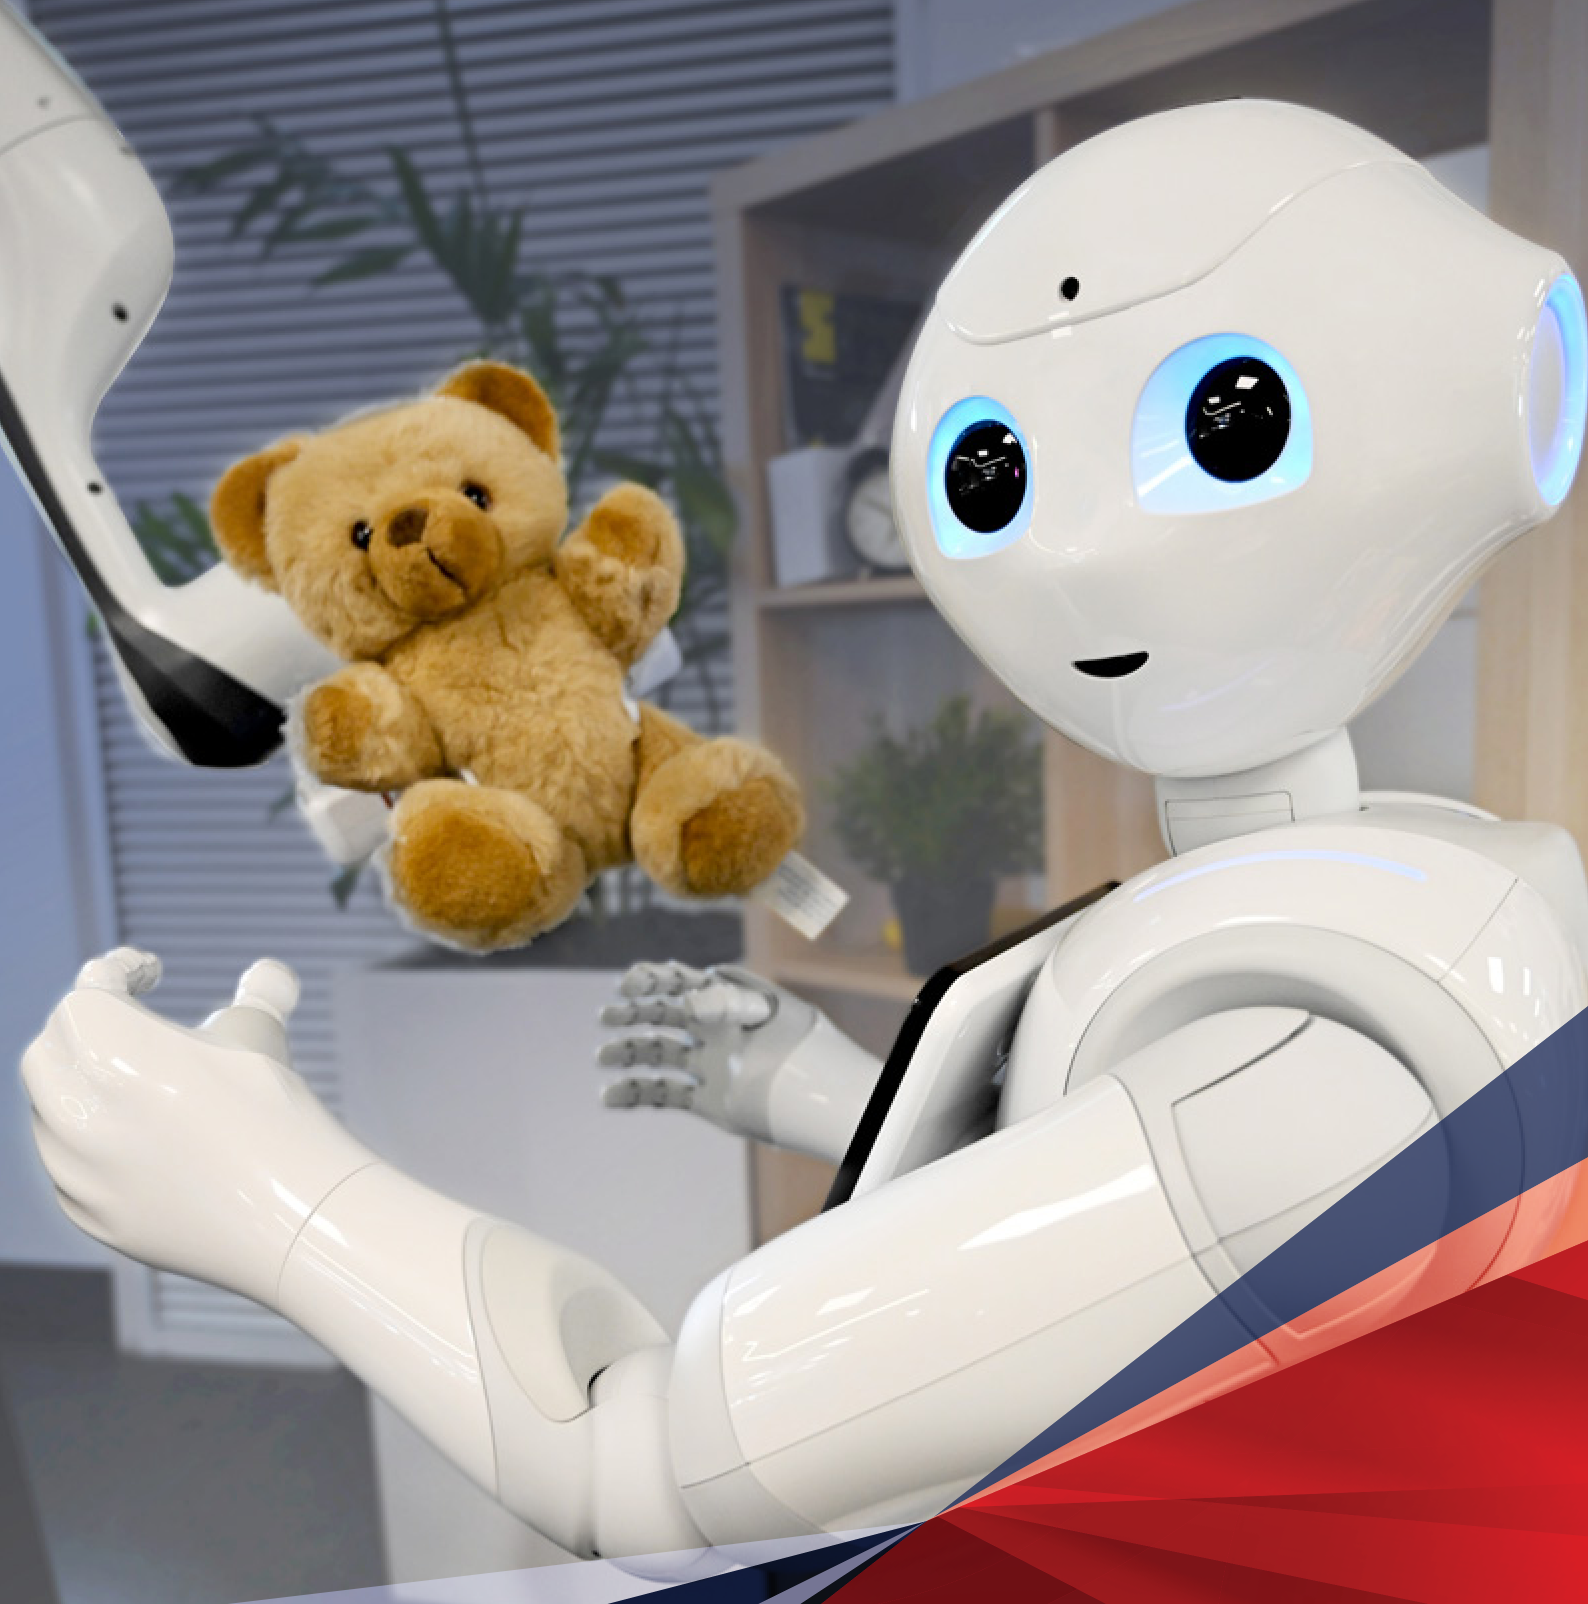

AUSTRALIAN CENTRE FOR ROBOTIC VISION

ANNUAL REPORT 2018

AN ARC CENTRE OF EXCELLENCE

Australian Government
Australian Research Council

We invent new things;
are open to crazy new ideas; and
encourage new ventures

Who we are

At the Australian Centre for Robotic Vision (the 'Centre') we strive to lead the world in the new discipline of robotic vision. Funded for $25.6 million over seven years, we are the largest expert body of our kind on the planet, bringing together the disciplines of robotics and computer vision.

Our formation recognises the fact that the breakthrough science and technologies needed to create a new generation of robots can only be achieved through concerted, large-scale and collaborative effort.

In answer to this, the Centre brings together more than 200 researchers from across Australia and the world. We run the world's biggest and most collaborative university-based Lab where, uniquely, robotics and computer vision experts work side by side. Our Lab spans four Australian universities: QUT, The University of Adelaide, The Australian National University (ANU) and Monash University. The Centre's interdisciplinary team further extends to include CSIRO's Data61 and overseas universities and research organisations including the French national research institute for the digital sciences (INRIA), Georgia Institute of Technology, Imperial College London, the Swiss Federal Institute of Technology Zürich (ETH Zürich), and the University of Oxford.

PHOTO by Anthony Weate

PHOTO by Anthony Weate

WORLD MAP HIGHLIGHTING PARTNER LOCATIONS

Our formation as an
ARC Centre of Excellence

The Australian Research Council (ARC) is a Commonwealth entity within the Australian Government, established as an independent body under the *Australian Research Council Act 2001.*

ARC Centres of Excellence, funded under the National Competitive Grants Program, are prestigious foci of expertise through which high-quality researchers collaboratively maintain and develop Australia's international standing in research areas of national priority.

Importantly, they facilitate strategic collaboration between universities, publicly-funded research organisations, other research bodies, governments and businesses in Australia and overseas.

As an ARC Centre of Excellence, we strive to deliver cultural, economic, social and environmental benefits to all Australians.

Distinguished Professor Peter Corke (QUT) and Professor Robert Mahony (ANU) identified the opportunity to bring the disciplines of robotics and computer vision together when the ARC invited applications in 2012 for new Centres of Excellence. After finalising our collaborative partner agreements in mid-2014, the Australian Centre for Robotic Vision was born. We formally launched on 9 March 2015.

VISION

Creating robots that see and understand for the sustainable wellbeing of people and the environments they live in.

MISSION

To develop new robotic vision technologies to expand the capabilities of robots.

VALUES

Our values represent our culture and the way we do things:

CREATE
We invent new things, are open to crazy new ideas and encourage new ventures

COLLABORATE
We work together and partner to solve grand challenges

EMPOWER
We energise, motivate and support our people to be knowledge leaders

IMPACT
We make a difference by applying our transformational research and turning our ideas into reality

About this Report

Our annual report covers the 2018 calendar year activities of the ARC Centre of Excellence for Robotic Vision.

It forms part of our official reporting and accounting requirements to the ARC. Activities encompass research, training, outreach, industry engagement, operations and finance.

Contents

Section 1 **About Us**

Centre **Performance**

Robots are vital to our future prosperity. However, without 'seeing' and 'understanding' capabilities approaching that of humans, a vast array of potential applications are closed to them. Our ultimate goal is for robotic vision to become the sensing modality of choice, opening the door to affordable real-world robots. We are developing new technologies so robots can see, understand and adapt to complex, unstructured and dynamically changing environments much like humans do.

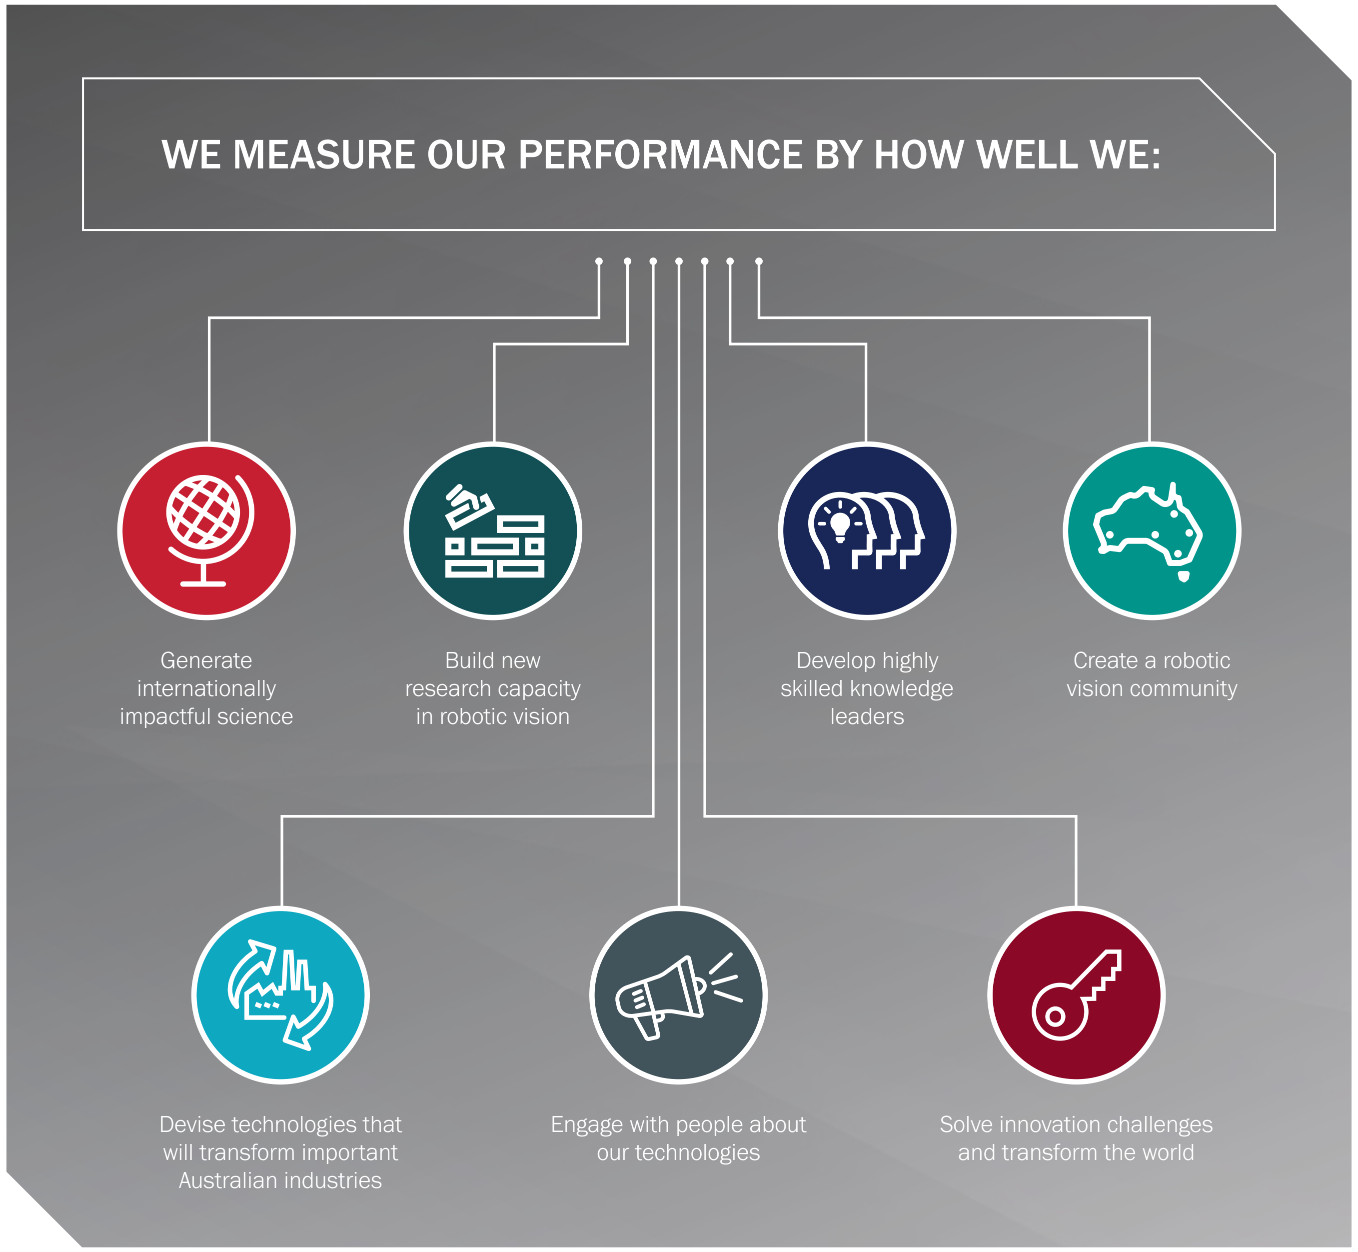

At a glance: **2018 achievements**

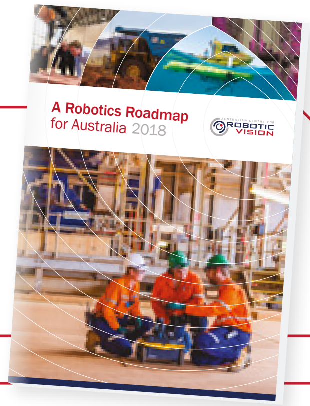

Released Australia's first Robotics Roadmap as a lasting Centre legacy and 'living guide' to how Australia can harness the benefits of a new robot economy. It forecasts robotics and automation will create, not replace jobs, and boost national productivity by $2.2 trillion over 15 years. See page 83

Brought the IEEE Robotics and Automation Society's flagship conference, the International Conference on Robotics and Automation (ICRA) to Australia for the first time. A coup largely thanks to the work of Centre Director Peter Corke and Advisory Board Chair Professor Alex Zelinsky. See page 98

Supported an Australia-first trial of SoftBank's popular social robot, Pepper in a real-world hospital setting. See page 54

Developed a 3D-printed camera system with nine cameras operating at different depths, enabling a harvesting robot (Harvey) to look around obstructing leaves much like a human. Centre researchers are separately developing an asparagus picking robot. See page 80

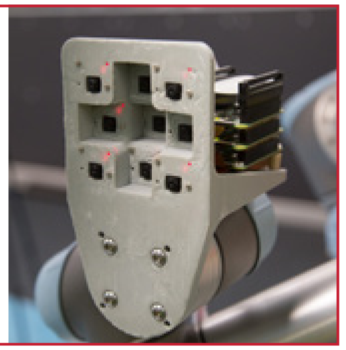

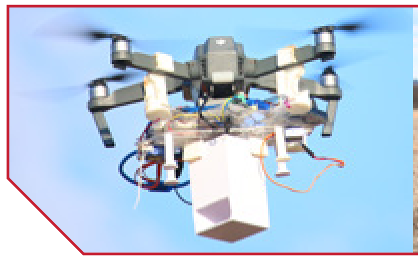

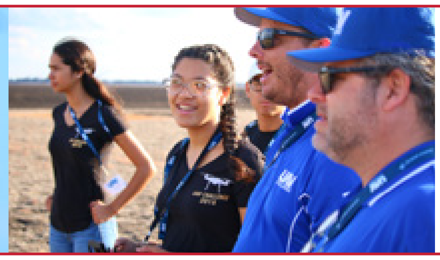

Organised the world's biggest outdoor flying robot challenge in partnership with QUT and CSIRO's Data61. See page 68

Assisted development of a smart phone app to help Australians self-diagnose skin cancer at the earliest possible stage. See page 70

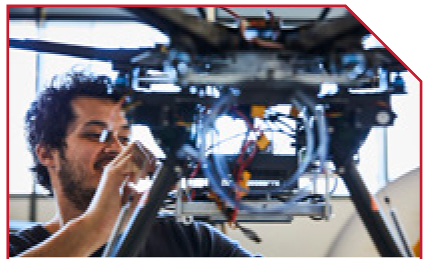

Demonstrated a prototype vision and sensing platform for unmanned aerial vehicles (UAVs) to help maintain a network of power lines as part of a FrontierSI (formerly CRCSI) and QUT collaboration with Ergon Energy. See page 78

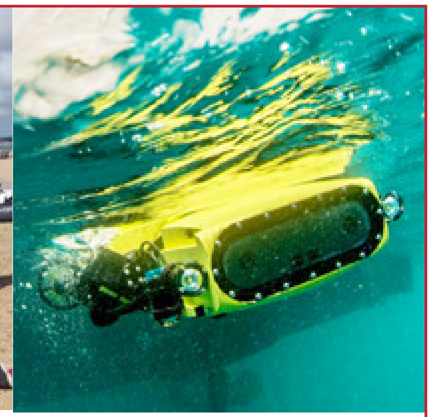

Centre-led student team came second in the coveted biennial Maritime RobotX Challenge in Hawaii. See page 74

Celebrated the first anniversary and half a million page views of Centre Director Peter Corke's QUT Robot Academy, an open online robotics education resource. See page 111

Staged the fourth successful Robotic Vision Summer School in the NSW beach hamlet of Kioloa. See page 111

Witnessed robotic vision as the superpower behind RangerBot, the world's first marine robot designed specifically for coral reef environments, put to work in a coral seeding project on the Great Barrier Reef. See page 72

Director's **Report**

All around the world there's a buzz of excitement, with new products and investment in the broad areas of robotics, computer vision and Artificial Intelligence (AI).

