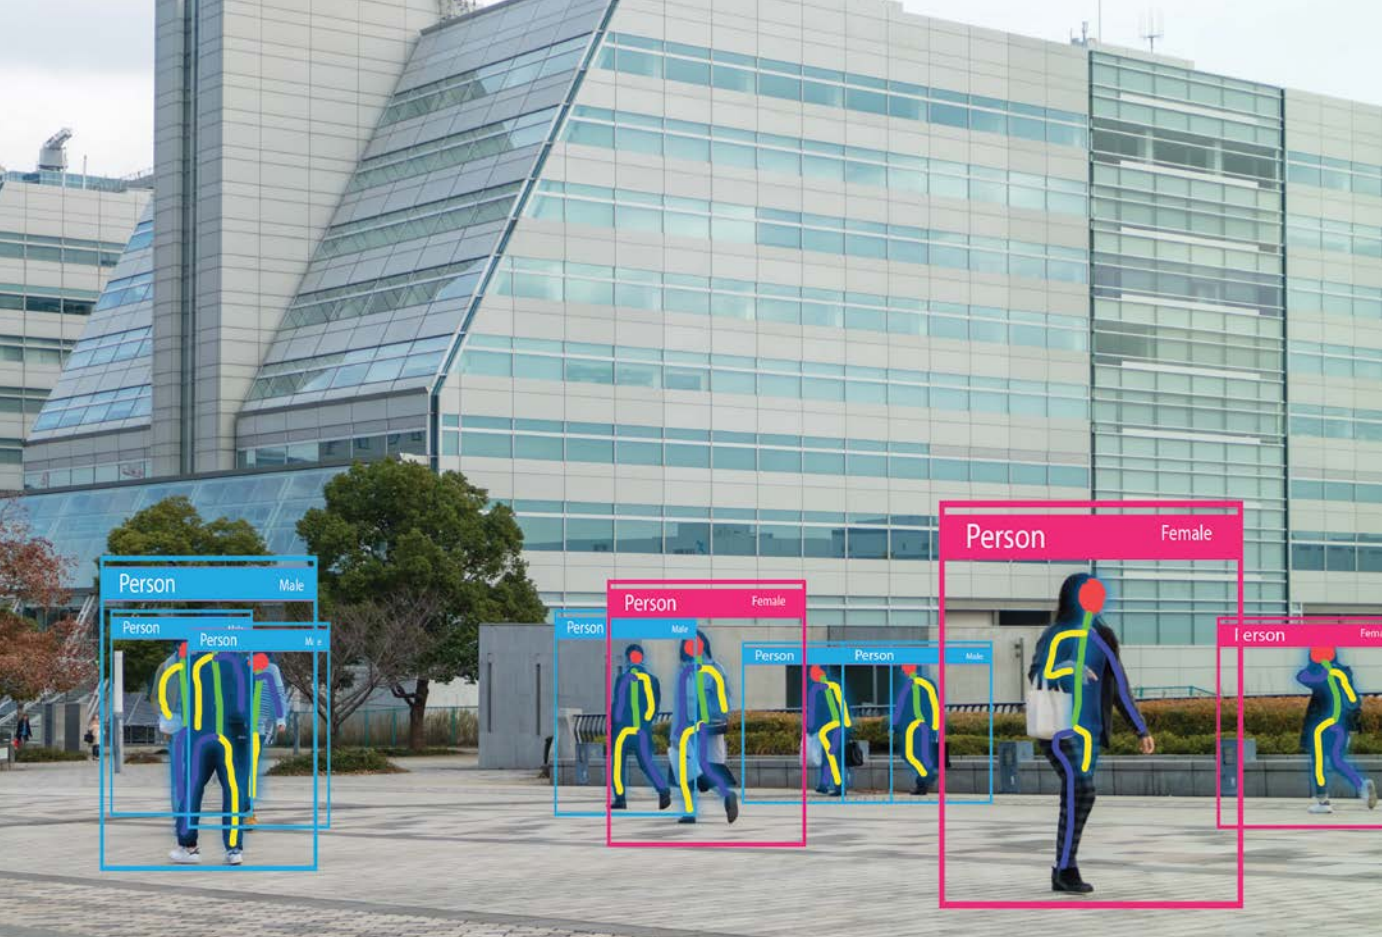

VISUAL QUESTION ANSWERING

Recognising cats and dogs is now an easy task for a computer. The next frontier? Computers that understand natural (human) language and vision (images) in order to answer questions.

Teaching with data

Machine learning starts with making each system learn from data; we do not hard code. A big challenge can be having access to enough data. Fortunately, AIML has huge data sets from partners, in which we have the question, the image and the correct answer. These are standardised data sets.

Each VQA system can be individualised to suit a particular task, from assisting with medical treatment plans, to assisting people who are vision impaired.

Research strengths at AIML

Our two major areas of focus are:

- **Dialogue**. How can we make VQA useful in real life? When people converse with each other, they don't ask just one short question; there is to-ing and fro-ing and ongoing interaction. A useful VQA system will need to keep track of what is being asked, know how to follow up, and maybe ask questions back to the person to clarify. To date, systems are only good at this when they have been heavily trained.
- **Data retrieval**. How can we enable VQA systems to retrieve information as it is needed? Let's say the system receives a question about zebras, but it doesn't know what a zebra is. We want it to be able to find out by itself.

PHOTOGRAMMETRY AND 3D MODELLING

Photogrammetry

The basis of photogrammetry is using photography to work out measurements, such as the distance between objects or the shape of an object.

To do this, we need to take multiple measurements from images so we can reconstruct the world. This is already an established field.

3D modelling

In the field of augmented reality, when synthetic content is added to a real image, such as 3D special effects to a photo, the synthetic content can 'float' over the photo and appear not believable. The challenge is to fix the synthetic geometry to the same frame as the objects around it.

Research strengths at AIML

AIML researchers are renowned for fundamental and applied research in photogrammetry, multiple-view geometry, and structure-from-motion.

- **Photogrammetry**. Our key areas of interest are facilitating accurate metric measurements of real-world objects from images, and researching novel statistical models and parameter estimation methods to make the resulting models and measurements as precise as possible. Our specialty is to work at the limits of resolution; where objects are so far away or so small that even the best cameras can't properly image them. It's not enough to merely measure things. We also need to know what exactly is being measured. So, we don't just want to know that the computer is measuring an object, we want to know what the object is.
- **3D modelling**. Our key areas of interest include real-time special effects in live video and automatically constructing high-fidelity 3D models from a collection of images.

DETECTING BREAST CANCER FASTER AND MORE ACCURATELY

ARC RESEARCH FELLOW GABRIEL MAICAS SUSO

Normally, you need a doctor to tell you if a medical scan contains a lesion. Well-trained doctors can tell just by looking at it.

Researchers at AIML wanted to know if machine learning could increase the speed and accuracy of diagnoses of cancerous lesions. As part of a bigger project with the Queensland University of Technology, Research Fellow Gabriel Maicas Suso and Professor Gustavo Carneiro designed a machine learning technique to diagnose breast cancer from MRI scans. MRI is commonly used to diagnose breast cancer in high risk patients.

The results were outstanding.

"We got very good results when deciding which images are suspicious," explains Gabriel.

"We were able to increase diagnostic accuracy, so that the machine can tell the doctor, for example, out of these 100 images don't look at these 90 as they definitely don't contain lesions."

The team also tried to localise an image, to spot exactly where the lesion is. "We were able to speed up detection by a factor of 1.78 while maintaining the same level of accuracy," says Gabriel.

The technique applies deep reinforcement learning, which enables machines to learn how to do complex tasks without being programmed by humans. And because they didn't have a lot of MRI data, the team also tried meta-learning, a promising new field in computer vision.

"In meta-learning, you teach a machine how to solve the problem," Gabriel explains.