The SpotMini robot from Boston Dynamics took the world by storm with its slick dog-like moves. There's no question it's a truly impressive robot, but what is often overlooked – and key to its ability – is its incredibly powerful vision system. John Deere, a very traditional manufacturer, completed its buyout of Blue River Technology which is a robotic vision start-up. The RoboticsClub competition in Queensland brought teams of students, some international, together to make a motley collection of junk yard cars self-driving – all in just a weekend – by leveraging a lot of commodity hardware and open-source software.

It seems that what was once the province of massive technology giants is now available to almost everybody. These are exciting times as we strive to create new robotic vision technologies to further expand the capabilities of 'truly useful' robots.

SCIENCE

Deep learning is transforming the capability of robots to make sense of visual data and the Centre continues to drive this fundamental science. Performance gains continue to impress, but we know that the results are obtained using images drawn from standard databases. How well does that performance transfer to the real world? We've read articles about how easily deep networks can be spoofed or tricked; is this a problem in practice? To help answer this question we've launched a new project to evaluate the performance of robotic vision systems.

Our other science projects – **Learning; Manipulation and Vision; Scene Understanding; Vision and Language; Robots, Humans and Action; Fast Visual Motion Control; and Robotic Vision, Evaluation and Benchmarking** – all continue apace and feed into our demonstrator projects in Self-driving Cars and Manipulation. In particular, our work in Learning is reaping dividends. In the Scene Understanding, Language, and Manipulation and Vision projects: our world-leading algorithms in depth prediction and semantic segmentation have formed important building blocks in Scene Understanding for learning scene structure and camera motion; likewise reinforcement learning is a core element of our novel approach to high-level robot navigation (in the Vision and Language project) and in learning to grasp (Manipulation and Vision project).

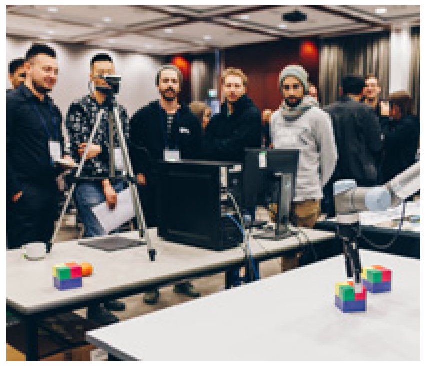

CULTURE

The Centre has two annual events that are key to building our culture, namely the Robotic Vision Summer School (RVSS) and our annual symposium, RoboVis.

Our fourth RVSS was held at the Australian National University (ANU) coastal campus set in the picturesque southern New South Wales beach hamlet of Kioloa. Becoming ever more popular each year, the 2018 event drew 50 students (16 international and 34 domestic) alongside four invited international experts who added to the unique and intensive week of lectures, technical deep dives and a hands-on robotics workshop. For many students, the workshop provided the first opportunity for them to step from behind the screen, developing algorithms, to experiencing the exhilaration and, at times, inevitable frustration of working with actual robots.

RoboVis 2018 touched down in Canberra and was a great success, providing a welcome opportunity for the entire Centre cohort to come together: Chief Investigators, Associate Investigators, Professional Staff, Advisory Board members; Research Fellows and PhD Researchers. The event featured a 'Town Hall' discussion about Centre outcomes and the important question of legacy and welcomed guest speakers: Nobel Prize-winning astrophysicist and ANU Vice-Chancellor Brian Schmidt on gender diversity initiatives and the importance of accountability in a Centre of Excellence and inspiring talks by two amazing female scientists, Elanor Huntington (the first female Dean of Engineering and Computer Science at ANU and still one of the few in the world) and cultural anthropologist, technologist and futurist Genevieve Bell.

Demos were a particular highlight of the event, underscoring the Centre's 2018 introduction of a new class of research project called demonstrator projects. These projects focus on showcasing to end users, in an accessible and compelling way, what robotic vision is all about. Taking advantage of being in Australia's capital, RoboVis ended with a playful yet potentially serious debate in Old Parliament House about legislating responsibility for AI agents!

Preparing the next generation of robotic vision experts for diverse career pathways across academia and industry sits at the heart of our culture. As part of this, the Centre strives to provide a rich training experience for early career researchers. We offer a range of development opportunities including a bespoke Knowledge Leadership program developed specifically for our needs, with 16 workshops delivered to date.

We also encourage our students to take advantage of opportunities to travel to other international laboratories (including those aligned to the Centre) and global technology giants like Google Brain and Amazon Robotics to gain new insights and expertise in research approaches.

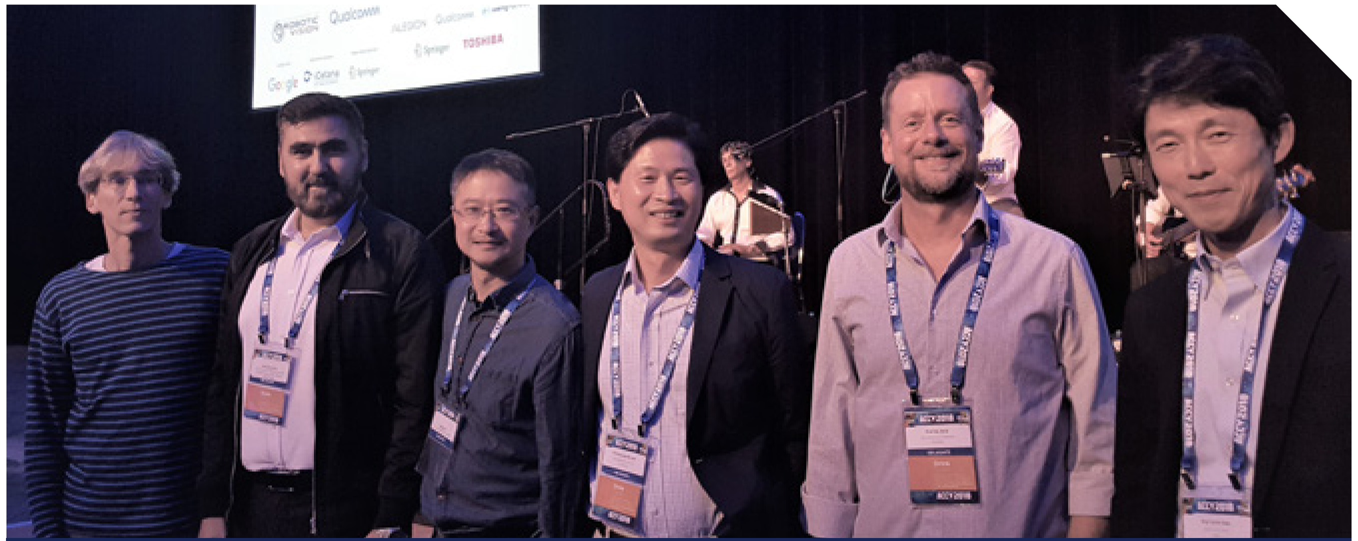

The ACCV 2018 organisers: Program Chair, Konrad Schindler (ETH Zürich); Local Organisation Chair, Ajmal Mian (University of Western Australia); Program Chair Hongdong Li (Centre/ANU); General Chair, Kyoung Mu Lee (Seoul National University); General Chair, Ian Reid (Centre/University of Adelaide); General Chair, Yoichi Sato (University of Tokyo)

INTEGRATE

A key goal of our Centre is integrating robotics and computer vision into the new field of robotic vision. We consider this problem at two levels: cultural and technological. At the cultural level we aim to create a new research community that overlaps, and draws from, the existing robotics and computer vision research communities. Like all communities, research communities have their own value systems. Computer vision values performance assessed against standard datasets while robotics values real-world experiments. Since its inception, the Centre has worked to connect these communities through annual robotic vision workshops associated with major annual robotics and computer vision conferences. This year we ran eight workshops at international conferences including the International Conference on Robotics and Automation (ICRA), the International Conference of Robots and Systems (IROS), Robotics Science & Systems (RSS), and the Asian Conference on Computer Vision (ACCV).

Building on our work this year to develop a unique global 'Challenge', we will launch several new robotic vision competitions in 2019 to bring these communities even closer together and cement the idea of robotic vision as a distinct topic. The competitions will allow robotic vision researchers to rigorously compare their results against ground truth, with the images drawn from cameras on a robot moving in a realistically rendered virtual world.

Our focus on technological integration is concerned with demonstrating robotic vision systems that can operate in the real world and perform useful tasks by linking visual perception of our world to robotic actions. Our Google-funded RangerBot – the world's first robotic vision-empowered reef protector – is an excellent example of the value of robots able to see and understand. RangerBot is not alone, as outlined in other groundbreaking case studies in this report that demonstrate tangible applications of robotic vision.

ENGAGE

The IEEE International Conference on Robotics and Automation (ICRA) is the world's largest general robotics conference, held annually since 1984. It travels around the globe but, until this year, has never ventured south of the Equator. In 2018, ICRA came to Brisbane; a coup for Australia and the Centre after four years' plotting. Records were broken for the number of papers presented (1,052) and attendees (3,000 delegates from 60 countries). As we prepare this annual report, the event which involved Centre researchers at every level, won the prestigious Australian Events Award for Best Congress or Conference of 2018.

In another 'win', the Asian Conference on Computer Vision (ACCV) returned to Australia for the second time in its history, spanning nearly three decades. The biennial event took place on the banks of the Swan River at Perth Conference and Exhibition Centre. Our Centre Deputy Director Ian Reid was ACCV General Chair, with Chief Invesigator Hongdong Li as Program Chair and Chief Invesigator Richard Hartley delivering the keynote speech at the conference opening. Like ICRA, we broke records for the number of papers submitted (979) and accepted (274) and delegates (550). Another great result for Australia and the Centre.

The Centre's reach is boosted via a number of online resources which continue to grow in popularity. In its first year, the open online robotics education resource, QUT Robot Academy (the brainchild of Centre Director Peter Corke) clocked up more than 70,000 international users and over 500,000 lesson views.

TRANSFORM

A key goal of our Centre is to effect positive transformation of society through tangible applications of robotic vision. In 2018, 'reef warrior', RangerBot, demonstrated the potential of robotic vision to transform management of precious marine environments, while our development of applications for SoftBank's humanoid robot, Pepper, used in an Australia-first hospital trial, underline the future value of social robots in health, mental health and aged care.

The Centre had its first foray into transformation at a national policy level, with the release of Australia's first Robotics Roadmap. This 212-page document is part of an audit of the current national capability and utilisation of robotic and computer vision technology. The Roadmap also contains a suite of policy recommendations. Designed as a 'living document' and Centre legacy, it forecasts robotics and automation will create, not replace jobs, and boost national productivity by $2.2 trillion over 15 years. It was launched at Parliament House, Canberra, on 18 June 2018.

We also need to periodically transform our Centre to ensure we are performing as well as we can and are adapting to changes in staffing and research trends. Over the course of this year, we have reduced the number of projects, favouring a smaller number of larger projects. We have also transformed the way we manage our projects – augmenting our Research Committee structure with twice-yearly, half-day reviews conducted by members of the Centre Executive. Finally, we have transformed our governance structure by combining our End User Advisory Board and Centre Advisory Committee into a single Centre Advisory Board.

FUTURE

In 2018, the Centre Executive Committee decided not to seek a second term as an ARC Centre of Excellence. While this means 2019 will be our penultimate year, we have no intention of slowing down. We have chosen to invest all our effort and resources into our current suite of projects. These are now well-established research projects with clear focus and positioned at the leading edge of important problems in robotic vision. Our job is to ensure they continue to be fully resourced, that interconnections between projects and Centre nodes continue to be made, and that the work is well disseminated to the global scientific community and general public.

A big focus for 2019 will be to define and articulate the Centre's legacy. Our legacy is more than the sum of individual scientific results. It is our people, connections made, and the founding of a new technical discipline – robotic vision.

We will continue to work with governments at all levels to address the issues uncovered in our Robotics Roadmap for Australia. And here, we acknowledge the work of our founding Centre Chief Operating Officer, Dr Sue Keay, who left the centre at the end of 2018. Sue has largely driven this work, helping to shape the national agenda, in the process drawing global recognition as a powerful advocate for transformational change and champion for diversity in all its forms. She has done much to build our ARC Centre of Excellence from scratch, its internal organisation and outward profile. She will be missed and we thank her for her significant contribution.

Peter & Ian

This Director's Report is jointly authored after Ian Reid took the reins as Acting Centre Director during Peter Corke's four-month sabbatical in the US (June-September 2018)

What is a **robot?**

A robot is a goal-oriented machine that can sense its environment, plan and act. Importantly, it is through sensors such as vision that robots understand their world and plan an action to achieve a goal.

Vision first evolved on planet Earth about 540 million years ago. It is the dominant sense in humans and also a game changer for robots – at least those working in uncontrolled real-world environments outside the closed world of the factory floor.

Robots have four main features: mobility, interactivity, communication and autonomy. In general we use the word 'robotics' to encompass ALL robotics-relevant fields such as computer and machine vision, artificial intelligence (AI) and machine learning as well as automation and autonomous systems.

Robots are often described in terms of two classes of robots – industrial or service – depending on their intended application. Industrial robots are used in industrial automation applications while service robots are not. Service robots may be for personal/domestic use or for use in professional settings (e.g., concierge robots in hotels).

Statistics on the number of robots produced in the world each year are divided into these two broad categories (industrial and service robots) by the International Federation of Robotics. The organisation's latest World Robotics Report shows a record high of 381,000 industrial robots, valued at US$16.2 billion, were sold in 2017 (up 30 per cent on 2016). Over the same period 8.5 million personal/domestic service robots and 109,500 professional service robots were sold (up 25 per cent and 85 per cent respectively).

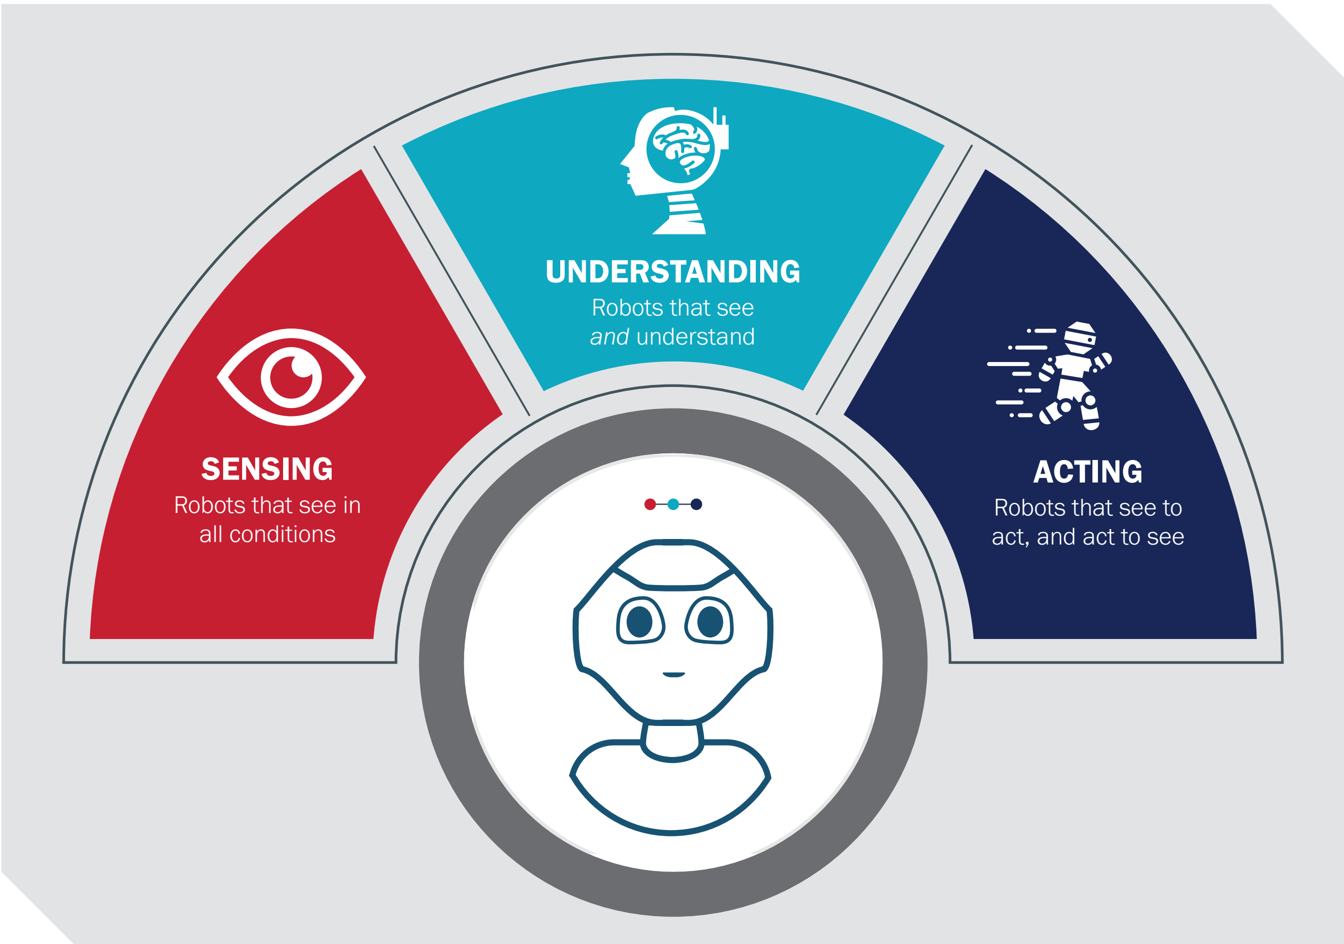

Meet our robots

The Centre is fortunate to have access to the following robotic platforms to advance our research

SoftBank's humanoid robot, Pepper ▸

Harvey, the robotic capsicum harvester; with work underway to develop an asparagus picking robot

Centre-built robots including our custom-built Cartesian robot, Cartman, which aced the 2017 Amazon Robotics Challenge in Japan, and the Gen3 miniature autonomous vehicle demonstrator

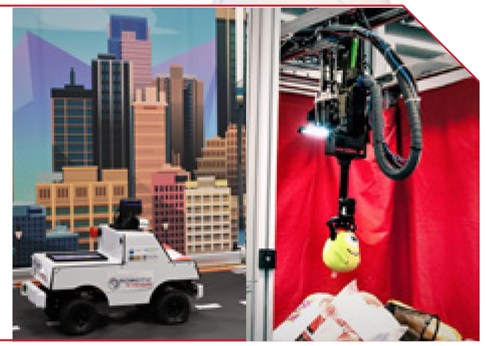

Powerful medical robotic platforms, including the RAVEN, for eye surgery research, and Snakebot, for robot-assisted arthroscopy

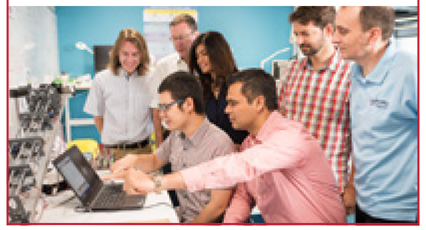

Vision-empowered marine robots COTSbot and RangerBot

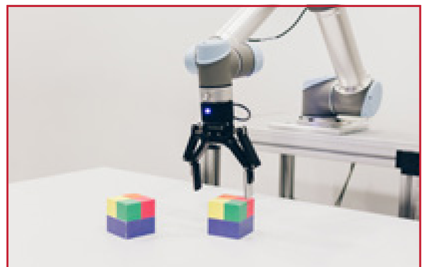

Robotic arms including Franka Emika's Panda, Kinova's Mico and Universal Robots' UR5

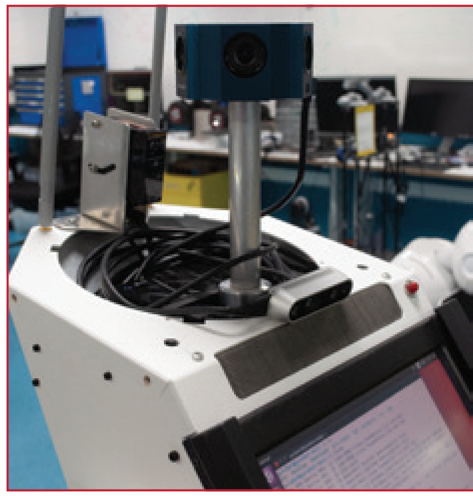

An Adept Guiabot mobile robot, used as a platform for the Centre's newest creation, BenchBot, which as its name suggests, is designed to help researchers without access to physical robots to remotely test code and benchmark research in the real world

What is vision and why is it important?

Vision or ‘seeing’ is an active process, described by early vision researcher David Marr as ‘a process of discovering from images *what* is present in the world and *where* it is’.

Robotic vision, like human vision, involves ‘seeing’ and ‘understanding’. It is a complex sense that extends beyond an understanding of the **where** and **what** of an object to its **why**. Namely, why an object is important in achieving a specific end goal and how it should be safely and effectively used.

For humans, everyday vision relies on an extraordinary range of abilities that most people take for granted.

Most adults have an ability to almost instantaneously understand what they see.

We can identify shapes; see colour; detect motion; gauge speed and distance; navigate; avoid obstructions; fill in blind spots; quickly correct distorted information; and even understand emotion.

Vision, quite simply, is an integral part of what makes us truly human. It is our dominant sense, involving our eyes and brain as equal partners.

In place of eyes, robotic vision involves primarily cameras as sensors.

Centre Director Peter Corke, in his QUT Robot Academy ‘Masterclass on Robotic Vision’, outlines two reasons why vision is an important and very practical sensor for robots:

1. Cameras are now very cheap because they are built into everything (mobile phones, laptops etc). What this means is the actual sensor – equivalent to the retina of a human eye – is now a device that perhaps costs less than a dollar.
2. Computation is also smarter and cheaper. We have access to very powerful computer chips with lots of memory, which allow researchers to run algorithms to process data that comes out of the sensor chip.

Visit **www.robotacademy.net.au** to access University-level, short video lessons and full online courses to help you understand and prepare for this technology of the future.

DID YOU KNOW?

At least a third of the human brain – our most complex organ, weighing around 1.4 kilograms and made up of about 86 billion neurons – is devoted directly or indirectly to visual processing. Further, 80 to 85 per cent of our perception, learning, cognition, and activities are mediated through vision.

How do robots ‘see’ the world?

Centre Chief Investigator Jonathan Roberts revisits this complex topic. The following article is modified from a piece he first penned for The Conversation not long after the Centre commenced operation.

So, how do robots ‘see’ the world? This is a simple question but the answer is complicated.

Seeing is a complex process and, in fact, scientists do not fully understand how we or other animals see. Science has shown how various parts of the seeing process in biology works, but given that at least a third of the human brain is involved in processing our sight, it is no wonder that animal vision research is still a very active scientific endeavour.

In the world of engineering and robotics, seeing starts with a stream of video images captured by cameras.

Since the late 1960s, researchers have thought about features in images. These might be lines, or interesting points like corners or certain textures. They write algorithms to find these features and track them from image frame to image frame in the video stream. This step is essentially reducing the amount of data from the millions of pixels in an image to a few hundred or thousand features.

In the recent past when computing power was limited, this was an essential step in the process. Researchers then think about what the robot is likely to see and what it will need to do. They write software that will recognise patterns in the world to help the robot understand what is around it.

In essence the robots are being programmed by a human to see things that a human thinks the robot is going to need to see. There have been many successful examples of this type of robot vision system. An example is the vision systems used in factories to sort or grade items as they move along a conveyor belt.

Continued over page ►

How do robots 'see' the world? Cont.

Did you know that much of the world's rice is now sorted by robotic machines that use vision to grade each grain of rice by size and colour? This is why you are likely to buy bags of rice with such uniform grains. It is robotic vision at work!

Sorting using robotic vision is one thing, but using vision for more complex tasks such as enabling a robot to move around the world is way more complicated. You will not find many robots today that can reliably navigate the world using vision alone. The lack of reliability is ultimately what holds back the use for robots in real-world applications.

In the past decade, a new approach to robotic vision has started to take shape. The breakthrough came in 2012 with deep neural networks: algorithms that learn from examples, much like humans do.

Researchers have developed robotic vision systems inspired by how scientists think animals see and think. That is, they use the concept of layers of neurons, somewhat like an animal brain. The researchers program the structure of the system and they provide the system with huge quantities of training images and video. The system then learns what is and what is not important in the images to achieve the desired outcome of identifying the relevant information in the images.

The recent success of robotic vision using deep learning can be partially attributed to the availability of significant computer power at a reasonable cost. Investment in these technologies is hence accelerating fast.

There are numerous applications for robots that can see. It's hard to think of a part of our life where such a robot could not help.

The first uses of robots that can see in complex environments are likely to be in industries that either have labour shortage issues, such as agriculture, or are inherently unattractive to humans and maybe hazardous.

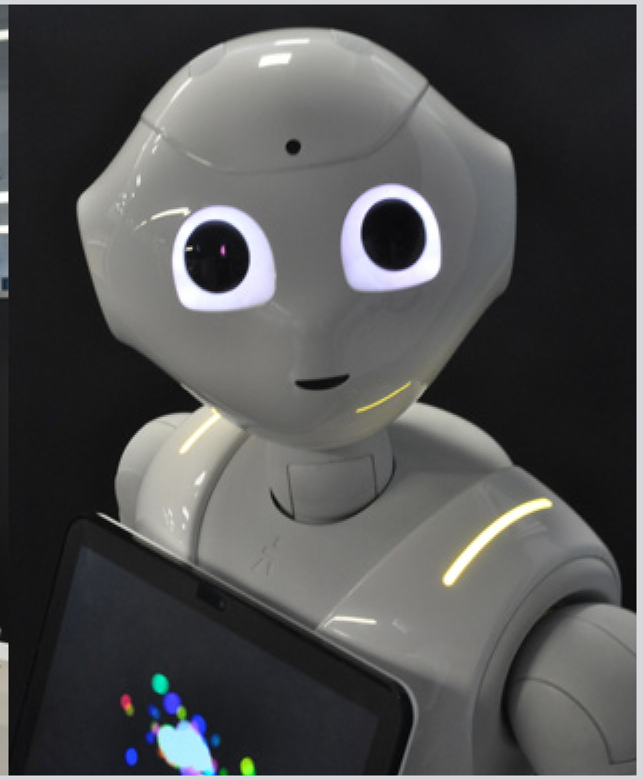

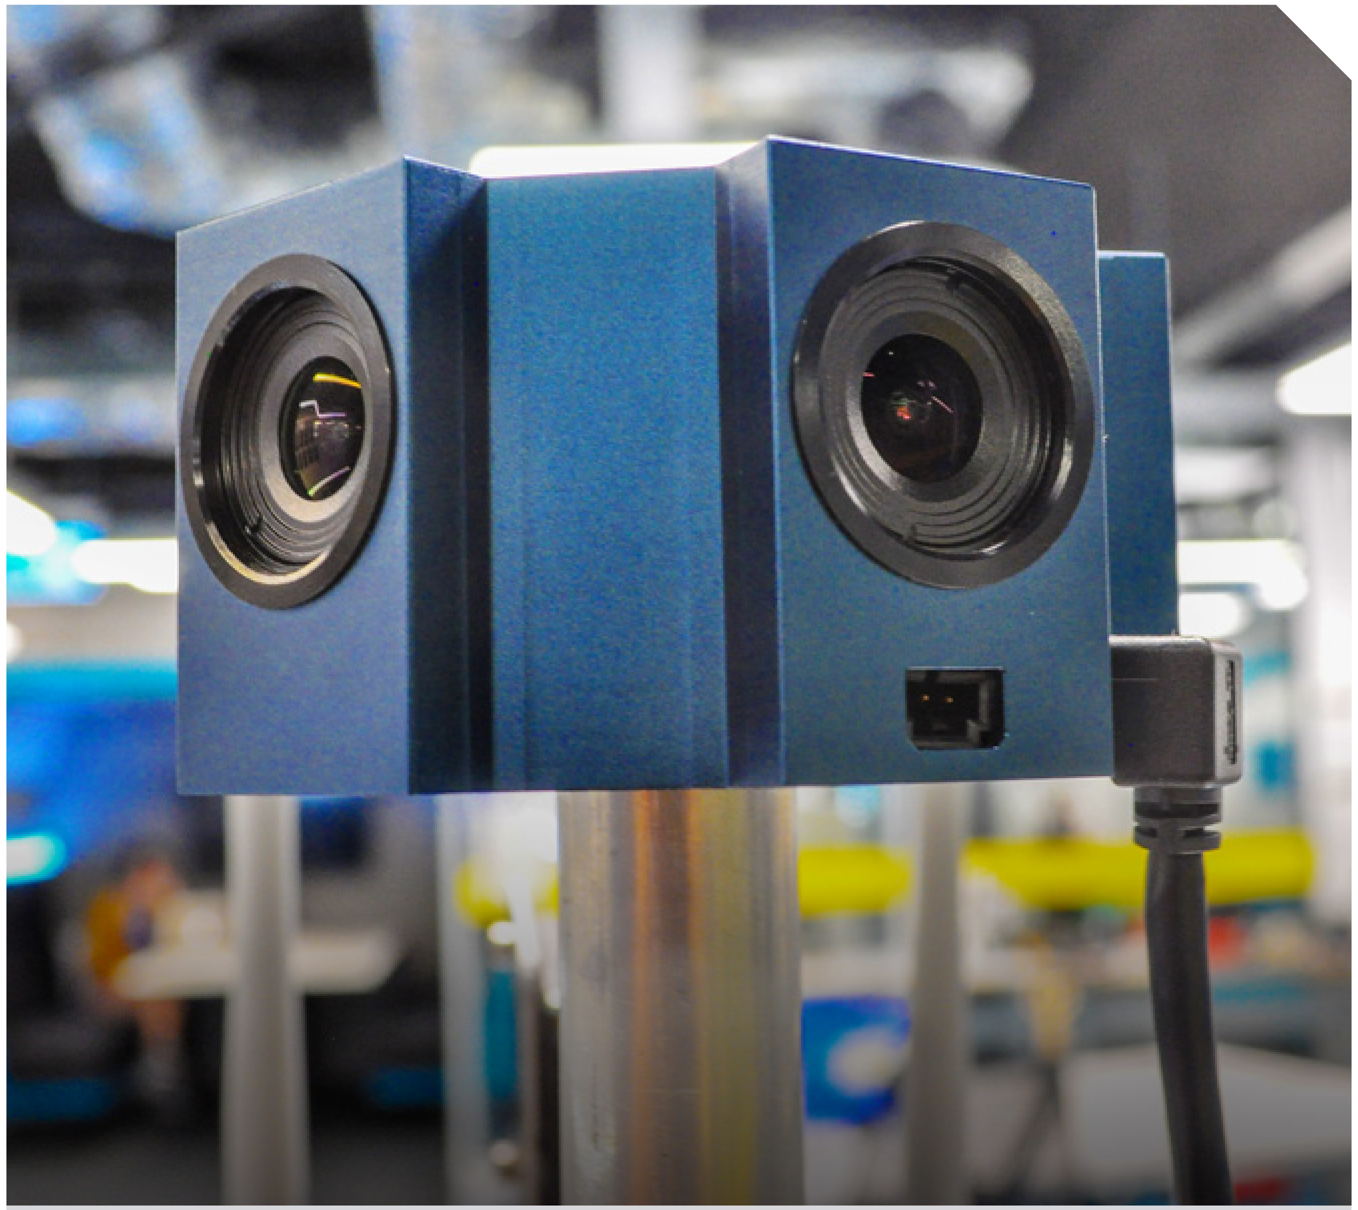

Tasks include searching through rubble after disasters, evacuating people from dangerous situations or working in confined and difficult to access spaces.

Applications that require a very long period of attention, something humans find hard, will also be ripe for a robot that can see. Our future home-based robot companions will be far more useful if they can see us.

And in an operating theatre near you, it is soon likely that a seeing robot will be assisting surgeons. The robot's superior vision and super precise and steady arms and hands will allow human surgeons to focus on what they are best at – deciding what to do.

Jonathan Roberts is a Centre Chief Investigator and Professor of Robotics at QUT. His main research interests are in the areas of Field Robotics and Medical and Healthcare Robotics. He is co-inventor of the UAV Challenge, the world's biggest (outdoor) airborne robotics challenge that doubles as a medical rescue mission (see page 68).

Section 2 **Research Performance**

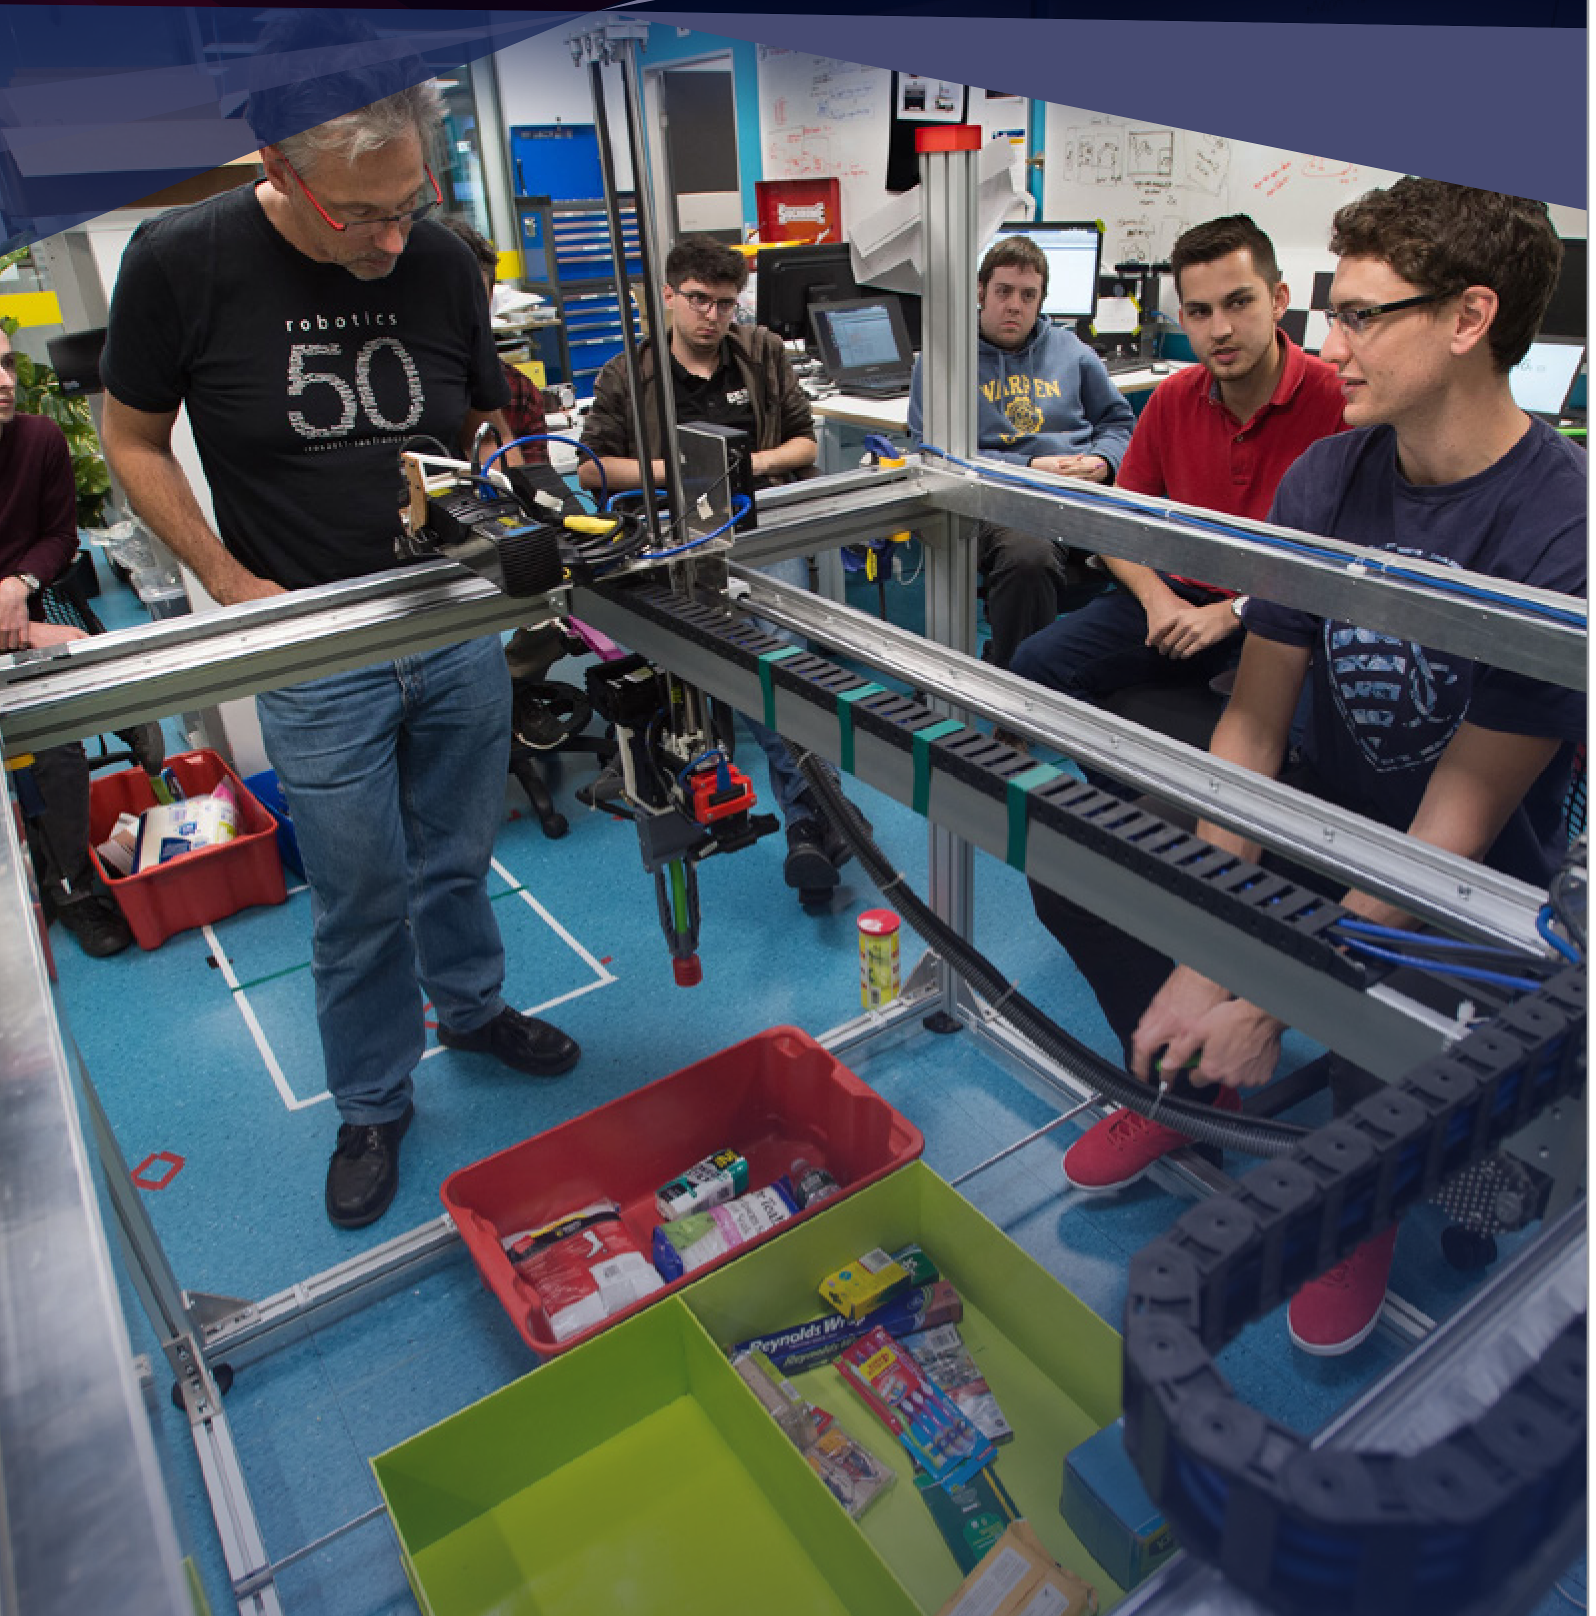

OUR VISION

Creating robots that see and understand for the sustainable wellbeing of people and the environments they live in.

Our ambitious research programs will be to develop technologies to harness the rich information from visual data to allow robots to perceive the world and be truly useful to humans. In 2018, the Centre sharpened its focus to seven research projects and the introduction of demonstrator projects as a new class of research project.