"We trained the machine using minimal data, which contained just a few images. Those few images allowed us to increase diagnostic accuracy. This really surprised us."

One day, computers will routinely help doctors to increase their accuracy in the detection of cancer, as a very reliable second opinion.

This project was funded by the ARC.

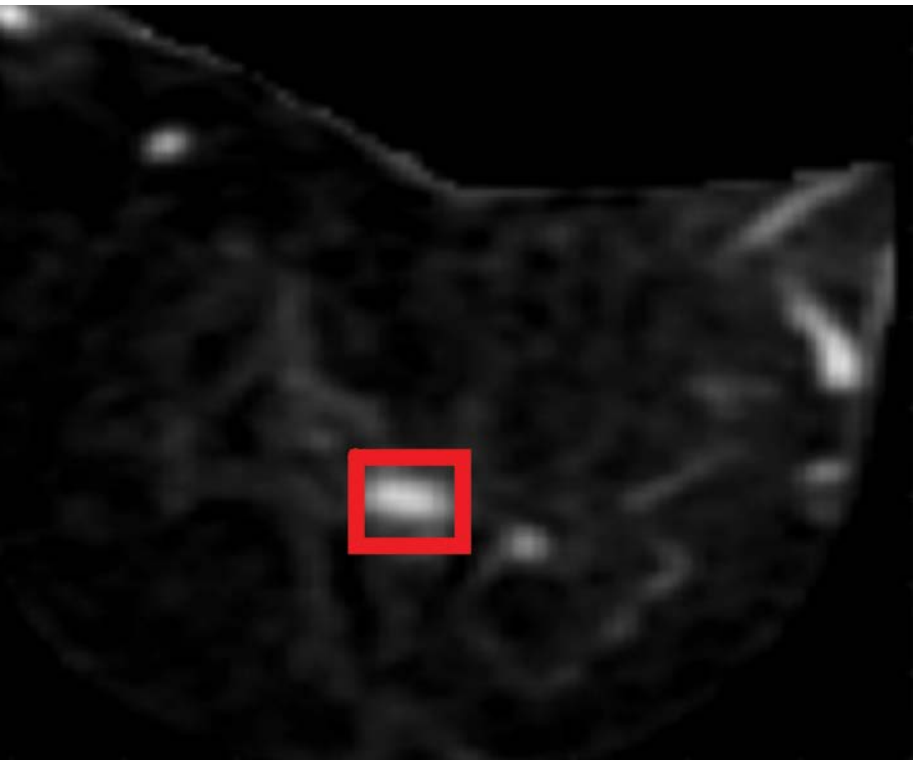

"We trained the machine using minimal data, which contained just a few images. Those few images allowed us to increase diagnostic accuracy. This really surprised us."

BEYOND RECOGNITION TO UNDERSTANDING

PROFESSOR IAN REID

To enable machines to see, we equip them with sensors and teach them to recognise objects. A great deal of research effort has been put into this in recent years.

But if robots are going to be driving cars and doing other important tasks one day, recognising objects is not enough. We must also teach them to understand which objects are important in a scene, and which are not. This is the science of scene recognition using semantic understanding.

"The aim of our project is to create semantic representations of the environment using visual observations," says Professor Ian Reid, Head of the School of Computer Science at the University of Adelaide, and Deputy Director of the Australian Centre for Robotic Vision.

"By this we mean that we want machines to apply meaning to objects like a human would."

The team started by assuming that visual SLAM and 3D geometry will soon be efficient and accurate enough to be useful.

"We then incorporated a lot higher level of understanding to the models in a series of projects," explains Ian.

> "The ultimate goal is for machines to be able to understand the 'affordance' of objects – what they're used for, and how – so that they can act appropriately."

The team worked to not only predict what machines see, but also apply meaning, all in a self-supervised manner. They are enabling machines to:

- understand that objects sitting on top of other objects are separate items, and understand what role each object plays
- understand that objects they don't recognise (like chewing gum on a floor) are separate entities to the things around them (like the floor)
- tell us the object class of every pixel in an image, as well as what each object is used for and thus how one should interact with it.

BABY STEPS: LEARNING WITH TRIAL AND ERROR

DR EHSAN ABBASNEJAD

When a toddler takes her first steps, she's taking a risk. She risks overbalancing and falling down. Her gait isn't quite right (too short, long or stilted), and her feet don't work in perfect alignment with her toes.

Just like small children, trusted autonomous systems learn through risk taking and trial and error.

AIML researchers recently developed a program that uses trial and error to quickly learn the best way to achieve goals, without instruction from a human.

"We encouraged the computer to take risks," explains Dr Ehsan Abbasnejad.

"We wanted the computer to learn how to confidently make good decisions on its own."

Rather than teaching the computer how to get better at a task, the team asked the computer to learn by trial and error by playing the game Pac-Man. In this 1980s video game, there is a reward (go up levels and win points) when you survive, choose the right paths and avoid the enemies.

But there is a penalty (lose a life or start again), if you make a mistake. The computer got better at the game the more it played it.

"We didn't provide any labelled examples. We didn't instruct it to, for example, 'turn left' when there was an enemy approaching; it had to figure it out," says Ehsan.

> "When it got something right there was a payoff. When it got something wrong it was penalised."

Trial and error learning for trusted autonomous systems has applications for driverless cars, robots, defence systems and any other situation where a robot is asked to achieve a task without any human guidance.

BRING ME A SPOON

DR QI WU

It's the ultimate futuristic fantasy. Robots living alongside us and carrying out tasks at our command. But implementing this in real life is currently impossible.

"To make it happen, we need to create new techniques for linking natural language to vision and action in unstructured, previously unseen environments," explains Dr Qi Wu, the project lead. "Our ultimate goal is to enable machines to behave more like humans."

So when a robot is asked to bring a spoon, it needs to be able to think about where a spoon might be, navigate to the right place, and ask clarification questions if needed – just like a human would.

In 2018, the team built the Matterport3D Simulator – the first ever simulator that uses a 3D image of a real house.

"Using the simulator, we taught the computer basic navigation like how to access the bathroom and kitchen," says Qi. "In future, the simulator will support a range of embodied vision and language tasks, and right now, it is the first benchmark dataset for visually-grounded natural language navigation in real buildings."

Using Matterport3D, humans need only provide short navigation instructions with the rough location of the spoon. Other simulators require very detailed instructions (such as turn left, go upstairs).

This process – called room-to-room navigation – is an essential first step in reaching the ultimate goal of robots that carry out natural language commands.

"The second step is for the machine to not only go to the right room, but also grab something," says Qi. "We've done work on this too, by training a machine to navigate through a previously unseen environment to select an object identified through general natural language."

After that, it will get even harder. Robots will have to learn how to find hidden objects (such as a spoon in a drawer) and ask clarification questions (such as whether you want a teaspoon or dessert spoon).

Instruction: Go to the stairs on level one and bring me the bottom picture that is next to the top of the stairs

| FUNDING: ARC DECRA

I CAN SEE CLEARLY NOW THE NOISE HAS GONE

DR ZYGMUNT SZPAK

When you look at a white elliptic shape on a black background very far away, such as in outer space, it appears blurry or pixilated and, unfortunately, not elliptic.

This problem occurs when we look at any object at the limits of resolution. Fields such as astronomy, microscopy and defence suffer from this problem.

"Our challenge was to determine the true shape of ellipses from low resolution images," explains Dr Zygmunt Szpak who, together with Professor Wojciech Chojnacki, developed a new image processing technique to solve the problem.

"We did this by creating an intricate mathematical model of the image formation process at the limits of resolution."

The team reasoned that if they could understand how the forward process works (that is, start with the perfect image and distort it until it is a blurred low-resolution ellipse), they could invert this to estimate the shape.