"The litmus test for any Centre project is whether it will develop a technology that showcases the capability of robotic vision that can be demonstrated by a robot, ensuring it will deliver an outcome that is relevant to our vision"

– **Centre Director Peter Corke**

We deliver our vision through five cross-connected research programs. Centre research projects are headed by a Chief Investigator or senior Research Fellow. A project team comprises a leader, with a team of Research Fellows **(RF)**, Associate Investigators **(AI)**, PhD researchers **(PhD)**, and Chief Investigators **(CI)**. Centre PhD projects are also included in our research programs. All new projects are established through our Research Committee, which determines duration, staffing and other resourcing matters.

SENSING
Robots that see in all conditions

UNDERSTANDING
Robots that see and understand

ACTING
Robots that see to act, and act to see

LEARNING
Robots that learn and improve

TECHNOLOGY
Robots that are fast and low cost

Sensing

For robots to be useful to humans, they need to act purposefully under an incredible range of adverse viewing conditions, including high glare, fog, smoke, dust, rain, sleet or snow. This program concentrates on the development of robotic vision algorithms and novel vision hardware to enable robots to see and act in all environments and conditions encountered in the real world. We are also working with innovative sensing hardware, developing it further to support robots operating under challenging viewing conditions such as poor light or through partial obscuration or at non-human-visible wavelengths.

PhD Research Project

Eventful invention! Why event cameras have the edge in our fast-paced world

One of the many challenges facing robotic vision researchers is how to interpret and understand motion quickly and robustly. A significant bottleneck in the visual processing pipeline is the computation of motion between two consecutive image frames.

"We hope to imbue some of the properties inspired by nature into our robots to enable collision avoidance in high-speed scenarios"

While this problem has largely been solved by computer vision researchers, state-of-the-art methods are typically computationally expensive, causing a delay which imposes a limit on how fast a robot can react to visual stimuli. Another limiting factor is the camera itself, which only captures images at fixed time intervals and is effectively 'blind' between image frames.

This is different to biological vision, which captures visual information continuously. To address this problem, researchers have invented biologically-inspired 'event cameras'.

Event cameras operate similarly to human eyes and do not suffer from delay. However, their output is very different to traditional cameras, meaning existing motion estimation algorithms cannot be directly applied.

Centre PhD Researcher, Cedric Scheerlinck (based at our ANU node in Canberra) is using event cameras to develop novel motion estimation algorithms for use in robotics. In September 2018, his research scored him a 12-month Swiss Government Excellence Scholarship, taking him to the University of Zürich.

Cedric's research aims to enable smarter and faster interactions between robots and their environment – important, for example, in driverless cars.

"Many animals use vision as their primary sense," Cedric said. "Fast flying creatures such as birds and insects have evolved to use lightning-fast motion cues to react to their environment. We hope to imbue some of the properties inspired by nature into our robots to enable collision avoidance in high-speed scenarios."

In place of images, event cameras output a sequence of electrical impulses called 'events' that represent local changes in brightness. For machines (or humans) to interpret this output, the first step is to reconstruct an image that summarises the information contained in these events.

Continued ▸

Cedric Scheerlinck in the Lab at the University of Zürich. He says: "Events are triggered by movement; local changes in brightness as the foreground passes over the background. Green indicates an increase in brightness and red represents a decrease."

The top line of images shows the output of a conventional camera, while the bottom shows an event camera real-time reconstruction.

Images reflect examples of challenging conditions including high dynamic range, low light or a fast-moving object such as a person running in front of car headlights in pitch black conditions. The conventional camera frame is delayed, under-exposed and motion-blurred compared to the real-time event camera reconstruction.

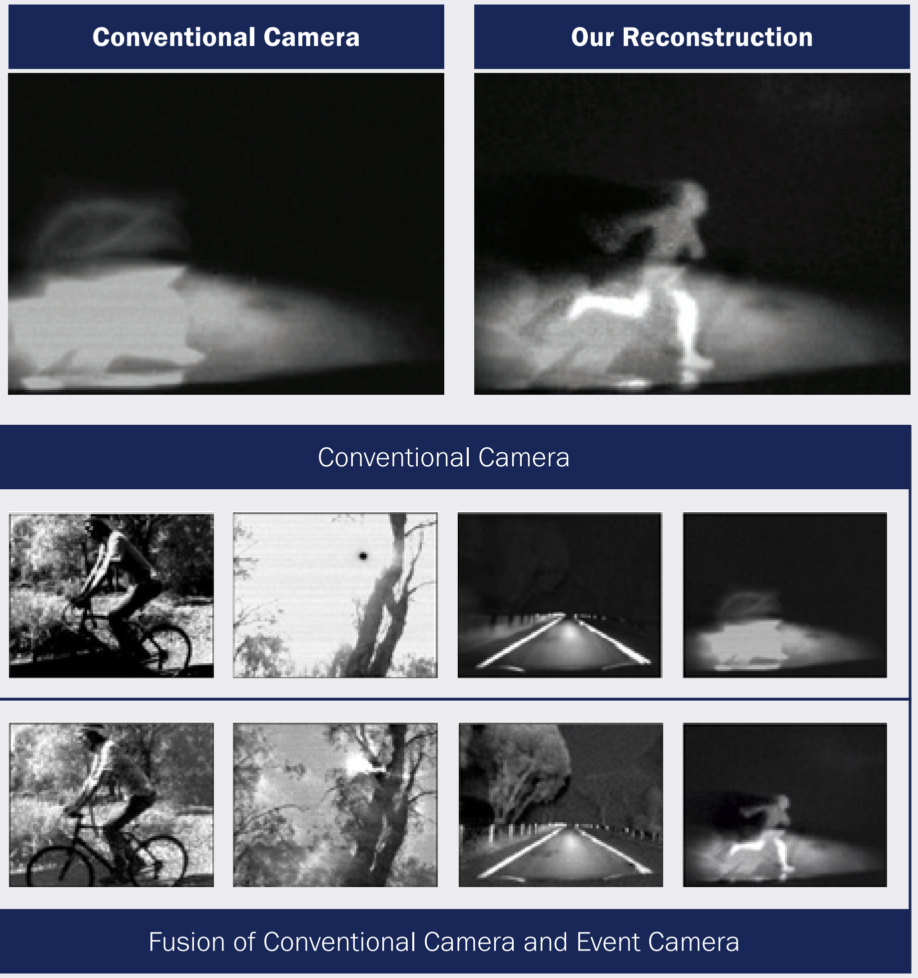

"To be usable in real-world robotic scenarios, this reconstruction needs to be performed as fast as possible to minimise the delay between reality and perception," said Cedric.

"We have achieved ground-breaking results in continuous-time image reconstruction and have released open-source code to encourage and support the research community.

"The ultimate objective is to achieve efficient (low-power, low-latency) perception, much like the capability of the human brain, to enable robots to extract meaningful information from the world around them."

Aside from driverless cars, event cameras will be beneficial in ongoing development of UAVs, particularly in search and rescue operations in mountain areas, forests or collapsed buildings.

This PhD research project is supervised by Centre Chief Investigator Rob Mahony.

EVENT CAMERAS: AT A GLANCE!

- / Event cameras are biologically inspired vision sensors.
- / Like the human eye, they capture visual information asynchronously and continuously without relying on a shutter to capture images.
- / Event cameras are much faster than conventional cameras and have microsecond temporal resolution.
- / They work well under a large range of illumination conditions, including low light and high dynamic range scenes.
- / Event cameras are not susceptible to motion blur; a big problem for conventional cameras especially in high-speed, low-light scenarios.

PhD Research Project

Why light field cameras have the edge

Robots for the real world will inevitably interact with refractive objects like wine glasses and other clear objects. The challenge for robotic vision researchers is to enable robots to detect refractive objects. For example, wine glasses and clear water bottles in domestic applications; glass and clear plastic packaging for quality assessment and packing in manufacturing; and water and ice in outdoor operations.

Refractive objects can cause many robotic vision algorithms to become unreliable or fail.

The key stumbling block lies in the fact that traditional robotic sensors, such as monocular/stereo/RGB-D cameras and LIDAR are unable to differentiate depth ambiguities introduced by refractive objects in a scene.

Enter, light field cameras! Unlike conventional cameras that only capture an image from a single perspective, light field cameras instantaneously capture multiple images of the same scene from slightly different perspectives.

For researchers, this better shows how objects are unevenly distorted by refractive objects as we change perspectives (encoded in a single light field image).

Centre PhD Researcher Dorian Tsai has developed an algorithm to distinguish between refracted and Lambertian (non-refracted) image features using light field cameras.

"Where previous methods were limited to large camera baselines relative to the refractive object, our method achieves comparable performance and further extends these capabilities to light field cameras with much smaller baselines than previously considered, while achieving much higher detection rates," Dorian said.

"We have also demonstrated that rejecting refracted features using our method yields more accurate position and orientation estimates when the robot is approaching refractive objects."

Looking ahead, Dorian's focus is on combining this method with previous work using light field cameras for visual servoing.

He said: "Our goal is to enable a robot to detect a glass of water and pick it up without relying on a prior model of the glass. Doing so would be a world-first and a critical step towards allowing robots to operate in close proximity with refractive objects."

This PhD research project is supervised by Centre Director Peter Corke, Associate Investigator Thierry Peynot and Research Affiliate Donald Dansereau, a postdoctoral scholar from Stanford Computational Imaging Lab, currently working as a senior lecturer at the University of Sydney. Of note, Dr Dansereau is the author of the open-source Light Field Toolbox for Matlab.

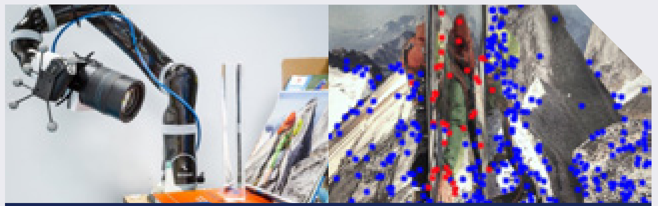

(Left) A light field camera mounted on a robot arm was used to distinguish refractive objects in a scene in structure from motion experiments. (Right) Image features that were distinguished as Lambertian (blue) and refracted (red), revealing the presence of the refractive cylinder in the middle of the scene.

LIGHT FIELD CAMERAS: AT A GLANCE!

- / Light field cameras are a new paradigm in imaging technology with the potential to greatly augment robotic vision. They can also be called plenoptic cameras.
- / The light field is the totality of all light rays in a 3D space, flowing through every point and in every direction.
- / Unlike conventional cameras that capture an image from a single perspective by recording each ray of light that passes through a point or 'pinhole' in space, light field cameras instantaneously capture multiple images of the same scene from slightly different perspectives.
- / Light field cameras still produce images through a single main lens, but record all light rays that pass through 'a volume' of space inside the camera using micro-arrays of lenslets attached to the CCD (charge-coupled device) sensor.
- / A robot can understand much more about the structure of a scene from a light field image.

Understanding

Understanding a scene enables a robot to interact with its environment by knowing what things are important to its task and how these things are behaving or are likely to behave. This program develops models, representations and learning algorithms that will allow robots to reason about their visual percepts, describe what they see and plan actions accordingly.

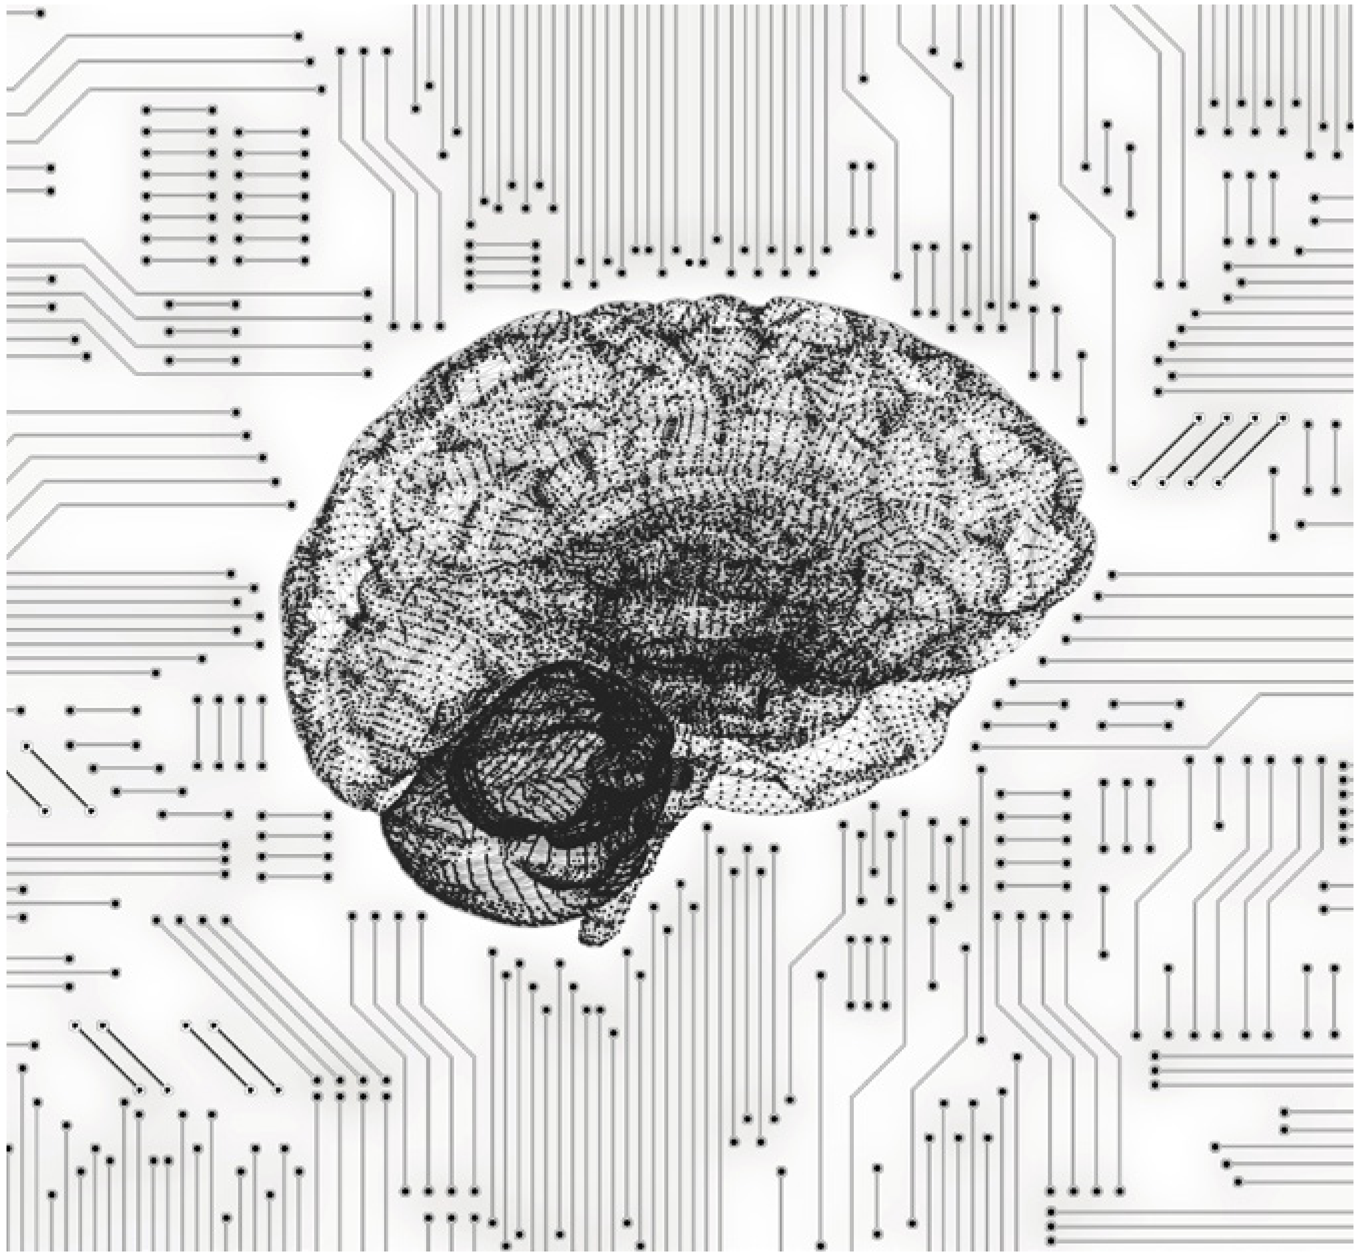

PROJECT

Scene Understanding

For robots to be truly useful in the real world, they need to recognise objects, including humans, and understand how they interact with each other. This project focuses on **geometric** and **semantic understanding** of a scene, including large-scale mapping. As a robot moves through its environment, geometric understanding enables it to build up an accurate picture of how far away every part of the scene is. The robot also requires a more human or semantic understanding of a scene, to interpret not just where bits of the scene are, but what they are — for example, what, in the scene, is a door, and what is a wall?

Project Leaders: Ian Reid (CI), Niko Sünderhauf (CI)

Team Members: Niko Sünderhauf (CI), Tat-Jun Chin (CI), Viorella Ila (RF), Andrew Spek (RF), Yasir Latif (RF), Saroj Weerasekera (RF), Huangying Zhan (PhD), Kejie Li (PhD), Lachlan Nicholson (PhD), Mina Henein (PhD), Jun Zhang (PhD), Natalie Jablonsky (PhD), Mehdi Hosseinzadeh (PhD), Sourav Garg (PhD), Ch'ng Shin Fang (PhD)

Project aim: Create geometric and semantic models and representations of the environment from visual observations. This ability is crucial for a robotic system to reason about a scene, its objects, their relationships – and their affordances (specific use or purpose that objects can have) – in order to plan actions effectively. Particular competencies to be developed include:

- semantically labelled 3D maps;
- object-based as opposed to point-cloud-based mapping of a scene;
- models for incorporating physically based rules, such as gravity and forces, for predicting the consequences of action; and
- the role of deep learning with geometry using 1000s or millions of observations of scenes to help create priors to resolve ambiguity.

KEY RESULTS:

Traditionally, SLAM (simultaneous, localisation and mapping) navigation systems have created maps based on isolated and usually sparse feature points. These systems excel at localising the camera, but produce maps that are not amenable to higher level interpretation or query. During 2018, the project team involving researchers from QUT, University of Adelaide and ANU nodes has been working towards building SLAM systems based on higher-order and potentially more semantically meaningful entities, involving geometric entities such as quadrics and planes, and even general objects. As a result, the project has developed systems that are more accurate and can handle dynamic objects more robustly and in a more natural way, also respecting natural meaningful relationships such as horizontal planar surfaces supporting objects in the scene.

Although localisation and mapping systems have traditionally been based on well-understood geometry, Centre researchers (especially at the University of Adelaide and Monash nodes) have been investigating how far a learning paradigm can succeed for SLAM. The project team has produced a system based on deep learning that can accurately predict the depth of a scene from a single camera; predict the optic flow between pairs of images in a video sequence; and use these predicted entities within a deep learning framework to estimate the pose difference. Combining all of these elements leads to state-of-the-art camera pose accuracy. Furthermore, the ability to predict the depth of a scene using a deep network alleviates the well-known scale-drift problem in SLAM because the network learns the implicit scene scale from typical sizes of objects. The team has also shown how this can all be achieved in a self-supervised manner, and how a single network can be used to predict depths, surface normals and a scene semantic segmentation.

Centre researchers have previously pioneered so-called single-view depth prediction using deep networks. Taking this further, the project explores the idea that the shape of an object can be predicted (using a network) from a single two-dimensional view. The team has shown how shape in the form of a dense point cloud can be accurately predicted. This has been incorporated into a SLAM system for self-driving vehicles that estimates the shape of cars in the scene as well as the dominant planes and point structure of the scene.

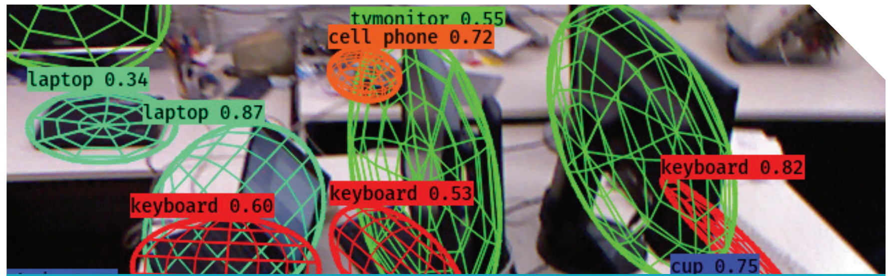

QuadricSLAM is a localisation and mapping system that represents objects using their coarse geometry – location and rough extent – via quadrics.

PROJECT

Robots, Humans and Action

Activities are a sequence of interactions between one or more agents (human or robot) with their environment in the pursuit of a goal. To fully understand activities, we need to model the dynamic relationship between an agent, its environment, objects in the environment, and other agents.

Project Leader: Stephen Gould (CI)

Team Members: Hongdong Li (CI), Richard Hartley (CI), Dylan Campbell (RF), Xin Yu (RF), Rodrigo Santa Cruz (PhD), Cristian Rodriguez (PhD), Amir Rahimi (PhD), Samira Kaviani (PhD), Zhen (Frederic) Zhang (PhD), Hamid Rezatofighi (AI), Miaomiao Liu (AI), Peter Koniusz (AI), Sadegh Aliakbarian (Associated PhD), Shihao Jiang (Associated PhD)

Project aim: Facilitate robot-human interaction and cooperation by developing models for robots to learn to recognise and describe activities from video. The project focuses on representations that allow a robot to monitor, understand, and predict both actions and intent of a human. Moreover, it investigates ways that understanding a visual environment can be used to predict the consequences of robot actions. This project links to other Centre research: *Scene Understanding* for detecting objects and modelling physical environments; *Vision and Language* for describing tasks and action sequences; *Learning* for improving underlying algorithms and models; and *Manipulation and Vision* for robot control and human-robot cooperation.

KEY RESULTS

The project team continued to innovate and demonstrate state-of-the-art performance on recognising human activity in images and videos from supervised training data. The models created allow robots to recognise up to 400 different activities in 10-second video clips. These activities, taken from the open source Kinetics database, cover everything from country line dancing to air drumming, applauding, arranging flowers, abseiling, kayaking, jogging, flipping pancakes, scrambling eggs and making tea. In 2018, research resulting from this work was published in top-tier computer vision conferences, including the Conference on Computer Vision and Pattern Recognition (CVPR) and the European Conference on Computer Vision (EECV).

The team developed techniques for encoding generic video sequences using unsupervised data. This gets around the limitation of needing large quantities of human labelled data for training deep learning models. For example, a recently compiled Kinetics dataset from Google required over 5000 human hours just to label the videos. Moreover, the dataset curators had to ensure that at least 400 examples of each activity category were sourced, without duplication, as well as performing quality control on the annotations. With an unsupervised approach the features can be trained on videos without any annotation and labels only needed for training the final classifier. With the push towards one-shot or low-shot learning we hope to then be able to classify videos with only a handful of annotated examples, resulting in a 100-fold reduction in labelling effort.

In 2018 early progress was made on developing models that anticipate human actions in the short-term future. This is important to facilitate future applications where robots need to react as early as possible, such as in autonomous driving, human-robot interaction and assistive robotics. This research is still in its infancy, but already the project team is able to successfully predict the evolution of an action by only observing a small portion of that action. For example, by watching a person walk, it is possible to anticipate the continued gait and cadence of their walk, or the extension of a hand as someone reaches for an object.

Moving beyond simple action classification, the team has developed models that recognise activities composed as a pattern of primitive actions. For example, holding a glass, followed by pouring water, followed by drinking from the glass.

> "In 2018, early progress was made on developing models that anticipate human actions in the short-term future. For example, by watching a person walk it is possible to anticipate the continued gait and cadence of their walk, or the extension of a hand as someone reaches for an object."

To further understand and describe activities in rich detail the project team developed state-of-the-art human pose tracking algorithms in 3D. This work, which appeared at CVPR 2018, sets the stage for understanding human-object interaction in the context of carrying out an activity, which, combined with intent forecasting, will be the main focus of our research in 2019. The key breakthrough here was the detection of recurrent motion patterns, which reduces the problem to the analysis of rigid motion, a departure from limited traditional low-rank shape models. The project's approach has been demonstrated on single-person walking and dancing sequences from a static camera and will be extended to the multi-person setting as well as moving cameras.

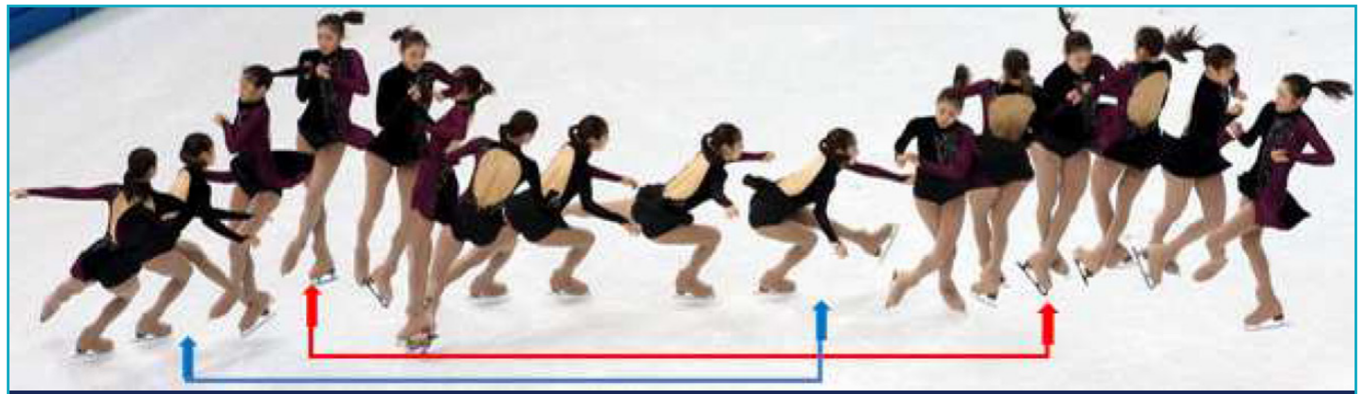

An example of an ice-skating dance sequence tracked in 3D

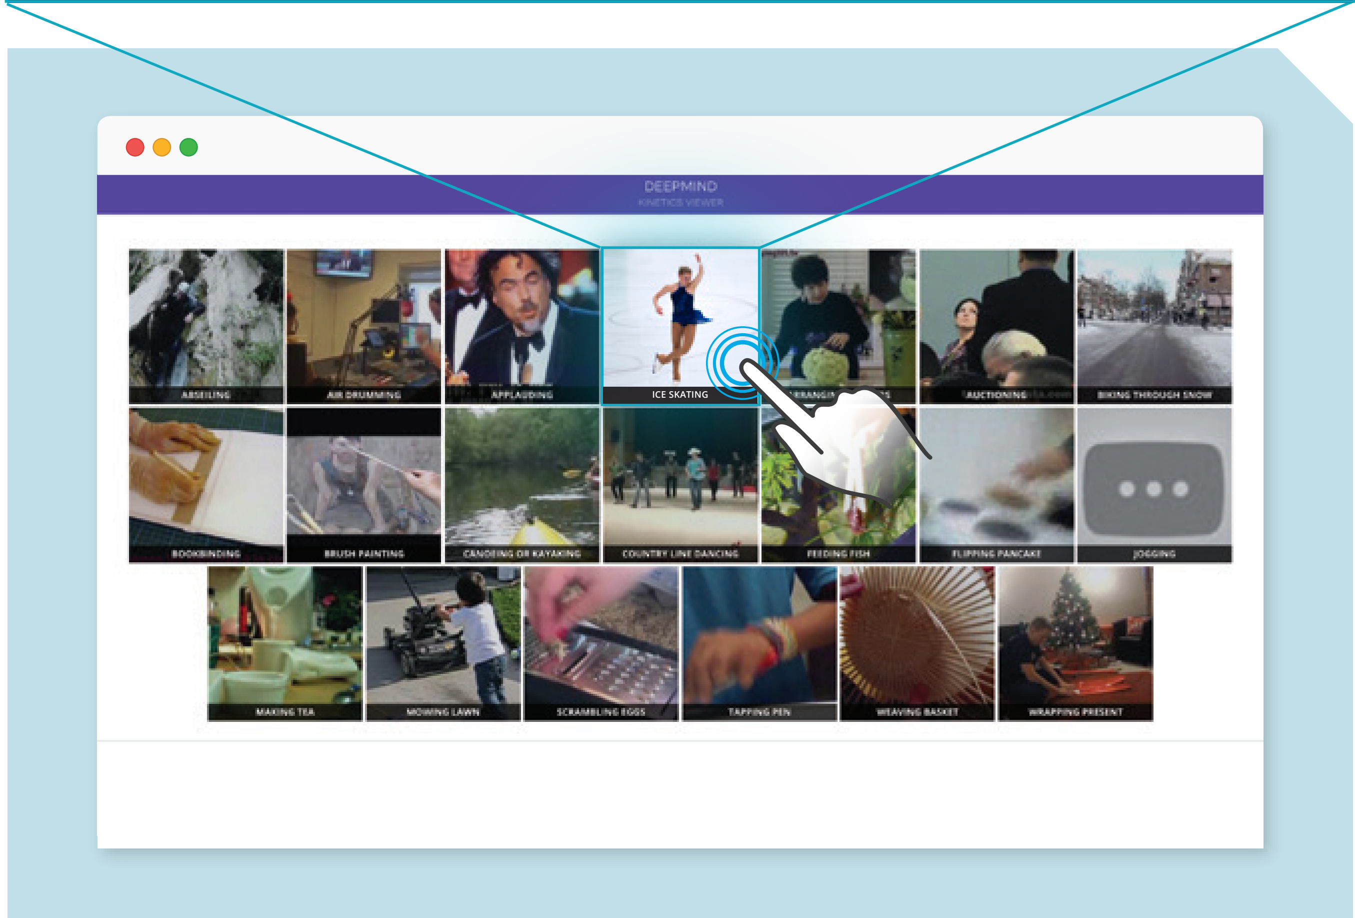

"The models the project team has created allow robots to recognise up to 400 different activities in 10-second video clips; everything from country line dancing to air drumming, applauding, arranging flowers, abseiling, kayaking, jogging, flipping pancakes, scrambling eggs and making tea."

PROJECT

Vision and Language

The ability to process vision takes up more of the human brain than any other function, and language is our primary means of communication. Any robot that is going to communicate flexibly about the world with a human will inevitably need to relate vision and language in much the same way that humans do. It's not that this is the best way to sense, or communicate, only that it's the human way, and communicating with humans is central to what robots need to be able to do.

Project Leader: Anton van den Hengel (CI)

Team Members: Stephen Gould (CI), Chunhua Shen (CI), Anthony Dick (AI), Qi Wu (AI), Hui Li (RF), Violetta Shevchenko (PhD)

Project aim: Develop technologies towards solving a subset of key problems in visual robotics primarily centred on the goal of natural language robot tasking. This extends beyond Visual Question and Answering (VQA) for robots or Dialogue for Tasking, but includes questions of what needs to be learned, stored, and reasoned over for a robot to be able to carry out a general task specified by a human through natural language. Recent progress in VQA has allowed the development of methods which are capable of learning to respond to unforeseen questions about unforeseen images. This is particularly interesting because it requires developing a method which is not designed for a single predefined task, but rather aims to respond in real time to unforeseen input. For the Centre's research this is important towards our objective of creating robots able to see and understand in changing, uncontrolled environments of the real world.

KEY RESULTS:

The project team developed a powerful new method to enable robots to learn to proactively exploit resources available to them. This is a particularly significant development for the application of vision and language technology to robotics because it is impossible to provide a robot with all possible information it will ever need to operate in a general unstructured environment. There will always be unexpected scenarios, objects, and human interactions for which the only option is to actively seek external information that might clarify the situation. The method devised achieves this using the latest meta-learning technology, giving the method the ability to learn how to learn.

The team generated and published one of the first datasets evaluating a robot's capacity to understand a general indoor navigation instruction. This is important because it is an essential capability for robots to move beyond repeating predetermined tasks in a fixed location to carrying out a broad range of tasks in general and unstructured environments.

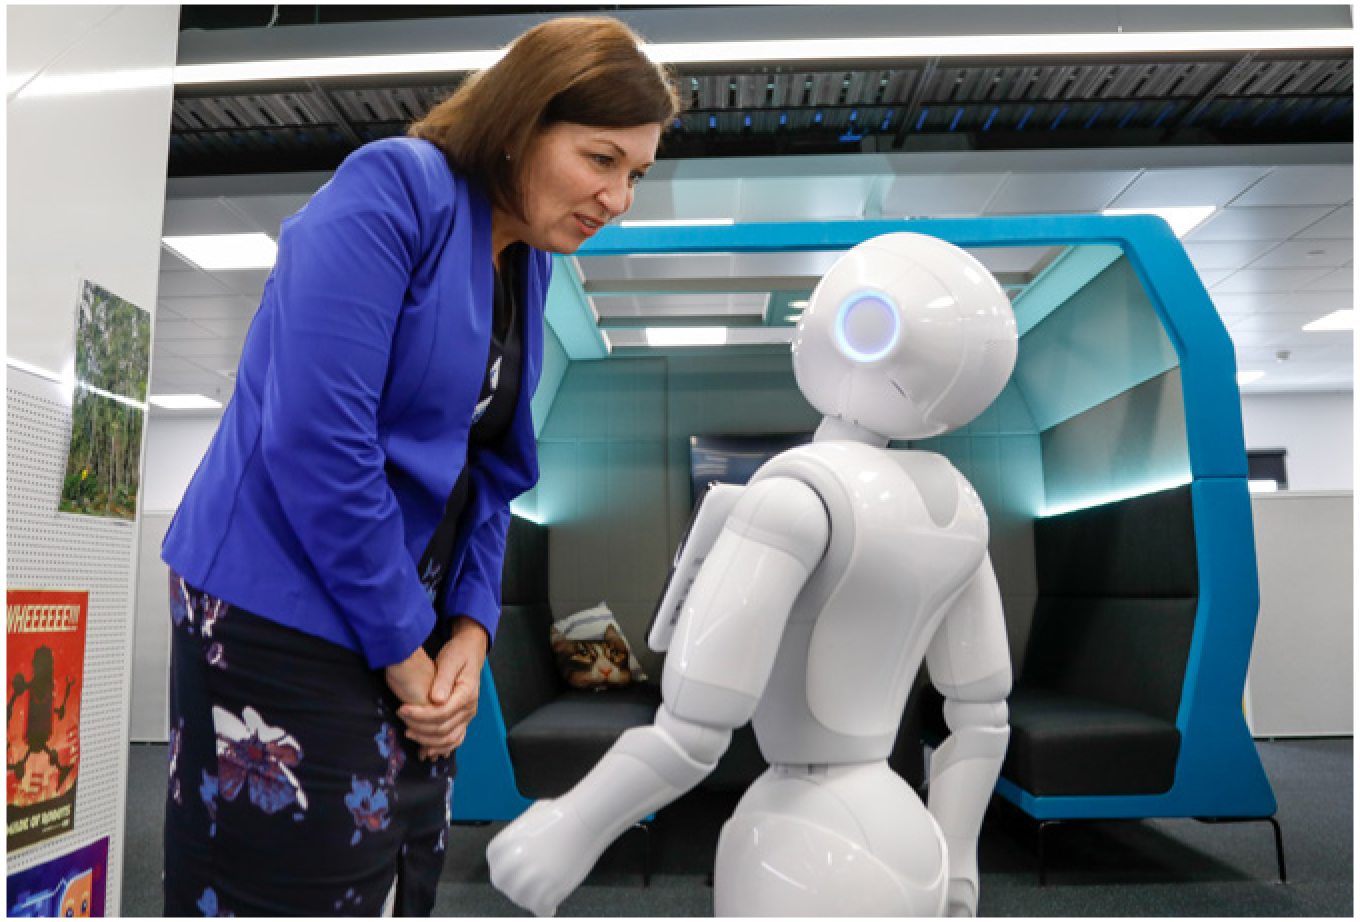

The challenge here is to reflect the generality of indoor environments and the incredible variety of ways that humans provide navigation instructions. Achieving this while enabling metric evaluation of performance requires particular attention. Facebook funding was secured to pay for the data annotation. The fact that the paper describing the dataset has already been cited 50 times indicates its impact, and work continues in this vein with a new object-centred version of the challenge being deployed on Amazon Mechanical Turk.

For full details: *Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments.* Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, Anton van den Hengel, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018.

Placed fourth globally in the Visual Question Answering (VQA) Challenge. Competition was extremely close with tight margins between the first four teams.

Prof van den Hengel presented a series of international talks, including:

- Plenary address at one of the primary international Natural Language Processing (NLP) conferences, ACL 2018. The keynote covered work carried out at the Centre bringing vision and language technologies to robotics, but also identified opportunities for further application of NLP technologies in robotic vision more broadly.
- Plenary address at the Visual Learning and Embodied Agents in Simulation Environments workshop at ECCV 2018 entitled "Problems with open set Vision and Language problems". The talk highlighted the changes required in common algorithm evaluation practice to reflect the fact (long recognised within the robotics literature) that the world is an open-set environment and thus it is impossible to prescribe all of the objects and situations that a robot might find itself in.
- An invited talk at the Australian Driverless Vehicles Summit 2018 on work carried out within the Centre's Vision and Language project and its application to driverless cars.

Acting

While picking up and moving objects is an unconscious activity – something we've learned over time through repetition and routine – for robots, grasping and manipulation is hard. The Centre strives to be the world leader in the research field of visually-guided robotic manipulation. Our goal is for robots to master manipulation in unstructured and changing environments that mirror the messiness and unpredictability of our world.

PROJECT

Manipulation and Vision

Hand-eye coordination in complex visual environments involves developing robust motion control of robotic platforms based on vision data that is capable of dealing with variation and complexities encountered in real-world tasks.

Project Leader: Jürgen 'Juxi' Leitner (RF)

Team Members: Jürgen Leitner (RF), Peter Corke (Centre Director), Robert Mahony (CI), Akansel Cosgun (RF), Douglas Morrison (PhD), Norton Kelly-Boxall (Master of Philosophy Researcher), Zheyu Zhuang (PhD), Robert Lee (PhD), Jordan Erskine (PhD)

Project aim: This project goes beyond engineered visual features and engineered environments to develop demonstrator systems that allow manipulation of real-world objects like capsicums, cups, pens, tools, etc. A key aspect of the project is robustness: to develop systems and architectures that can deal with a wide variety of operating conditions and can adapt easily to new tasks, new objects and new environments. This project aims to create more adaptable, robust solutions by building on previous experiences. These might be experiences previously encountered by the robot itself, for example during exploration or previous task executions, but might also be transferred from other systems or even human demonstrations.

Building on recent advances in deep learning, the project's scientific goals aim to understand the following:

/ How to scope learning algorithms and architectures for hand-eye coordination tasks in complex visual environments.

/ How to interface the resulting robust visual perception system into a control framework and provide analysis of the system performance.

/ How to transfer working architectures from simulation to real-world, and from task-to-task, efficiently and effectively.

KEY RESULTS:

Exhibited Cartman at ICRA 2018.

Successfully demonstrated results of reactive reaching and grasping research at the 2018 International Conference on Intelligent Robots and Systems (IROS). The aim of this research was to achieve robust outcomes with pretty much anything placed in front of a robot to pick up. To achieve this, the system needs to reason about local geometric features to learn the best grasp point to pick up an object. Results were first presented at Robotics Science and Systems in Pittsburgh and also demonstrated on Kinova's new Gen 3 Ultra lightweight robotic arm at IROS 2018 in Madrid.

It is difficult to train a robotic system to perform certain tasks, especially in the real world, where things can change or even break. Simulation provides a wonderfully safe environment to learn behaviours. The stumbling block comes in attempting to transfer these behaviours back into robotic systems as simulators are not perfect. The project team showed that by using a limited number of labelled images, together with cheap-to-collect unlabelled images, a reaching task can be quickly transferred from simulation to real-world use on the robot, Baxter. Success in this area was captured in an IJRR paper by 2018 PhD graduate Fangyi Zhang.

Playing with sand is not just for kids. It can translate into useful scientific results, as demonstrated by Centre Research Fellow Valerio Ortenzi and visiting academic Andrea Cherubini. Their paper, *Towards vision-based manipulation of plastic materials* (published at IROS), reveals that a robot can learn models of deformable materials, such as sand, which are usually very hard to model. Using neural networks, the motions performed were linked with the visual changes observed, allowing the robot to then create sand figures based on a target image provided.

Getting a grip on human-robot cooperation

The Centre collaborated in a ground-breaking study with The BioRobotics Institute of Scuola Superiore Sant'Anna (Pisa, Italy) to reveal guiding principles that regulate choice of grasp type during a human-robot exchange of objects.

During preparation of this annual report, the 2018 study entitled 'On the choice of grasp type and location when handing over an object' was published in Science Robotics (February 2019).