"We looked at what happens under different conditions, such as when a camera receives a limited number of photons, or stores a small number of colours," says Zygmunt. "The ultimate aim is to provide a machine learning solution to any imaging scenario."

"Incredibly, and rather unexpectedly, we worked out how to fit an ellipse accurately, even when it was barely visible in the data", says Zygmunt.

> "The computer works out for itself the level of noise and how to correct it."

| FUNDING: DSTG CERA GRANT

AIML STAFF

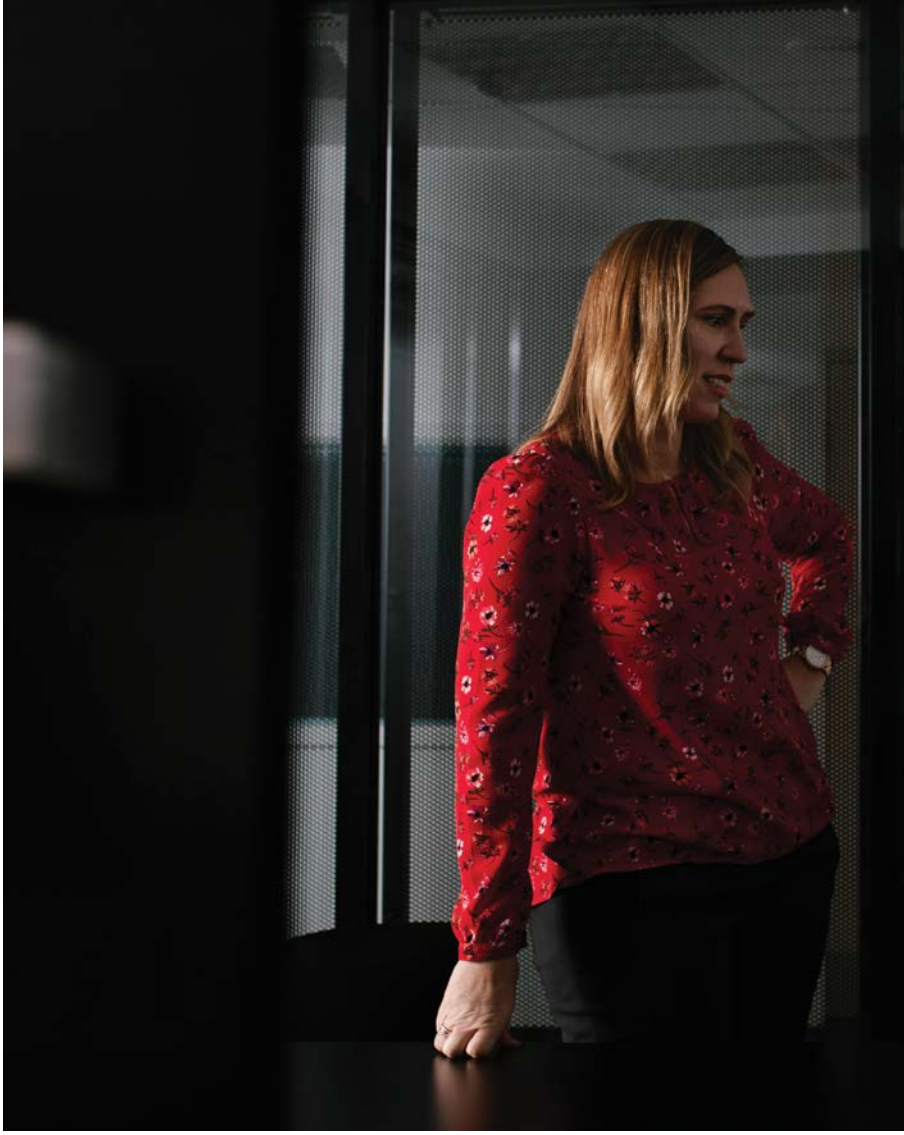

AIML Director

Professor Anton van den Hengel

AIML Academics

Professor Ian Reid
ARC Laureate Professor

Dr Lingqiao Liu
ARC DECRA Research Fellow

Associate Professor Anthony Dick

Associate Professor Javen Qinfeng Shi
Director, Advanced Reasoning and Learning

Associate Professor Tat-Jun Chin
Director, Machine Learning for Space

Professor Gustavo Carneiro
Director, Medical Machine Learning

Dr Qi Wu
ARC DECRA Fellow

Professor Chunhua Shen

AIML Research Team

Post-Doctoral Researchers

Dr Dong Gong
Dr Peng Wang
Professor Wojciech Chojnacki
Dr Zifeng Wu
Dr Zygmunt Szpak
Dr Damien Teney
Dr Ben Ward
Dr John Bastian
Dr Jamie Sherrah
Dr Guansong Pang
Dr Yuankai Qi
Dr Jeiwei Cao
Dr Yuanzhonhan Cao
Dr Peng Chen
Dr Hamid Rezatofighi
Dr Bohan Zhuang
Dr Saroj (Chamara) Weerasekera
Dr Hui Li
Dr Yasir Latif
Dr Michele Sassdelli
Dr Pulak Purkait
Dr Ravi Garg
Dr Wei Liu
Dr Bo Chen
Dr Alvaro Parra Bustos
Dr Ehsan Abbasnejad
Dr Gabriel Maicas