Centre Research Fellow Valerio Ortenzi and The BioRobotics Institute's PhD Researcher Francesca Cini analysed the behaviour of people when they have to grasp an object and when they need to hand it over to a partner.

The researchers looked at the way people picked up the objects – including a pen, a screw driver, a bottle and a toy monkey – passed them to another person and how that person then grasped those objects.

"Perhaps surprisingly, grasping and manipulation are regarded as very intuitive and straightforward actions for us humans," said Dr Ortenzi.

"However, they simply are not. We intended to shed light on the behaviour of humans while interacting in a common manipulation task and handover is a perfect example where little adjustments are performed to best achieve the shared goal to safely pass an object from one person to the other."

The study revealed passers tend to grasp the purposive part of the objects and leave 'handles' unobstructed to the receivers. Intuitively, this choice allows receivers to comfortably perform subsequent tasks with the objects.

The findings will help in the future design of robots that have the task of grasping objects and passing them.

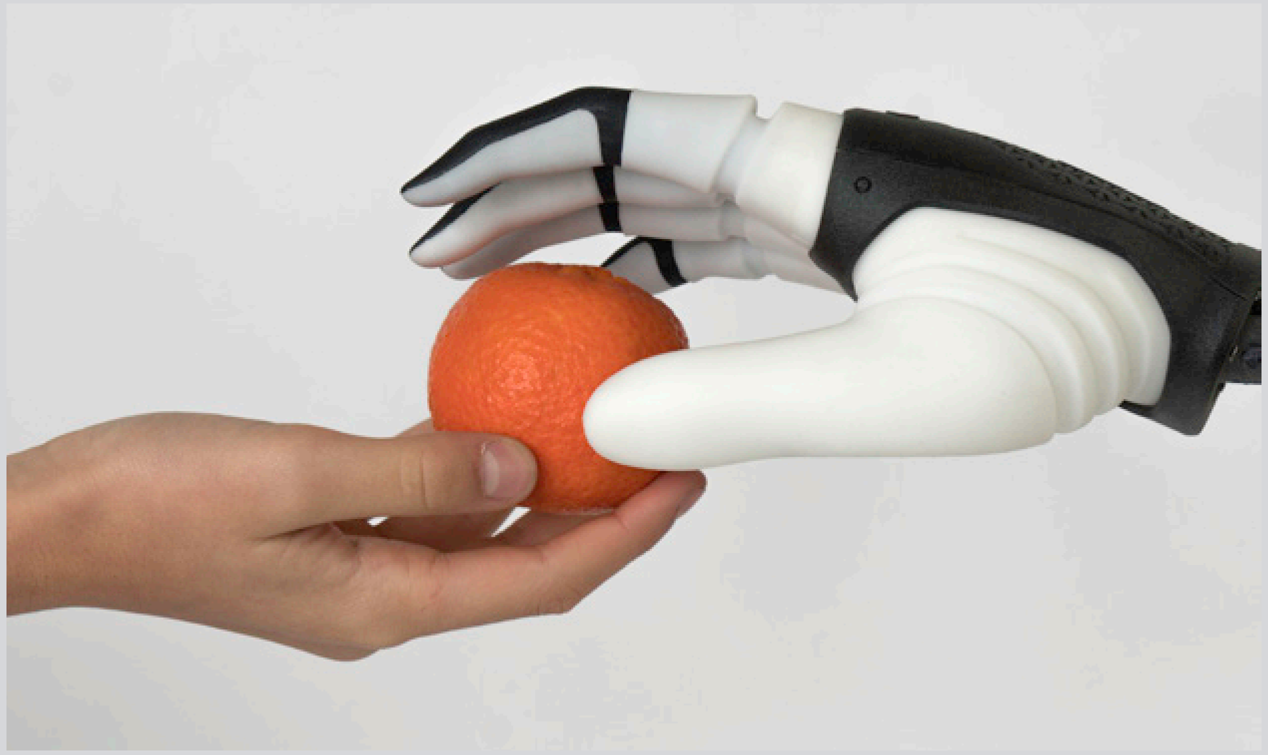

PHOTO by Elastico Disegno www.elasticofarm.com

DEMONSTRATOR PROJECT

Manipulation

Real-world manipulation remains one of the greatest challenges in robotics. In 2018, we introduced a manipulation demonstrator project to better engage with end users and compellingly demonstrate what robotic vision is all about. This demonstrator takes the form of a robotic workstation where we can showcase various robotic manipulation technologies developed within the Centre.

Project Leader: Peter Corke (Centre Director)

Team Members: Jürgen 'Juxi' Leitner (RF), Steve Martin (Research Engineer)

Project aim: The centrepiece of the manipulation demonstrator is a new Panda robotic arm from Franka Emika, a Munich-based start-up. This robot has a number of advanced features, such as force sensing and control, which means it can detect when it has touched the environment and is therefore safe to use in the presence of people. The robot's gripper is fitted with an Intel RealSense RGBD camera to sense the world. It is connected to a dedicated computer system with GPUs to run deep networks and a touch screen.

KEY RESULTS:

The algorithms in the collection so far include:

- deep grasping network of PhD Researcher Doug Morrison;
- visual reaching network of PhD Graduate Fangyi Zhang;
- reinforcement learning to sort objects into piles based on their similarity; and
- a voice-commanded visual pick and place system.

PROJECT

Fast Visual Motion Control

Robots of the future will need to move about in a range of complex, cluttered and unstructured environments full of other autonomous agents, including humans and animals. In order to do this effectively and seamlessly, they must behave much like humans and other animals, moving quickly and competently through cluttered environments. The relationship between motion in the world and changes in images are at the heart of robotic vision, providing rapid and continuous feedback for control.

Project Leader: Robert Mahony (CI)

Team Members: Tom Drummond (CI), Nick Barnes (AI), Jochen Trumpf (AI), Andrew Spek (RF), Thanuja Dharmasiri (RF), Jean-Luc Stevens (PhD), Cedric Scheerlinck (PhD), Timo Stoffregen (PhD), Sean O'Brien (PhD), Pieter van Goor (PhD), Alex Martin (Research Engineer)

Project aim: Develop real-time (augmented) vision processing pipelines that generate dynamic dense structural representations of the environment surrounding a robotic system and couple these representations into robust base-level motion control. This capability is critical for a robot to move dynamically and naturally in a scene independently of the high-level task and action planning routines. The key is for robots to have a real-time dense understanding of the environment in which they move (better than 300Hz refresh rate). The representation doesn't need to be accurate, however, it needs to be reliable. Consider for example, a person crossing a road. It doesn't matter if a robot doesn't know to the millimetre where this person is, when crossing a road. It is, however, a matter of life and death that the robot does know that there is a human present. The project is structured around two themes: "Real-time algorithms for spatio-temporal vision" and "Visual control". The first Centres on vision processing algorithms that run fast enough for highly dynamic mobile vehicles. The second concerns the processing of dense structure representations of the environment into base level vehicle control algorithms.

KEY RESULTS:

Event cameras are a recent imaging technology whereby individual pixels in the camera process incoming light and asynchronously trigger 'events' when the change in measured irradiance exceeds fixed thresholds. They are ideally suited for robotic vision applications since they combine high dynamic range (work outside in sunny conditions and inside in dark conditions), low latency (no time-delay makes integrating vision and control easier), along with low bandwidth (less data clogging up internal networks of the robot). However, the raw data generated by an event camera is incompatible with classical vision sensor data and existing algorithms do not work. That said, a key achievement of the project in 2018 is the development of a new image processing framework for event cameras based on a continuous image state; conceptually the instantaneous value of irradiance illuminating each pixel. The project team demonstrated an asynchronous implementation of a continuous time filter that integrates events to reconstruct an estimate of the image state based on observer theory drawn from the systems and control field. The approach can be extended to compute convolutions of images and allows computation of image gradients, image pyramids, etc.

The approach allows the event camera to be used as a front-end sensor for a wide range of real robotic applications in a natural and straightforward manner.

Drones flying through trees! How can flying Speeder bikes of *Star Wars* fame manoeuvre through forests at break-neck speed without colliding into trees? Only in movies, of course. In the real world, they would have crashed in the first 20 seconds. The Fast Visual Motion Control project is developing the vision algorithms and control technology to achieve similar feats with small scale aerial robots. The key is real-time dense optic flow, coupled with predictive path planning in the flow domain. Optical flow provides key time-to-contact information and dense optic flow, computed at every pixel, provides dense perception of the obstacles that a robot must avoid, even small branches or thin wires. By predicting where the vehicle will fly in the future, the speed at which the vehicle avoids obstacles can be increased, at least until the thrust limitations of the platform lead to the inevitable. In 2018, the project demonstrated a quadrotor vehicle flying through a forested area using dense flow algorithms running on a TEGRA TX2 at 100Hz.

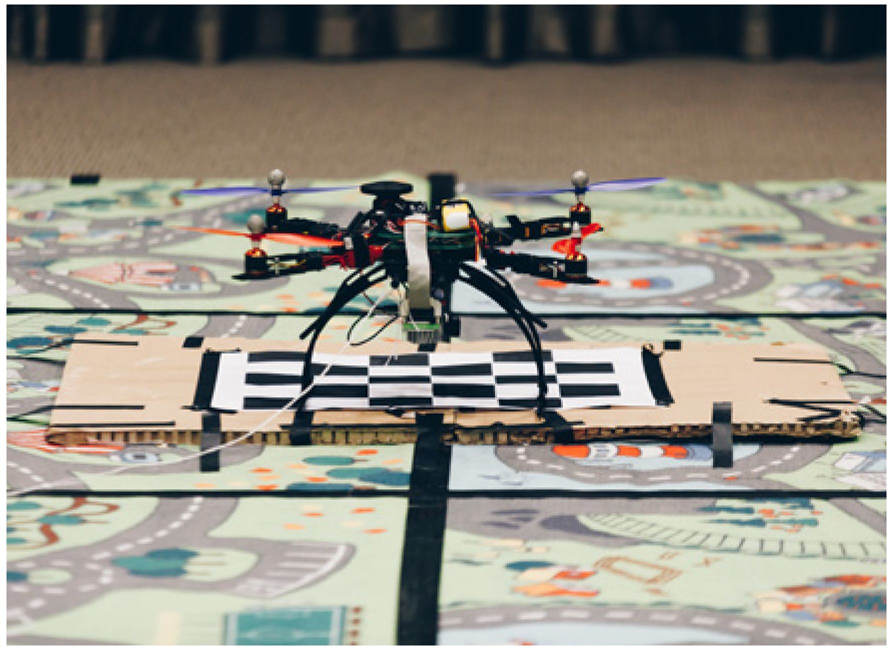

Vision-based Simultaneous Localisation and Mapping (SLAM) has been extensively researched in computer vision and robotics literature. Indeed, so much work has been done on this topic that most experts believe there is little more that can be done. Not so. Advances in this project are developing new algorithms expected to significantly improve state-of-the-art performance for small-scale, real-time implementation of visual SLAM systems. In 2018, a key advance came in the discovery of an abstract symmetry action that models the visual SLAM problem in a natural manner. Thanks to this, the problem of visual SLAM can now be formulated as a dynamic optimisation problem on a set of symmetries rather than in the raw pose-map coordinates that are classically used. Although the formulation is abstract, the results are clear: algorithms developed using this approach have lower computational complexity, use less memory and are more robust for the same performance (accuracy of estimate). Initial results are highly promising and demonstrate comparable performance (accuracy) to state-of-the-art algorithms at a fraction of the compute and memory requirements. The approach has the potential to revolutionise the way low level Visual Inertial Odometry and embedded visual SLAM are performed on real-world, small-scale dynamic robots.

Conventional or monocular cameras produce images by recording the colour of each ray of light that passes through a point or 'pinhole' in space. The resulting image discards distance information and means that a single camera is a poor choice of sensor to enable visual manipulation of nearby objects. Animals and humans get around this by having more than one eye. On a robot, however, every camera takes up space and incurs costs for weight and energy budgets. Single camera solutions are also important for small-scale, embedded robotic applications.

Plenoptic or light field cameras produce images through a single main lens, but record all light rays that pass through 'a volume' of space inside the camera using micro-arrays of lenslets attached to the CCD (charge-coupled device) sensor. It is possible to understand much more about a scene's geometric structure from a light field image than from a conventional image. This enables a robot to act on nearby objects in a scene with confidence about their true location, using only a single camera. However, the accuracy of any 3D reconstruction of a scene depends on the precision of camera calibration. It is more difficult to calibrate a plenoptic camera than a conventional camera. To address this, the Fast Visual Motion Control project proposed a new class of plenoptic features (image structures that represent a specific point in space) along with advanced algorithms that provide state-of-the-art calibration parameters for all plenoptic camera geometries. The algorithms not only work for all camera geometries, they outperform existing algorithms specifically developed for particular cameras. Properly calibrated plenoptic cameras have the potential to provide high quality depth data for near camera robotics applications.

Flashback to Cartman

Remember Cartman? Not US sitcom *South Park*'s most pugnacious character, but a super-savvy robot, built from scratch by a Centre team, acing the 2017 Amazon Robotics Challenge in Japan.

Staking its claim as the only custom-built Cartesian robot at the global challenge – packing a doubly powerful punch as the cheapest – Cartman knocked out 16 major academic and industry research teams from 10 countries to claim the US$80,000 first prize.

The success story features in the Australian Research Council's 2018 snapshot, *Making a difference – Outcomes of ARC supported research*, casting a spotlight on 'remarkable research delivering social, cultural, economic and environmental benefits' to all Australians.

Centre Research Fellow Jürgen 'Juxi' Leitner led the Centre's 22-member team to victory, solving a real-world challenge for global giant Amazon via use of a smart Cartesian robot to pick and stow warehouse items in an unstructured environment.

Taking Cartman's legacy further, Juxi; fellow Amazon Robotics Challenge team member, Master of Philosophy Researcher, Norton Kelly-Boxall; and Research Fellow Nicole Robinson have ventured into the world of start-ups. They established LYRO Robotics (www.lyro.io) as a separate entity to commercialise the next generation of smart robotic systems.

"Robotics is about enhancing human life in some way and, ultimately, a way for communities and individuals to truly connect," Juxi said.

"The development of an innovative and meaningful robotics industry, with robots able to see and understand environments they work in, including in the home, will give people precious time back.

"Like Bill Gates' vision of 'a computer on every desk and in every home', think about the positive impact of smart, affordable robots in every home. A utopian not dystopian future!"

Excited by the prospect of translating research into real-world applications via a start-up – with robotics and automation also positioned to shoulder the three D's in industry, namely 'dull, dangerous and dirty' work – Juxi said the venture would not have been possible without a dynamic culture of innovation fostered at the Centre.

"Incredible passion flows from the Centre's Director Peter Corke, whose philosophy is simple: make something awesome!

"The Centre also has possibly the biggest university robotics Lab in the world, with the largest pool of PhD students working alongside the world's top researchers. So, it's a place where awesome things happen every day."

Learning

Humans learn from experience. The richer our experiences, the more we learn; hence the saying, 'live and learn'. For robots, deep learning – a subset of machine learning loosely modelled on human brain structure – takes the place of learning from experience.

The success of deep learning is attributed to the vast computational resources available and large annotated datasets containing millions of images. In spite of the excitement generated by recent developments in deep learning, there is a lack of understanding about how deep learning works, which invites questions about convergence, stability and robustness of such models.

This program explores the enormous potential that still exists towards solving previously impossible problems in machine perception.

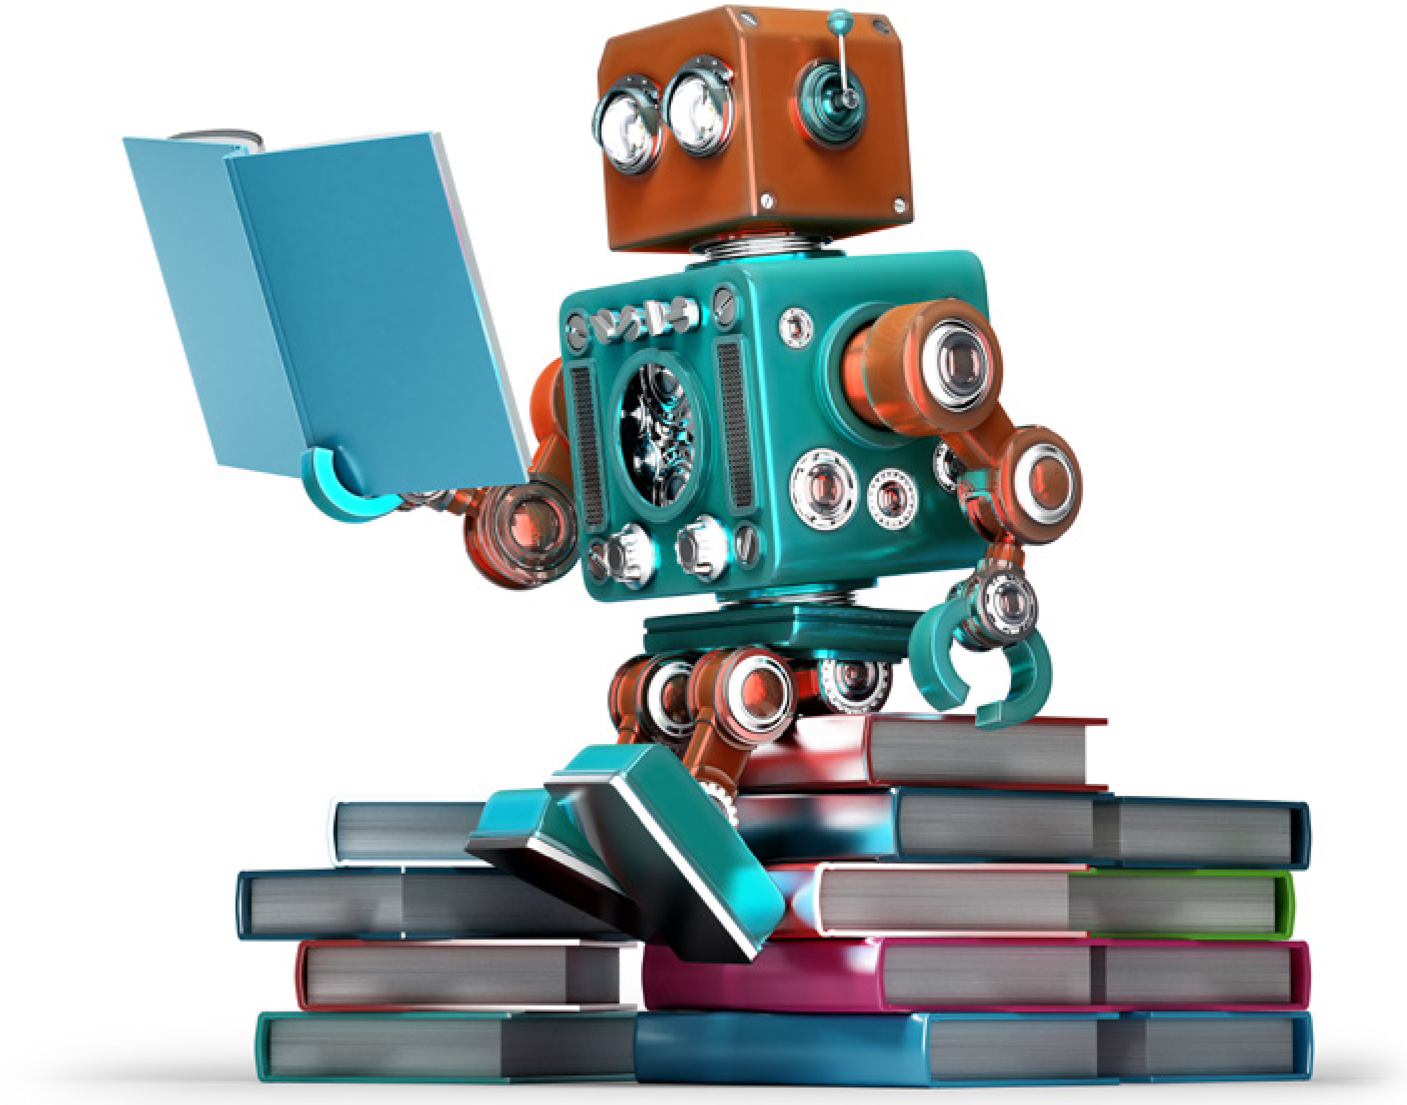

PROJECT

Learning

The Centre is creating new 'game-changing' learning technologies for robotic vision. A key goal is life-long learning so that a robotic system can adapt and continually improve performance over its lifetime.