Software Engineers

Alex Cichowski
Phil Roberts

Machine Learning Engineers

Karl Hornlund
Sebastian Parkitny

Research Programmers

Thomas Rowntree
Kiet To
Sam Bahrami

AIML Executive Team

Ian Will
Operations Manager

Dr Paul Dalby
Business Development Manager

Emily Holyoak
Communications and Engagement Manager

Ali Anderson
Outreach Programs Manager

Rachel Kontic
Senior Administrative Officer and PA to the Director

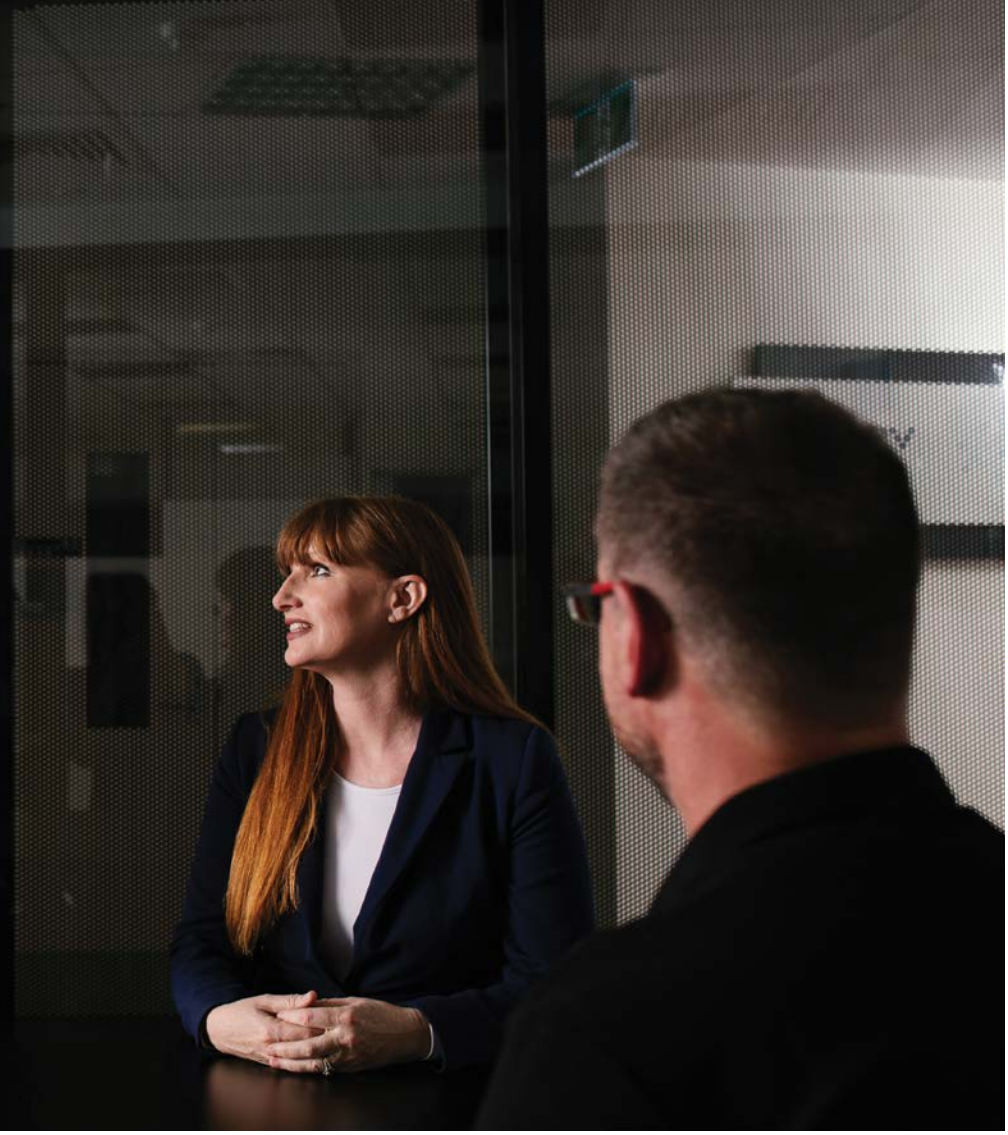

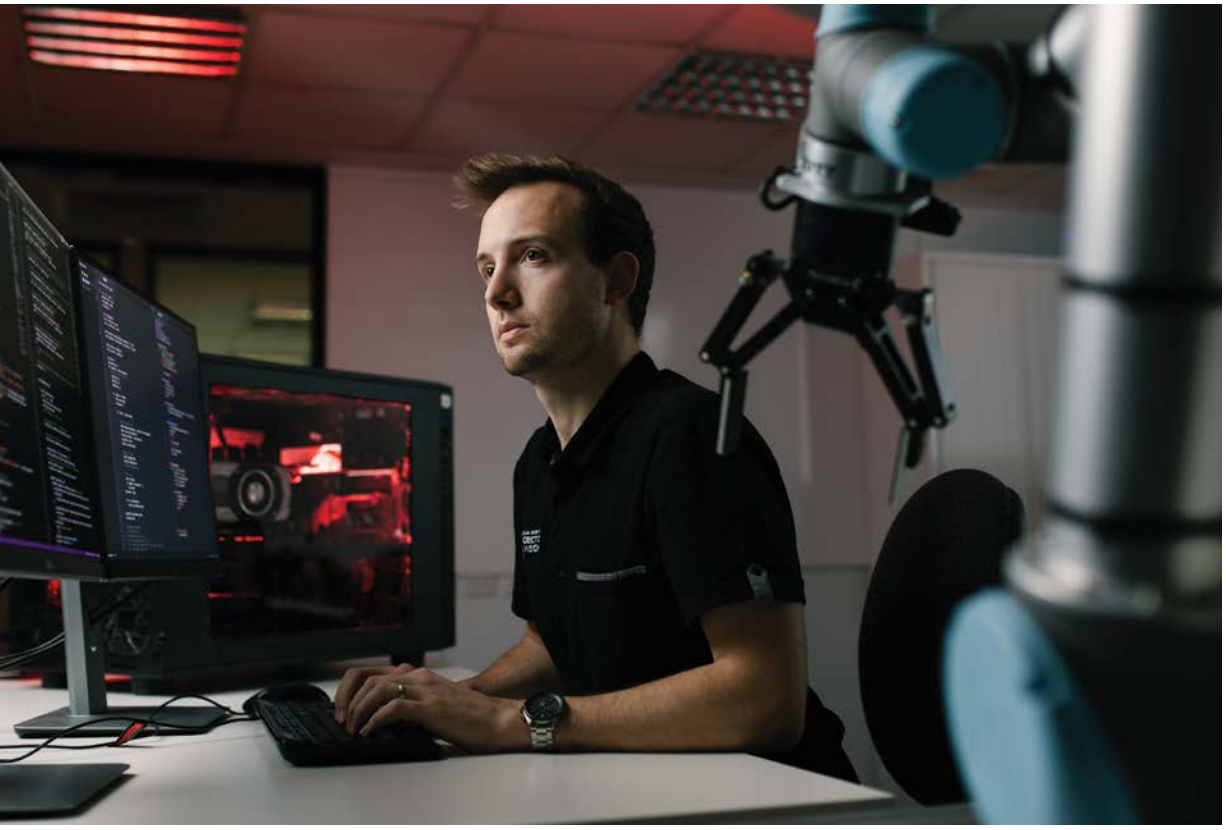

COLLABORATION

We achieve great things because we collaborate.

Partnerships

Lockheed Martin

Technology and innovation company Lockheed Martin Australia (LMA) announced in 2018 that they would be AIML's first Foundation Partner. The partnership is delivering machine learning research for national security, the space industry, business, and the broader community. LMA will move a team of researchers from its STELaRLab (Science Technology Engineering Leadership and Research Laboratory) in Melbourne to be co-located with AIML researchers in Adelaide. Premier Steven Marshall made the official partnership announcement at an event held at Lot Fourteen.

AIML have also secured $100k in funding from Lockheed Martin for the Adelaide Design and Engineering Practical Team (ADEPT), a self-organised group of University of Adelaide students. This funding will enable ADEPT to expand its activities, and generate higher value opportunities.

Lenovo

AIML signed a strategic partnership agreement with Lenovo (an internationally renowned hardware supplied facing increased demand for machine learning based solutions). The partnership will see Lenovo generate high-value collaborative opportunities for AIML. Professor van den Hengel and Dr Paul Dalby attended Lenovo's Transform 2.0 event in Sydney in December.

Defence

The Institute's long collaboration with the defence industry, in particular DST Group, continues. Although multiple projects have been funded, the focus is on developing a long-term partnership that will see AIML contribute to Australia's defence capability. The proposal has been met with enthusiasm from defence broadly.

Projects

Riverland grape grower's project

In early 2018, we met with Wine Australia and Riverland Wine at the Loxton Research Centre to consider how researchers could collaborate with wine growers to use AI, machine learning and robotic vision to develop a Digital Vineyard Guidance System. The aim is to increase productivity and efficiency in vineyards. Together, the team applied for and received a grant worth $134,000 from the South Australian Wine Industry Development Scheme. The team also includes members from the Faculty of Engineering, Computer and Mathematical Sciences, and the School of Agriculture, Food and Wine.

Defence research projects

AIML secured Defence Innovation Partnership funds to collaborate on two defence research projects:

- $150,000 for human-machine interfaces for detecting, monitoring and managing psychological stress, led by the University of Adelaide with partners the University of South Australia, Flinders University, ElectroAutoMedics and Defence Science and Technology.
- $150,000 for AI Enabling Australia's Future Submarine, led by Acacia Systems with partners The University of Adelaide, Flinders University, Defence Science and Technology and Lockheed Martin.

Multiple projects in Electronic Warfare have also been secured, although the details are confidential.

Clockwise from top left: Lockheed Martin partnership announcement, Lenovo signing, collecting data in the Riverland.

CONFERENCES

Conferences

Artificial Intelligence – The Future of Work

This workshop for the Committee for Economic Development of Australia (CEDA) took place in early November. Anton was the keynote speaker.

Submarine Industry Alliance conference

Anton van den Hengel and Horden Wiltshire (from Acacia Research) spoke about information processing in the Next Generation Submarine, and at the Centre for Defence and Strategic Studies about machine learning to the next generation of defence leaders.

AVDI Driverless Vehicles Summit

Anton van den Hengel spoke at the SA Government Department of Industry and Skills mining industry-research engagement day.

Hybrid World Tech Conference

Anton van den Hengel spoke at Hybrid World and AIML had a stand. His presentation on 'Some Ideas About the Future of Machine Learning' formed part of the theme around Machines that Think and Autonomous Machines.

MICCAI 2018

Research Fellow Gabriel Maicas presented on 'Training Medical Image Analysis Systems like Radiologists' at the 2018 international conference on Medical Image Computing and Computer Assisted Intervention in Granada, Spain.

ICCVG 2018

Zygmunt Szpak was invited to deliver the plenary talk at the International Conference on Computer Vision and Graphics 2018 in Warsaw, Poland.

ECCV 2018

PhD student Zhipeng Cai had two first-authored papers accepted as oral presentations at the European Conference on Computer Vision in Munich, Germany.

IDVS3

Anton van den Hengel spoke about the latest AIML machine learning research findings at the International Driverless Vehicle Summit in Adelaide.

CVPR2018

AIML was well represented at the Conference on Computer Vision and Pattern Recognition held in Salt Lake City. Dr Damien Teney, Dr Qi Wu, Professor Anton van den Hengel, Dr Bohan Zhuang, Professor Ian Reid, Dr Peng Wang, Associate Professor Anthony Dick and Professor Chunhua Chen all had papers accepted to the conference.

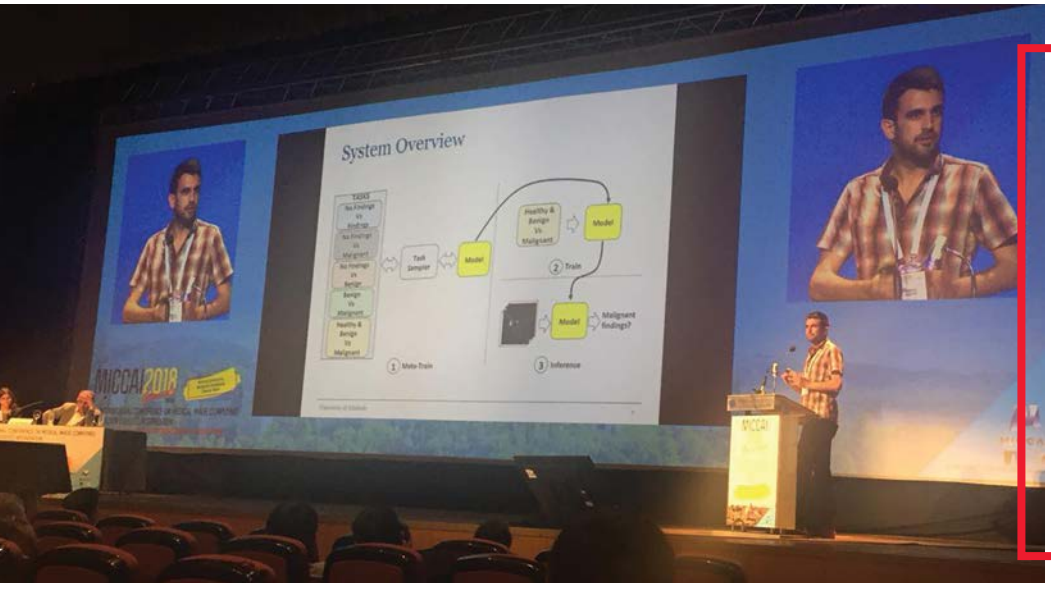

AIML hosted these special visitors

Dr Christyl Johnson
NASA

Associate Professor Simon Lucey
Carnegie Mellon

Professor Mike Ford
Institute for Nanoscale Technology at UTS

Assistant Professor Stuart James
UCL

Alessio Del Blue
Istituto Italiano di Tecnologia

ENGAGEMENT

We have impact beyond our walls.

Inspiring students

ADEPT partnership

AIML and Lockheed Martin's STELaRLab partnered with the Adelaide Design and Engineering Practical Team. The agreement gives students opportunities to work with industry.

Ingenuity Exhibition

In October volunteer staff from AIML had an information stand at the University's Ingenuity event, an exhibition showcasing current student projects.

University of Adelaide Open Day

AIML had a stand at Open Day next to the School of Computer Science. Guests had the opportunity to chat to our staff about careers in machine learning and AI, and watch Ur5sula the robot arm in action.

'Beer and pizza' nights

We hosted regular beer and pizza nights throughout 2018 to chat to undergraduate students about PhD opportunities.

Summer scholarships program

Undergraduate students can work with us under the Adelaide Summer Research Scholarship program run by ECMS. We offer project ideas and also encourage other suggestions.

Maths Education and Industry Outreach Forum

AIML was invited to speak at the forum by organisers the Australian Association of Mathematics Teachers. It was a great opportunity to inspire teachers and maths coordinators about machine learning.

School tours

We hosted our first school tour, students from Hallett Cove School. We plan to do more of these in future.

Connecting with the public

Scope TV

Dr Zygmunt Szpak filmed a segment about gathering data on cyclists to improve road safety, which aired May 2018.

AiLab

We sponsored two AiLab Meetups in 2018 and were involved in AiLab events including the launch night. AiLab brings businesses together with researchers, academia and the community to help everyone navigate the AI landscape.

RiAus roundtable discussion

AIML Director Anton van den Hengel participated in the discussion on AI with Chief Scientist Alan Finkel and Ellen Broad.

Lot Fourteen Open Day

We had an exhibition stand at the site of the former Royal Adelaide Hospital, which is being turned into Australia's first creation and innovation neighbourhood.

AI Collaborative Network

Our staff participated in the networking and panel discussion event, at which AIML's Research Programmer Thomas Rowntree spoke.

Wear It Purple Day

AIML took part in this important day on which diversity it celebrated, and respect and acceptance of young people of all identities is encouraged.

Space agency talk

Dr Tat-Jun Chin and Research Programmer Sam Bahrami were invited to be part of the press conference when Premier and Andy Thomas visited the University to talk about Adelaide's new space agency.

Research Tuesdays

Professor Anton van den Hengel and Professor Ian Reid hosted a sold-out public address on the subject of Machine Learning.

Media and public engagement

Much of the coverage of artificial intelligence in the media has been misguided. AIML has engaged with the media to provide informed and factual information.

In addition to those listed above, some examples include:

- InDaily article "What to expect from the robot uprising" 5/3/2018
- Nine News interview on 13/3/18
- Presentation to the Senate Select Committee on the Future of Work, May 2018
- Interview for Human Resources Magazine on the impact of machine learning on HR in September 2018
- Video on The Science Channel about machine learning in September 2018
- Professor van den Hengel gave a keynote at the Centre for Economic Development Australia conference on artificial intelligence - the future of work
- ABC News interview on AI and the future of work in November 2018.