Project Leader: Gustavo Carneiro (CI)

Team Members: Chunhua Shen (CI), Tom Drummond (CI), Stephen Gould (CI), Niko Sünderhauf (CI), Rafael Felix (PhD), Tong Shen (PhD), Vladimir Nekrasov (PhD), Ming Cai (PhD), Benjamin Meyer (PhD), Ben Harwood (RF), Yan Zuo (PhD), Gil Avraham (PhD), Luis Guerra (PhD), Adrian Johnston (PhD), Michelle Sasdelli (RF), Masoud Faraki (RF), Feras Dayoub (RF), Xian Wang (PhD), Alan Zhu (PhD), Yunyan Xing (PhD), William Hooper (PhD), Dimity Miller (PhD), Serena Mou (PhD), Quazi Rahman (PhD), Jordan Erskine (PhD), Brendan Tidd (PhD), Rodrigo Santa Cruz (PhD), Xin Yu (PhD), Tong Zhang (PhD), Medhani Menikdiwela (PhD), Toan Minh Tran (Affiliated PhD researcher), Jerome Williams (Affiliated PhD researcher), Cuong Cao Nguyen (Affiliated PhD researcher), Amir Rahimi (Associated PhD researcher), Jing Zhang (Associated PhD researcher), Kartik Gupta (Associated PhD researcher), Sadegh Aliakbarian (Associated PhD researcher)

Project aim: This program addresses important challenges in deep learning, such as: effective transfer learning, role of probabilistic graphical models in deep learning, efficient training and inference algorithms. Answering these questions will allow us to design and implement robust visual learning systems that will help our robots fully understand the environment around them.

KEY RESULTS

Six papers were accepted at leading international conferences in the fields of computer vision and machine learning: four at the European Conference on Computer Vision (EECV) and two at the Computer Vision and Pattern Recognition (CVPR) conference. The papers were accepted across several areas, including:

- segmentation in open set world;
- cycle consistent generalised zero-shot learning;
- face super resolution;
- visual question answering;
- webly supervised semantic segmentation; and
- deep learning model compression.

Niko Sünderhauf and Gustavo Carneiro won a Google Faculty Research Award for working on the implementation of a benchmarking dataset and methods to assess open world robotic systems. Relevant to this key result, Niko, Gustavo, Juxi Leitner and Anelia Angelova (Google) are editing an *International Journal of Computer Vision* special issue on deep learning for robotic vision, having received a large number of submissions from world leading researchers in this field.

Two papers were accepted for publication in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, a monthly peer-reviewed scientific journal established by the IEEE Computer Society. They are: "Visual Permutation Learning" by Rodrigo Santa Cruz, Basura Fernando, Anoop Cherian and Stephen Gould; and "Approximate Fisher Information Matrix to Characterise the Training of Deep Neural Networks" by Zhibin Liao, Tom Drummond, Ian Reid and Gustavo Carneiro.

2018 resulted in a number of successful inter-node collaborations. For example, the work between University of Adelaide and Monash on Real-Time Joint Semantic Segmentation and Depth Estimation Using Asymmetric Annotations. Other strong collaborations occurred between University of Adelaide and QUT on generalised zero-shot segmentation and, separately, outlier/anomaly detection.

Some members of the project received an ARC grant of $726,921.00 (funding commences 2019) under its Linkage, Infrastructure, Equipment and Facilities (LIEF) scheme for "A world-class machine learning facility for Australia". Involved in the grant are: Professor Anton van den Hengel; Professor Svetha Venkatesh; Professor Ian Reid; Associate Professor Stephen Gould; Professor Chunhua Shen; Associate Professor Simon Lucey; Associate Professor Hongdong Li; Associate Professor Anthony Dick; Professor Massimo Piccardi; and Professor Gustavo Carneiro.

Technology

We are striving to create advanced algorithms and techniques that will allow robotic systems deployed in large-scale, real-world environments to run computer vision in real time.

It's no secret that competitions or 'challenges' are deeply rooted in robotics research and for good reason, driving significant breakthroughs in technological advancement. Competitions are often thought of as 'zero-sum' games – for one person to win, others must lose. In applied research, this is not the case. Technology is the big winner with positive off-shoots for all. Competitions also inspire collaboration; test and build skills; generate new ideas; foster communication and boost public knowledge and acceptance of robotics. Think RoboCup; Amazon Robotics Challenge (won by a Centre team in 2017); the DARPA Grand Challenge; and Australia's UAV Challenge, co-organised by the Centre. Even the toughest test in AI, the 'Turing Test', devised by Alan Turing in 1950 as a test of a machine's ability to exhibit intelligent behaviour equivalent to (or indistinguishable from) a human, is a form of competition. It was adapted from a Victorian-style competition called the imitation game.

In an exciting new initiative, we are launching our first Robotic Vision Challenge in early 2019. This challenge is destined to become a Centre legacy. Its ultimate objective: to bring global communities together and encourage new thinking on problem-solving.

PROJECT

Robotic Vision Evaluation and Benchmarking

This project will develop new standardised benchmark tasks, evaluation metrics, and a new robotic vision challenge competition. It is being run in two parts:

ROBOTIC VISION CHALLENGE

Big benchmark challenges and competitions like Pascal VOC, ILSVRC ('ImageNet') or COCO supported much of the remarkable progress in computer vision and deep learning over the past few years. The Centre aims to recreate this success for robotic vision. The competition is planned to run annually in conjunction with a major computer vision and robotics conference.

BENCHBOT

This innovative online portal will allow researchers around the world to remotely test their machine learning systems on-board various types of real robots in real environments.

Project Leader: Niko Sünderhauf (CI) / Feras Dayoub (RF)

Team Members: David Hall (RF), John Skinner (Research Associate), Rohan Smith (Research Associate), Ben Talbot (Research Associate), Haoyang Zhang (RF)

Project aim: To overcome the current lack of meaningful standardised evaluation protocols and benchmarks which are a significant roadblock for the evolution of robotic vision and impede reproducible and comparable research in robotic vision globally.

The Centre's first **Robotic Vision Challenge** requires participants to detect objects in video data produced from high-fidelity simulations. The novelty of this challenge is that participants are rewarded for providing accurate estimates of both spatial and semantic uncertainty for every detection using probabilistic bounding boxes. Robots will greatly benefit from a sense of spatial and semantic uncertainty to 'switch on' robotic vision like any other sensor and avoid becoming over-confident about *where* objects are and *what* they are, potentially resulting in catastrophic failure (like crashing into things).

The competition aims to motivate researchers to develop object detection methods that work robustly and reliably on robots.

Robots will greatly benefit from a sense of spatial and semantic uncertainty to 'switch on' robotic vision like any other sensor and avoid becoming over-confident about where objects are and what they are, potentially resulting in catastrophic failure.

While much progress has been made in improving the performance of computer vision within the robotics domain, it is still very difficult for researchers to transition their solutions from desktop computers to robots. Reasons for this include a steep learning curve and variable hardware and software configurations which are typically individual to each robotic platform. **BenchBot** aims to alleviate this by providing a standard benchmarking test relying on the same hardware and realistic conditions to evaluate different machine learning systems and compare their performance in a fair manner. BenchBot also aims to provide access to robots to researchers in institutions that are not able to maintain their own inventory of robots.

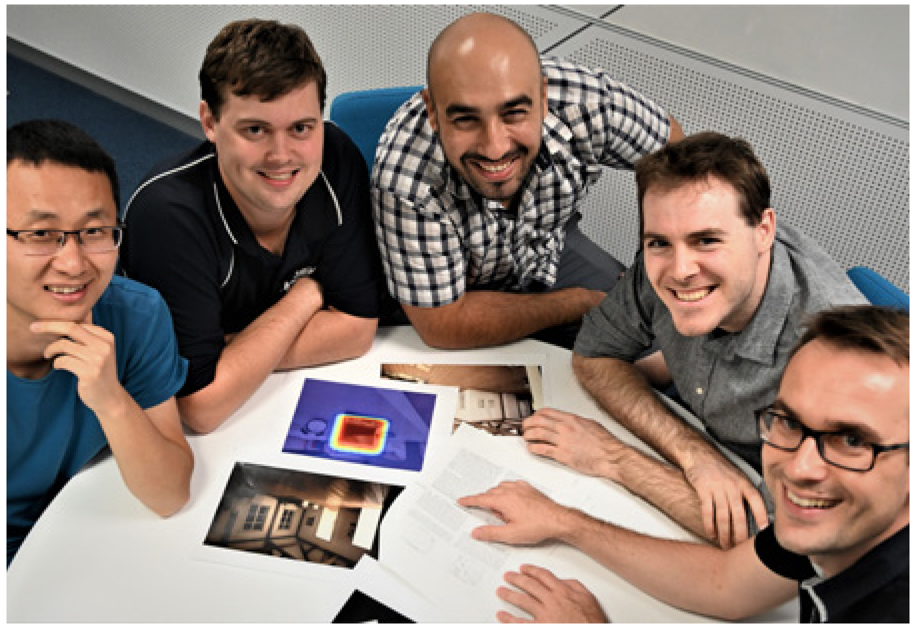

Simulated video sequences

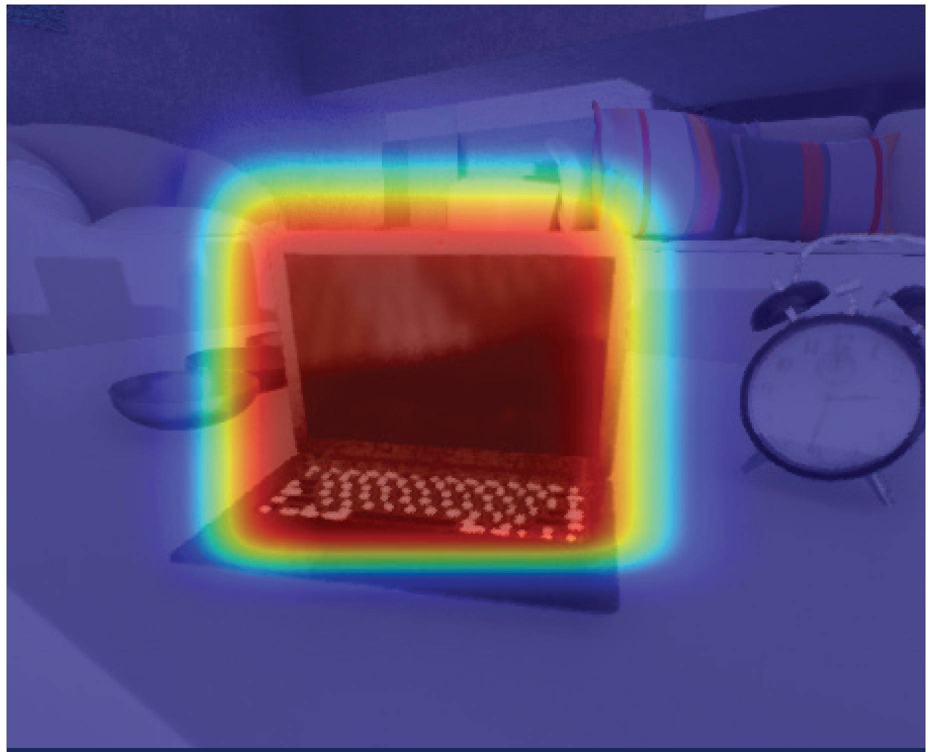

Probabilistic bounding box detecting a laptop

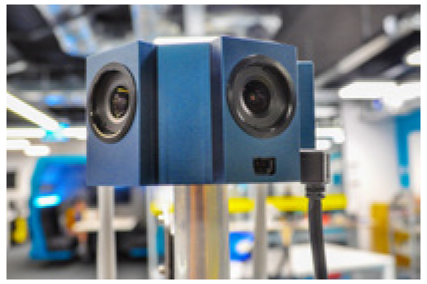

KEY RESULTS

ROBOTIC VISION CHALLENGE

- The Centre's first Robotic Vision Challenge was released online in December 2018 (www.roboticvisionchallenge.org). Results and competition winners will be announced during a workshop at the world's largest computer vision conference,the Computer Vision and Pattern Recognition (CVPR) conference in 2019.
- The project team's work to create the competition included development of a new evaluation measure, evaluation protocol and dataset to assess how well object detection methods can estimate spatial and label uncertainty. The new probabilistic detection quality (PDQ) measure has been described in a paper and submitted to CVPR. The dataset consists of over 56,000 images from 18 simulated indoor video sequences, featuring day and night scenes and simulating three different types of domestic service robots.
- Niko Sünderhauf received a $72,000 Google Faculty Research Award to support the project.
- Workshops were run at CVPR and the prestigious Robotics: Science and Systems (RSS) conference. Importantly, ideas and insights about early concepts of the challenge from robotics and computer vision research communities at these forums helped shape the design and scope of the Robotic Vision Challenge.
- The project team started a collaboration with colleagues from Imperial College in London to extend the challenge dataset.

BENCHBOT

- Development and implementation of BenchBot, running on a dedicated server at the Centre's QUT node.
- Successfully demonstrated BenchBot using two different types of robot: Pepper and Guiabot.

DEMONSTRATOR PROJECT

Self-driving Cars

Project Leader: Michael Milford (CI)

Team Members: Niko Sünderhauf (CI), Chunhua Shen (CI), Ian Reid (CI), Hongdong Li (CI), James Mount (PhD), Sourav Garg (PhD), Steven Martin (Research Engineer)

Project aim: Develop a suite of mini autonomous vehicles that can be replicated across the Centre and used at each node (QUT, The University of Adelaide, ANU and Monash University). The primary objective is to enable Centre researchers, no matter where they are located, to better demonstrate and develop their research.

Ultimately, these platforms will be used to engage with members of the public, industry and government to demonstrate self-driving car technology and associated ethical and technological problems.

Driving this development has been the dedicated work of research engineer Steven Martin and PhD student James Mount, assisted by a team of research assistants and students including Jason Queen, Christopher Lenton, Lucas Browne, Quinlan Barthelme, Alexander Erickson, Hasini Charuka and Tim Quelch.

KEY RESULTS

The project team progressed a miniature self-driving car from a first generation prototype to a polished third generation vehicle. The team is close to finalising a replicable, mature platform for research demonstration and fundamental research.

In 2018, the car traveled from its QUT base to other Centre nodes in Melbourne (Monash University) and Canberra, where it was demonstrated at RoboVis and featured in discussions around usage and collaboration opportunities.

Project members published autonomous car-related research at leading conferences (International Conference for Robotics and Automation; International Conference on Intelligent Robots and Systems; and Robotics: Science and Systems) and a number of other top-tier publication outlets. Topics included work on government and industry autonomous vehicle projects led by Centre students including Sourav Garg and research engineer Adam Jacobson. Autonomous vehicle-related publications were also accepted for ICRA 2019.

Another highlight of 2018 was a visit to the Centre by Lisa Teunissen, a Masters student in Automotive Technology at Eindhoven University of Technology in the Netherlands. During a four-month internship, Lisa worked on a range of robotic vision projects, developing lane detection and pedestrian crossing detection capabilities for the miniature autonomous car. She demonstrated its performance and delivered a final report. Lisa's visit led to an additional internship ongoing in 2019 with Attila Lengyel, an Electrical Engineering student from Delft University of Technology, also in the Netherlands.

Project members showcased autonomous car-related research and provided expert commentary at a number of organisations, with invitations to talk at Uber, Argo.AI, Nuro.AI, Nvidia, Bosch Autonomous Driving, AutoX and local organisations including Queensland Department of Transport and Main Roads. Presentations were also given at the Autonomous Vehicles Summit, Multiple Future Transport Panels with KWM, Myriad, CEBIT, and the Australian Electric Vehicle Expo. Of 66 presentations in 2018 across 12 cities and five countries, a significant number touched on autonomous vehicles.

The self-driving cars demonstrator project featured on national children's TV program Brain Buzz. Centre PhD student James Mount appeared in studio, demonstrating to children how autonomous cars can work and fail.

Project members also provided extensive media commentary around autonomous vehicles, especially in the wake of the Uber fatality in March 2018, with six media interviews in one day, including a feature piece on Stan Grant's Matter of Fact.

Lane marker detection and pedestrian detection using a multi-stage visual processing system

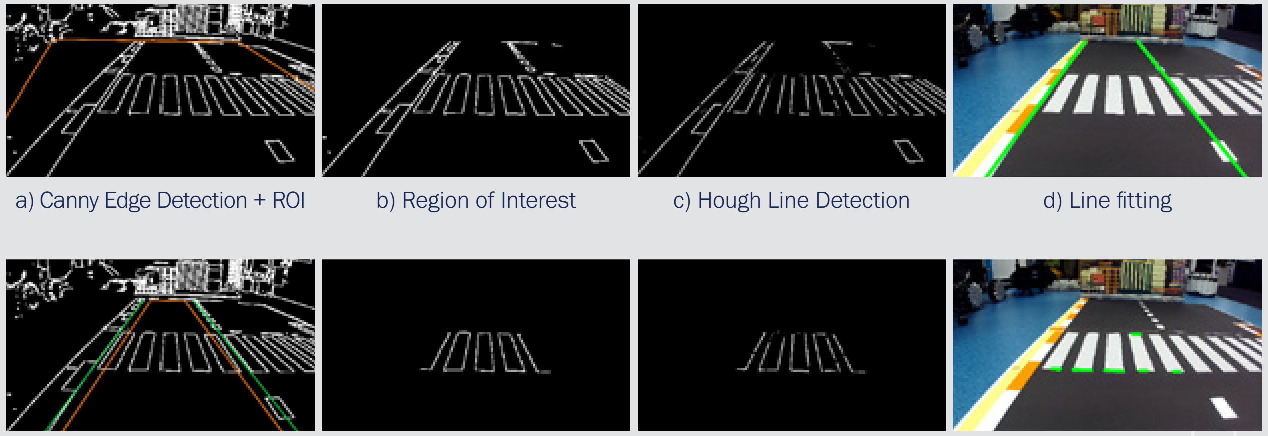

a) Canny Edge Detection + ROI | b) Region of Interest | c) Hough Line Detection | d) Line fitting

a) Canny Edge Detection + ROI | b) Region of Interest | c) Hough Line Detection | d) Crosswalk detection

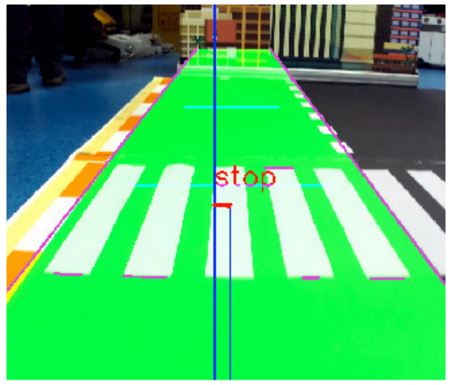

Activity Plan for 2019

CULTURE

We are creating a vibrant, high energy, future-focused, collaborative robotic vision community, developing knowledge leaders for both industry and academia

STRATEGIC OBJECTIVES

- Develop the next generation of robotic vision experts through effective recruitment, retention and training
- Ensure the Centre functions as a cohesive organisation of interactive, collaborative and highly effective research teams

KEY TASKS

- Maintain full complement of Research Fellows.
- Aim for 90 per cent retention of PhD enrolments.
- Provide knowledge leadership training to all early career researchers to help in their career development.
- One annual symposium each year to build Centre culture.
- Regular (monthly) meetings of the Research Committee and project teams with quarterly face-to-face meetings.
- Centre Executive meetings at least monthly with quarterly face-to-face meetings held in conjunction with Research Committee meetings.
- Implement initiatives and monitor progress towards meeting our Centre's Gender KPI targets.

INTEGRATE

We are bringing the disciplines of robotics and computer vision together to create new robotic vision technologies

STRATEGIC OBJECTIVES

- Connect research organisations, governments, industry and the private sector to build critical mass in robotic vision
- Lead robotic vision in Australia and overseas

KEY TASKS

- Promote our Robotic Vision Resources Hub.
- Visit international partners and host visits.
- Organise robotic vision workshops at key international conferences.
- Hold an annual Robotic Vision Summer School targeted at international attendees.

SCIENCE

We are leading the world in transformational research in the new field of robotic vision

STRATEGIC OBJECTIVES

/ Deliver internationally recognised research in robotic vision

/ Create and implement projects based on collaboration and innovation that enhance research outcomes

KEY TASKS

SENSING

PhD Research Project (Event Cameras)

/ Develop novel asynchronous neural network algorithms for event cameras that can perform tasks such as classification at low-latency and low-power.

/ Smart fusion of event cameras with multiple sensing modalities for image reconstruction, segmentation and optical flow.

PhD Research Project (Light Field Cameras)

/ Develop a visual servoing algorithm using a light field camera that will allow a robot to control its pose relative to the refractive object.

UNDERSTANDING

Research Project: Scene Understanding

/ Incorporate notion of Scene Graph (learned from data) into object-based SLAM. This will enable maps to be interpreted automatically at a higher level of abstraction; lead to improved data association; and a more principled way of proposing relationships between objects.