Jumpgate VR

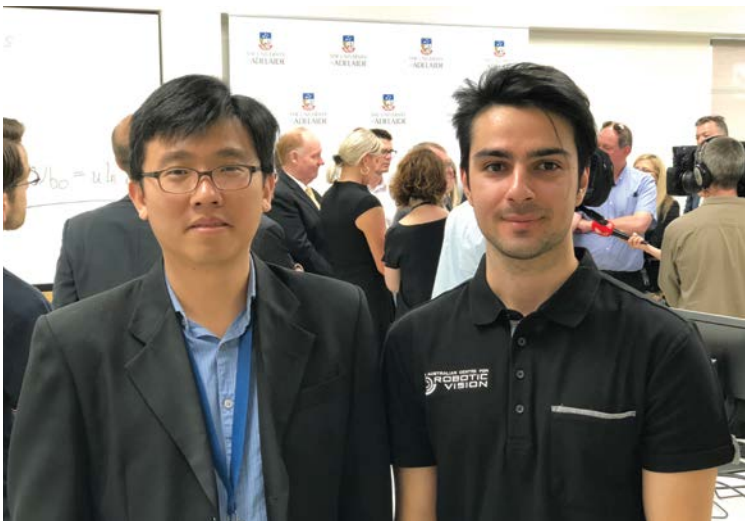

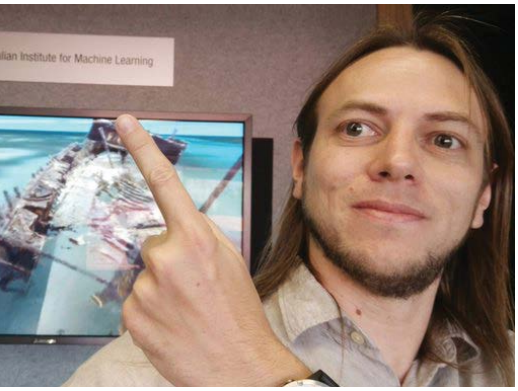

Defence engagement

Defence engagement has been a priority. Some examples include

- The AIML team won the ISPRS 2D semantic labelling challenge. This is hosted by US Defence, and demonstrates a key capability in extracting value from overhead images. This is of particular interest to Australian defence and intelligence agencies, and the win represents a validation of AIML's leading capacity in this critical area.
- Professor van den Hengel gave spoke at the Submarine Industry Association Conference in Canberra November 2018
- DST/Army Podcast on Robotics and Autonomous Systems 12/6/18
- Invited talk at the Centre for Defence & Strategic Studies in November 2018. This was an opportunity to present the potential impact of Machine Learning in Defence to the next generation of Defence leaders in Australia and the region. We have been asked to continue this engagement for every group that goes through the program.

COMPETITIONS

2ND PLACE IN XVIEW CHALLENGE

A team from AIML featuring Jamie Sherrah and Phil Roberts, and DST Group's Victor Stamatescu beat more than 4000 submissions from around the world to gain second place in the Defense Innovation Unit's (DIU) xVIEW Challenge.

GOOGLE SUMMER OF CODE

We were involved in our very first Google Summer of Code. Dr Zygmunt Szpak mentored student Arijit Kar from the Indian Institute of Technology Kharagpur. Together they worked on a Julia Language project under the umbrella of the NumFOCUS organisation.

FIRST PLACE!

AIML achieved success in the REFUGE Retinal Fundus Glaucoma Challenge, with a first place in the Segmentation Leaderboard and also in the Segmentatino of Nuclei competition.

NVIDIA PIONEER AWARD

Dr Qi Wu (ACRV / AIML) was awarded the prestigious NVIDIA Pioneer Award for his paper 'Learning semantic concepts and order for image and sentence matching' at the Computer Vision and Pattern Recognition conference in Salt Lake City. Dr Wu joins an elite group of researchers to receive this accolade; it is only the second time an Australian university has been involved in winning this award.

WE'RE NUMBER ONE IN VQA 2.0

A team led by Dr Damien Teney (AIML) and Peter Anderson (ACRV, ANU, and Microsoft) placed first in the VQA 2.0 challenge. Other members of the team include David Golub from Stanford, Po-Seng Huang, Lei Zhang and Xiaodong He from Microsoft, and Professor Anton van den Hengel from AIML.

AWARD FOR EXCELLENCE IN COLLABORATION

The collaboration between AIML and LBT Innovations won the SA Science Excellence award for Research Collaboration for the development of that APAS technology. APAS is an automated agar plate reader, which promises to greatly reduce the cost of this critical pathology procedure. APAS achieve first-in-class FDA approval last year.

NUMBER ONE IN SEMANTIC SEGMENTATION

Congratulations to Dr Zifeng Wu and Professor Chunhua Shen on having made it to the top of the Cityscapes leaderboard again.

Cityscapes is a semantic segmentation dataset of city scenes, and a hotly contested international challenge. The challenge is to separate the pixels belonging to different classes of objects. Semantic Segmentation is one of the fundamental challenges in computer vision, and underpins a variety of important practical applications.

The winning approach is based on a single Convolutional Neural Network (rather than an ensemble), and a shallow one at that. This approach is based on a new interpretation of the unravelled view of deep residual networks which explains some of the behaviours that have been observed experimentally. As a result, we have been able to derive a new, shallower, architecture of residual networks which significantly outperforms much deeper models.

2018 PUBLICATIONS

Abbasnejad ME, Dick AR, Shi QR, van den Hengel A, Active learning from noisy tagged images. BMVC Workshop, 2018.

Abbasnejad E, Wu Q, Abbasnejad I, Shi J, van den Hengel A. An Active Information Seeking Model for Goal-oriented Vision-and-Language Tasks. CoRR 2018;abs/1812.06398.

Abbasnejad E, Wu Q, Shi J, van den Hengel A. What's to know? Uncertainty as a Guide to Asking Goal-oriented Questions. CoRR 2018;abs/1812.06401.

Abdi M, Lim CP, Mohamed SMK, Nahavandi S, Abbasnejad E, van den Hengel A. Discriminative Clustering of High-Dimensional Data Using Generative Modeling. 2018. p 799-802.

Alfaro-Almagro F, Jenkinson M, Bangerter NK, Andersson JLR, Griffanti L, Douaud G, Sotiropoulos SN, Jbabdi S, Hernandez-Fernandez M, Vallee E, Vidaurre D, Webster M, McCarthy P, Rorden C, Daducci A, Alexander DC, Zhang H, Dragonu I, Matthews PM, Miller KL, Smith SM. Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank. Neuroimage 2018;166:400-424.

Anderson P, Wu Q, Teney D, Bruce J, Johnson M, Sünderhauf N, Reid ID, Gould S, van den Hengel A. Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments. 2018. p 3674-3683.

Ansdell M, Ioannou Y, Osborn H, Sasdelli M, Smith J, Caldwell D, Jenkins J, Räissi C, Angerhausen D. Scientific Domain Knowledge Improves Exoplanet Transit Classification with Deep Learning. Astrophysical Journal Letters 2018;869(1).

Battaglini M, Jenkinson M, De Stefano N, Neuroimaging AsD. SIENA-XL for Improving the Assessment of Gray and White Matter Volume Changes on Brain MRI. Human Brain Mapping 2018;39(3):1063-1077.

Bay A, Sidiropoulos P, Vazquez E, Sasdelli M. Hide and Seek tracker: Real-time recovery from target loss. CoRR 2018;abs/1806.07844.

Bozek J, Makropoulos A, Schuh A, Fitzgibbon S, Wright R, Glasser MF, Coalson TS, O'Muircheartaigh J, Hutter J, Price AN, Cordero-Grande L, Teixeira RPAG, Hughes E, Tusor N, Baruteau KP, Rutherford MA, Edwards AD, Hajnal JV, Smith SM, Rueckert D, Jenkinson M, Robinson EC. Construction of a neonatal cortical surface atlas using Multimodal Surface Matching in the Developing Human Connectome Project. Neuroimage 2018;179:11-29.

Bustos A, Chin T. Guaranteed Outlier Removal for Point Cloud Registration with Correspondences. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018;40(12):2868-2882.

Cao Y, Wu Z, Shen C. Estimating Depth from Monocular Images as Classification Using Deep Fully Convolutional Residual Networks. IEEE Transactions on Circuits and Systems for Video Technology 2018;28(11):3174-3182.

Cao Y, Zhao T, Xian K, Shen C, Cao Z, Xu S. Monocular Depth Estimation with Augmented Ordinal Depth Relationships. IEEE Transactions on Image Processing 2018.

Carneiro G, Tavares J, Bradley A, Papa J, Belagiannis V, Nascimento J, Lu Z, Conjeti S. DLMIA 2018 Preface. In: Stoyanov D, Taylor Z, Carneiro G, Syeda-Mahmood T, editors. Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018 and 8th International Workshop, ML-CDS 2018 Held in Conjunction with MICCAI 2018 Granada, Spain, September 20, 2018 Proceedings; 2018. p VII-VII.

Carneiro G, Tavares J, Bradley A, Papa J, Nascimento J, Cardoso J, Lu Z, Belagiannis V. 1st MICCAI workshop on deep learning in medical image analysis. Computer Methods in Biomechanics and Biomedical Engineering: Imaging and Visualization 2018;6(3):241-242.

Chance SA, Torso M, Jenkinson M. Cortical Disarray Measurement Using Diffusion Mri Correlates with Braak Staging and Improves Amyloid Pet Prediction of Ad Progression. Alzheimer's & dementia : the journal of the Alzheimer's Association 2018;14(7):P1245.

Chen B, Ting KM. Neighbourhood Contrast: A Better Means to Detect Clusters Than Density. Advances in Knowledge Discovery and Data Mining; 2018. p 401-412.

Chen B, Ting KM, Washio T, Zhu Y. Local contrast as an effective means to robust clustering against varying densities. Machine Learning 2018;107(8-10):1621-1645.

Chojnacki W, Szpak ZL. Determining ellipses from low resolution images with a comprehensive image formation model. CoRR 2018;abs/1807.06814.

Chojnacki W, Szpak ZL. Full explicit consistency constraints in uncalibrated multiple homography estimation. CoRR 2018;abs/1805.02352.

Connally EL, Ward D, Pliatsikas C, Finnegan S, Jenkinson M, Boyles R, Watkins KE. Separation of trait and state in stuttering. Human Brain Mapping 2018;39(8): 3109-3126.

Do T, Cheung N. Embedding based on function approximation for large scale image search. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018;40(3):626-638.

Gale W, Oakden-Rayner L, Carneiro G, Bradley AP, Palmer LJ. Producing radiologist-quality reports for interpretable artificial intelligence. CoRR 2018;abs/1806.00340.

Gong D, Tan M, Shi Q, van den Hengel A, Zhang Y. MPTV: Matching Pursuit Based Total Variation Minimization for Image Deconvolution. CoRR 2018;abs/1810.05438.

Gong D, Zhang Z, Shi Q, van den Hengel A, Shen C, Zhang Y. Learning an Optimizer for Image Deconvolution. CoRR 2018;abs/1804.03368.

Griffanti L, Jenkinson M, Suri S, Zsoldos E, Mahmood A, Filippini N, Sexton CE, Topiwala A, Allan C, Kivimaki M, Singh-Manoux A, Ebnaeier KP, Mackay CE, Zamboni G. Classification and characterization of periventricular and deep white matter hyperintensities on MRI: A study in older adults. Neuroimage 2018;170:174-181.

Guo G, Wang H, Shen C, Yan Y, Liao H. Automatic Image Cropping for Visual Aesthetic Enhancement Using Deep Neural Networks and Cascaded Regression. IEEE Transactions on Multimedia 2018;20(8):2073-2085.

Harston GWJ, Carone D, Sheerin F, Jenkinson M, Kennedy J. Quantifying Infarct Growth and Secondary Injury Volumes: Comparing Multimodal Image Registration Measures. Stroke 2018;49(7):1647-1655.

Hodgson J, Mott R, Baylis S, Pham T, Wotherspoon S, Kilpatrick A, Raja Segaran R, Reid I, Terauds A, Koh L. Drones count wildlife more accurately and precisely than humans. Methods in Ecology and Evolution 2018;9(5):1160-1167.

Hu Q, Wang P, Shen C, van den Hengel A, Porikli F. Pushing the Limits of Deep CNNs for Pedestrian Detection. IEEE Trans Circuits Syst Video Techn 2018;28(6): 1358-1368.

Li H, Wang P, Shen C. Toward End-to-End Car License Plate Detection and Recognition With Deep Neural Networks. IEEE Transactions on Intelligent Transportation Systems 2018.

Li H, Wang P, Shen C, van den Hengel A. Visual Question Answering as Reading Comprehension. CoRR 2018;abs/1811.11903.

Li H, Wang P, You M, Shen C. Reading car license plates using deep neural networks. Image and Vision Computing 2018;72(C):14-23.

Li X, Zhao L, Ji W, Wu Y, Wu F, Yang M, Tao D, Reid I. Multi-task structure-aware context modeling for robust keypoint-based object tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018;OnlinePubl:1-13.

Liao Z, Drummond T, Reid I, Carneiro G. Approximate Fisher Information Matrix to Characterise the Training of Deep Neural Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018.

Lin G, Shen C, van den Hengel A, Reid ID. Exploring Context with Deep Structured Models for Semantic Segmentation. IEEE Trans Pattern Anal Mach Intell 2018;40(6):1352-1366.

Liu F, Lin G, Qiao R, Shen C. Structured Learning of Tree Potentials in CRF for Image Segmentation. IEEE Transactions on Neural Networks and Learning Systems 2018;29(6):2631-2637.

Liu H, Ji R, Wang J, Shen C. Ordinal Constraint Binary Coding for Approximate Nearest Neighbor Search. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018:1-1.

Liu W, Zhang P, Chen X; Shen C; Huang X, Yang J. Embedding Bilateral Filter in Least Squares for Efficient Edge-preserving Image Smoothing. IEEE Transactions on Circuits and Systems for Video Technology 2018.

Lu H, Shen C, Cao Z, Xiao Y, Van Den Hengel A. An embarrassingly simple approach to visual domain adaptation. IEEE Transactions on Image Processing 2018;27(7):3403-3417.

Ma C, Shen C, Dick AR, Wu Q, Wang P, van den Hengel A, Reid ID. Visual Question Answering With Memory-Augmented Networks. 2018. p 6975-6984.

Makropoulos A, Robinson EC, Schuh A, Wright R, Fitzgibbon S, Bozek J, Counsell SJ, Steinweg J, Vecchiato K, Passerat-Palmbach J, Lenz G, Mortari F, Tenev T, Duff EP, Bastiani M, Cordero-Grande L, Hughes E, Tusor N, Tournier JD, Hutter J, Price AN, Teixeira RPAG, Murgasova M, Victor S, Kelly C, Rutherford MA, Smith SM, Edwards AD, Hajnal JV, Jenkinson M, Rueckert D. The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. Neuroimage 2018; 173:88-112.

Nekrasov V, Dharmasiri T, Spek A, Drummond T, Shen C, Reid ID. Real-Time Joint Semantic Segmentation and Depth Estimation Using Asymmetric Annotations. CoRR 2018;abs/1809.04766.

Nekrasov V, Shen C, Reid ID. Diagnostics in Semantic Segmentation. CoRR 2018;abs/1809.10328.

Nguyen D, Vu N, Do T, Nguyen T, Yearwood J. Improving Chamfer template matching using image segmentation. IEEE Signal Processing Letters 2018;25(11):1635-1639.

Page A, Li H. Meal-sensing signaling pathways in functional dyspepsia. Frontiers in Systems Neuroscience 2018;12(ARTN 10).

Pallebage-Gamarallage M, Foxley S, Menke RAL, Huszar IN, Jenkinson M, Tendler BC, Wang CY, Jbabdi S, Turner MR, Miller KL, Ansorge O. Dissecting the pathobiology of altered MRI signal in amyotrophic lateral sclerosis: A post mortem whole brain sampling strategy for the integration of ultra-high-field MRI and quantitative neuropathology. Bmc Neuroscience 2018;19.

Robinson EC, Garcia K, Glasser MF, Chen ZD, Coalson TS, Makropoulos A, Bozek J, Wright R, Schuh A, Webster M, Hutter J, Price A, Grande LC, Hughes E, Tusor N, Bayly PV, Van Essen DC, Smith SM, Edwards AD, Hajnal J, Jenkinson M, Glocker B, Rueckert D. Multimodal surface matching with higher-order smoothness constraints. Neuroimage 2018;167:453-465.

Saillant A, Platsch G, Shah V, Zuehlsdorff S, Jenkinson M, Chappell M. Reliability of myocardial blood flow estimates in abnormal scans. Journal of Nuclear Medicine 2018;59.

Schmid AB, Campbell J, Hurley SA, Jbabdi S, Andersson JL, Jenkinson M, Bangerter NK, Bennett DL, Tracey I, Frost R, Clare S. Feasibility of Diffusion Tensor and Morphologic Imaging of Peripheral Nerves at Ultra-High Field Strength. Investigative Radiology 2018;53(12):705-713.

Sheng B, Shen C, Lin G, Li J, Yang W, Sun C. Crowd counting via weighted VLAD on a dense attribute feature map. IEEE Transactions on Circuits and Systems for Video Technology 2018;28(8):1788-1797.

Snaauw G, Gong D, Maicas G, van den Hengel A, Niessen WJ, Verjans J, Carneiro G. End-to-End Diagnosis and Segmentation Learning from Cardiac Magnetic Resonance Imaging. CoRR 2018;abs/1810.10117.

Teney D, Anderson P, He X, van den Hengel A. Tips and Tricks for Visual Question Answering: Learnings From the 2017 Challenge. 2018. p 4223-4232.

Teney D, van den Hengel A. Visual Question Answering as a Meta Learning Task. ECCV 2018. p 229-245.

Wang L. Image and Sentence Matching via Semantic Concepts and Order Learning. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018.

Wang P, Wu Q, Cao J, Shen C, Gao L, van den Hengel A. Neighbourhood Watch: Referring Expression Comprehension via Language-guided Graph Attention Networks. CoRR 2018;abs/1812.04794.

Wang P, Wu Q, Shen C, Dick AR, van den Hengel A. FVQA: Fact-Based Visual Question Answering. IEEE Trans Pattern Anal Mach Intell 2018;40(10):2413-2427.

Wang X, Gao L, Wang P, Sun X, Liu X. Two-stream 3-D convNet Fusion for action recognition in videos with arbitrary size and length. IEEE Transactions on Multimedia 2018;20(3):634-644.

Wei X, Xie C, Wu J, Shen C. Mask-CNN: Localizing parts and selecting descriptors for fine-grained bird species categorization. Pattern Recognition 2018;76(C):704-714.

Williamson W, Lewandowski AJ, Forkert ND, Griffanti L, Okell TW, Betts J, Boardman H, Siepmann T, McKean D, Huckstep O, Francis JM, Neubauer S, Phellan R, Jenkinson M, Doherty A, Dawes H, Frangou E, Malamateniou C, Foster C, Leeson P. Association of Cardiovascular Risk Factors With MRI Indices of Cerebrovascular Structure and Function and White Matter Hyperintensities in Young Adults. Jama-Journal of the American Medical Association 2018;320(7):665-673.

Williamson W, Lewandowski AJ, Huckstep O, Visser E, Betts B, Jenkinson M, Dawes H, Foster C, Leeson P. Thalamic Volume is a Predictor of Cardiac Function at Rest and During Exercise in Young Adults. Circulation : journal of the American Heart Association 2018;138(Suppl_1):A16361.

Wu Q, Shen C, Wang P, Dick A, van den Hengel A. Image captioning and visual question answering based on attributes and external knowledge. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018;40(6):1367-1381.

Wu Q, Wang P, Shen C, Reid ID, van den Hengel A. Are You Talking to Me? Reasoned Visual Dialog Generation Through Adversarial Learning. 2018. p 6106-6115.

Xu W, Liu W, Huang X, Yang J, Qiu S. Multi-modal self-paced learning for image classification. Neurocomputing 2018;309(C):134-144.

Yao R, Lin G, Shen C, Zhang Y, Shi Q. Semantics-Aware Visual Object Tracking. IEEE Transactions on Circuits and Systems for Video Technology 2018;Online:1-14.

Yao R, Lin G, Shi Q, Ranasinghe D. Efficient dense labelling of human activity sequences from wearables using fully convolutional networks. Pattern Recognition 2018;78(C):252-266.

You M, Zhang Y, Shen C, Zhang X. An Extended Filtered Channel Framework for Pedestrian Detection. IEEE Transactions on Intelligent Transportation Systems 2018;19(5):1640-1651.

Zakizadeh R, Sasdelli M, Qian Y, Vazquez E. FineTag: Multi-label Retrieval of Attributes at Fine-grained Level in Images. CoRR 2018;abs/1806.07124.

Zakizadeh R, Sasdelli M, Qian Y, Vazquez E. Improving the Annotation of DeepFashion Images for Fine-grained Attribute Recognition. CoRR 2018;abs/1807.11674.

Zhang J, Wu Q, Shen C, Zhang J, Lu J. Multilabel Image Classification with Regional Latent Semantic Dependencies. IEEE Transactions on Multimedia 2018;20(10):2801-2813.

Zhang J, Wu Q, Shen C, Zhang J, Lu J, van den Hengel A, Goal-oriented visual question generation via intermediate rewards, ECCV 2018.

Zhang L, Wang P, Liu L, Shen C, Wei W, Zhang Y, van den Hengel A. Towards Effective Deep Embedding for Zero-Shot Learning. CoRR 2018;abs/1808.10075.

Zhang L, Wang P, Shen C, Liu L, Wei W, Zhang Y, van den Hengel A. Adaptive Importance Learning for Improving Lightweight Image Super-resolution Network. CoRR 2018;abs/1806.01576.

Zhang L, Wang P, Wei W, Lu H, Shen C, van den Hengel A, Zhang Y. Unsupervised domain adaptation using robust class-wise matching. IEEE Transactions on Circuits and Systems for Video Technology 2018;OnlinePubl:1-12.

Zhang L, Wei W, Zhang Y, Shen C, van den Hengel A, Shi Q. Cluster Sparsity Field: An Internal Hyperspectral Imagery Prior for Reconstruction. International Journal of Computer Vision 2018;126(8):797-821.

Zhou X, Liu Z, Gong C, Liu W. Improving video saliency detection via localized estimation and spatiotemporal refinement. IEEE Transactions on Multimedia 2018;20(11):2993-3007.

Zhuang B, Wu Q, Shen C, Reid ID, van den Hengel A. HCVRD: A Benchmark for Large-Scale Human-Centered Visual Relationship Detection. 2018. p 7631-7638.

Zhuang B, Wu Q, Shen C, Reid ID, van den Hengel A. Parallel Attention: A Unified Framework for Visual Object Discovery Through Dialogs and Queries. 2018. p 4252-4261.

Zhuang N, Yan Y, Chen S, Wang H, Shen C. Multi-label learning based deep transfer neural network for facial attribute classification. Pattern Recognition 2018;80(C):225-240.

Zsoldos E, Filippini N, Mahmood A, Mackay CE, Singh-Manoux A, Kivimaki M, Jenkinson M, Ebmeier KP. Allostatic load as a predictor of grey matter volume and white matter integrity in old age: The Whitehall II MRI study. Scientific Reports 2018;8.

Zulqarnain Gilani S, Mian A, Shafait F, Reid I. Dense 3D Face Correspondence. IEEE Transactions on Pattern Analysis and Machine Intelligence 2018;40(7):1584-1598.

FOR FURTHER ENQUIRIES

The University of Adelaide SA 5005 Australia

ENQUIRIES aiml@adelaide.edu.au

adelaide.edu.au/aiml

twitter.com/TheAIML

SOUTH AUSTRALIA