/ Further exploration of the boundaries between geometry and deep networks. In the first instance we will explore more comprehensive integration of uncertainty from deep networks as a means to move from deep Visual Odometry to deep SLAM.

/ Build geometric and semantic models and integrate these with social robotics and self-driving demonstrations.

Research Project: Robots, Humans and Action

/ Demonstrate 3D object and human pose tracking in videos of a person performing a simple task, such as assembling a piece of Ikea furniture.

/ Develop state-of-the-art models for inferring specific human-object interactions and tracking of these interactions over time.

/ Forecast the movement of a person to infer future object interactions and anticipate intent.

Research Project: Vision and Language

/ Extend the Room2Room dataset to evaluate a robot's ability to identify a specific object in another room. This is a significant step towards the 'Bring me a spoon' challenge, and an important extension of the existing dataset because it develops the challenge from merely navigating to the right location into that of identifying a specific object in a particular location.

/ Develop a demonstrator on the robotic arm located at the Centre's University of Adelaide node. We have demonstrated V2L technology previously on the Pepper robot. This demonstrator will extend this work, with the aim of enabling the robotic arm to follow novel natural language instructions.

/ Develop technology to enable a robot to identify information that it needs to specify and then complete its task. Moving from VQA into Visual Dialogue will provide the capability to ask questions that seek information necessary to complete a task, and to identify when enough information has been gathered and an action should be taken.

SCIENCE CONT.

ACTING

Research Project: Manipulation and Vision

/ Grasping with Intent: combine task-level description and constraints with the reactive grasping approaches developed in 2018.

/ Simulation to real world: how to efficiently learn (including RL policies) more complex tasks, such as sorting or manipulation, including quantifying the gap between what we can simulate and how the real world behaves.

/ Investigate how mobile bases can be helpful to perform higher level manipulation tasks like cleaning up a room.

Demonstrator Project: Manipulation

/ Table top manipulation demo with touch screen driven algorithm selection.

/ Mobile manipulation demo for indoor tasks.

/ Integration of Interaction (voice and gesture) into manipulation.

/ Sim2Real Transfer.

Research Project: Fast Visual Motion Control

/ Development of an image processing pipeline for event camera data for robotic applications. The processing pipeline will generate dense optical flow, find and track image features, and compute convolutional networks.

/ Development of a low-level Visual Inertial Odometry system that runs in real-time on embedded hardware for monocular cameras.

/ Algorithms for plenoptic depth in real time. The proposed approach will reduce computational cost of depth estimates by an order of magnitude and be ideal for embedded robotic applications.

LEARNING

Research Project: Learning

/ At least two inter-node publications involving researchers in the learning program and researchers in other programs and demonstrators.

/ A record number of Computer Vision and Pattern Recognition (CVPR) conference and International Conference on Computer Vision (ICCV) papers in 2019 is expected to position the Centre's learning program as one of the most prolific in the world. State-of-the-art results are expected in generalised zero-shot classification and detection; meta-learning; distance metric learning; weakly supervised segmentation; visual question answering; deep learning model compression; anomaly detection; reinforcement learning; and GANs.

/ Establish a dataset and benchmark to assess open world robotic systems. The aim is to publish work that will be used by the community in the assessment of robots that operate in real-world environments.

SCIENCE CONT.

TECHNOLOGY

Research Project: Robotic Vision Evaluation and Benchmarking

Robotic Vision Challenge

- / Develop a second challenge that supports active vision so that participants control the movements of the camera inside the simulation as if it was on a robot. This will require close collaboration with partners from NVIDIA.
- / Attract competition participants from around the world, including via publicity, outreach and workshops at computer vision and robotics conferences such as the Computer Vision and Pattern Recognition (CVPR) conference; International Conference for Robotics and Automation (ICRA); and Robotics: Science and Systems (RSS).

BenchBot

- / Tour all Centre nodes to showcase the platform and invite participants.
- / Make BenchBot publicly available online.
- / Publish a research paper about the project.
- / Design new tasks for robotic arms or mobile manipulators.

Demonstrator Project: Self-driving Cars

- / Replicate the initial 3rd generation car to form a fleet across the Centre's four nodes, providing each node with a high quality demonstration platform for research as well as a tool for further research and development.
- / Work across nodes to continue initial discussions around research collaborations using both the miniature and full-size autonomous research platforms available to the Centre.
- / Utilise the miniature autonomous cars to provide technology demonstrations to regular tours of the Centre headquarters at QUT and for select public outreach and engagement events with various stakeholders.

ENGAGE

We engage with people about the potential of robotic vision technologies by developing accessible robotic vision resources

STRATEGIC OBJECTIVES

- / Identify and engage with key stakeholders on the potential applications of robotic vision
- / Establish vibrant national and international robotic vision communities
- / Increase robotic vision educational opportunities

KEY TASKS

- / Host visits and tours of robotic vision facilities by government, industry and the community (including school groups).
- / Create media releases by partners related to robotic vision.
- / Create public lectures on robotic vision.
- / Promote the Robot Academy.
- / Continue to build network of industry contacts and maintain a customer relationship management (CRM) system.
- / Advocate recommendations contained in the Robotics Roadmap and continue its legacy via a website platform.

TRANSFORM

We solve innovation challenges by applying robotic vision technologies to transform the world

STRATEGIC OBJECTIVES

- / Demonstrate how research can advance products and services
- / Generate downstream investment to take robotic vision technology into industry
- / Foster innovation, entrepreneurship and new enterprises to advance robotic vision

KEY TASKS

- / Provide knowledge leadership training to all early career researchers to help in their career development.
- / Engage Centre's Advisory Board at least twice yearly.
- / Present on robotic vision at key industry events.
- / Popular science articles on Centre work in targeted publications.
- / Support spin-outs from the Centre by providing flexible work arrangements for researchers keen to explore entrepreneurship.

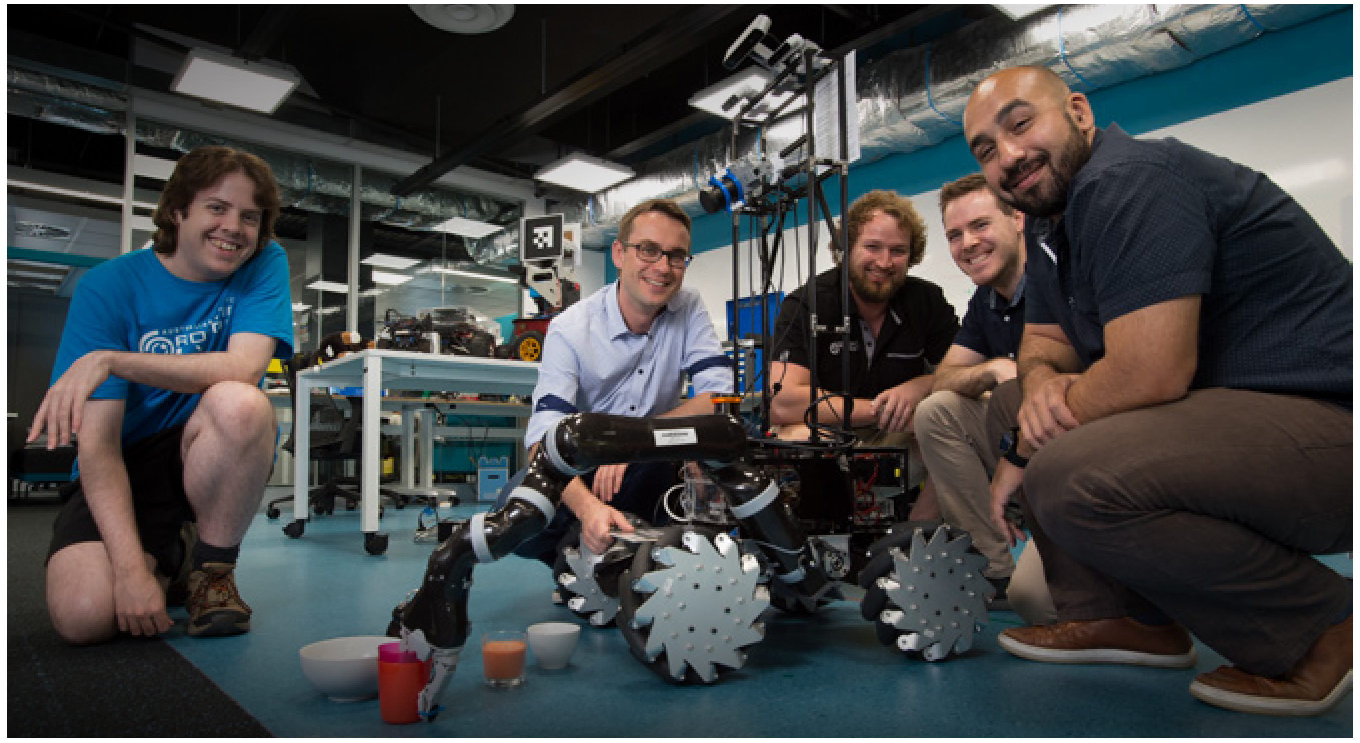

PHOTO by Anthony Weate

Section 3 **Research Impact**

Our transformational research outcomes will be applied to three critical challenges facing the world:

ENVIRONMENT
Robots to manage, protect and repair our natural and built environments

HEALTHCARE
Robots for improved and affordable healthcare in the home and hospital

RESOURCES
Robots to safely and effectively harness our natural resources, in particular, sustainable food production

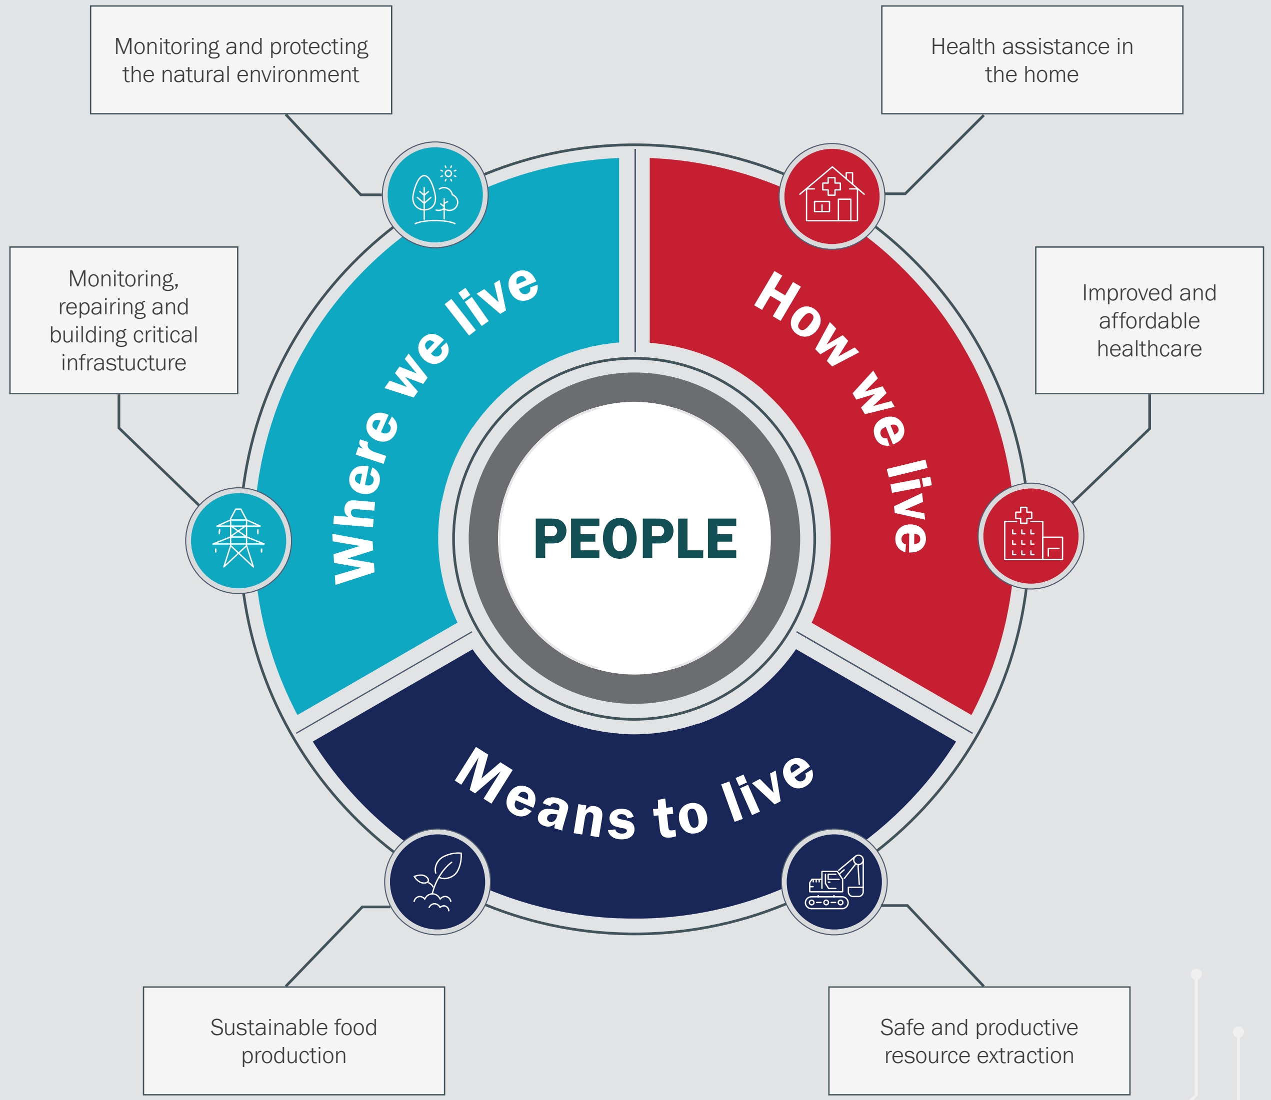

How we live

A key concern for any society is the health of its people. Good health, however, stretches beyond the presence or absence of disease to our physical, mental and social wellbeing.

Robotics and automation can play a key role across diverse areas of health and welfare, from preventative care to assisted living; primary healthcare (GPs); hospital-based care; and emergency assistance.

In Australia, health dominates public expenditure and employment. This makes it a prime candidate for innovation to both reduce costs and improve outcomes for patients and healthcare workers.

According to *Australia's Health 2018* (a biennial report by the Australian Institute of Health and Welfare), on an average day $467 million is spent on health. This adds up to $170.4 billion annually, or more than 10 per cent of our national GDP.

One in two Australians is estimated to have at least one of eight common chronic conditions: cancer; cardiovascular disease; mental health conditions; arthritis, back pain; chronic obstructive pulmonary disease; asthma; and diabetes.

As a result, on an average day, 406,000 visits are made to a GP; 21,400 presentations are made to public hospital emergency departments; 6,000 elective surgeries are performed; and 26,000 specialised community mental healthcare services are provided.

Not surprisingly, this places pressure on healthcare workers and resources. As detailed in Australia's first Robotics Roadmap, released by the Centre in June 2018, the healthcare (and social assistance) industry accounted for 16 per cent of serious workers' compensation claims in 2014-15, totalling 17,566 claims. Healthcare and social workers are a key group for injury due to the nature of their daily work. Healthcare workers may be exposed to a range of hazards including highly toxic drug and chemical agents, workplace stress and violence. They also perform physically demanding repetitive tasks such as lifting patients.

And here, truly useful robots step in. Not purely for their ability to take on the 'Three D's' – namely dull, dangerous and dirty tasks – but to enhance the healthcare system, covering everything from surgical assistance to internal logistics, patient handling, rehabilitation and care.

"Australians are living longer than ever before, but half of us are living with at least one chronic condition, which can affect the quality of our lives, as well as those of our families and carers"

–*Australia's Health 2018*; 16th biennial report by the Australian Institute of Health and Welfare

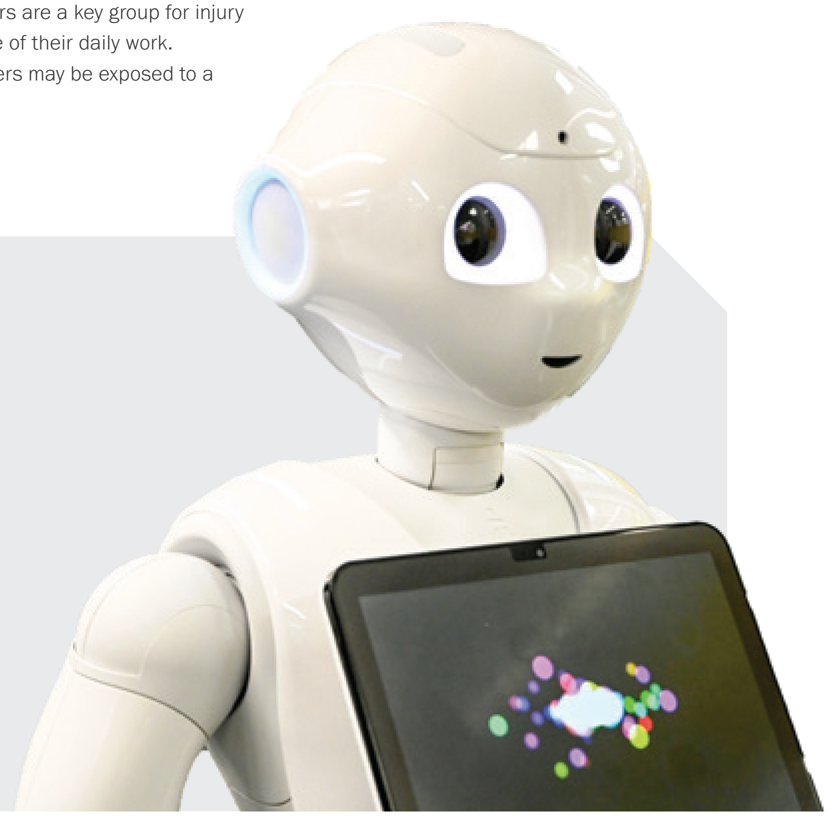

Maximising our research impact

To fulfil our leadership role in the new field of robotic vision while maximising our research impact, we focus on building sustainable partnerships across the research sector, together with public and private enterprises.

An important part of this work involves generating additional income and sources of funding. By increasing our network of end users, we are given the opportunity to conduct research that attracts investment outside our ARC funding.

Since its inception in 2014, the Australian Centre for Robotic Vision has attracted over $32 million in additional external funding including from industry, government, other ARC grant funding and university contributions.

PROJECT
Social Robots

Project Leader: Belinda Ward

Team Members: Gavin Suddrey (Software Engineer), Suman Bista (RF), Nicole Robinson (RF), Sue Keay (Chief Operating Officer).

In recent years, humanoid robots have stepped (or rolled) out of the laboratory to become commercially available products. These robots are designed to be engaging and to interact with people via natural language and visual interfaces. They are not designed to have complete physical capabilities such as lifting, grasping and manipulation like the universally-loved humanoid robots of science fiction. To explore how social robots can enhance our lives, the Centre secured a grant of $1.5 million from the Queensland Government to enhance the capabilities of SoftBank's Pepper robot using robotic vision. The research identifies key applications where social robots can deliver value beyond their novelty appeal. The Centre's social robotics team has focused on use cases for social robots in health, mental health and aged care.

PHOTO by Anthony Weate

Socially Assistive Robots in Health and Aged Care

In the health context, digital communication methods have increased due to the prevalence of technology, user familiarity and convenience of being able to engage in a conversation outside traditional face-to-face methods.

Communication with an embodied agent, such as a social robot, has potential to augment the cost-effectiveness of face-to-face assessment and treatment, and even become the primary delivery mode for some interventions. However, the inclusion of robots in healthcare requires their acceptance by potential users, and by healthcare professionals.

Meanwhile, our aging population is growing, with 3.7 million Australians now aged over 65. Social isolation is a growing problem for people living in their own homes. Many residents in aged care facilities receive less than an hour per day of human contact with staff and may have no visitors at all. Social robots offer a possible source of interaction or may be used to check on human wellbeing, encourage positive health outcomes and ensure that human-to-human interactions are well informed. Their role is not to replace but assist and enhance the work of human staff.

PEPPER, UP CLOSE AND PERSONAL:

In 2018, the social robotics team initiated a set of trials measuring the efficacy of social robots in one-on-one interactions in healthcare. The aim is to understand the factors that determine whether a person will engage with the robot beyond satisfying initial curiosity. A future step would be to determine whether meaningful interaction continues after the initial novelty wears off. Importantly, the trials build on the results of Centre Research Fellow Nicole Robinson's PhD thesis, which produced a scale predicting a person's ongoing intention to interact with a social robot.

The research trials will identify key areas for use of the Pepper robot (and future generations of social robots) to support Queensland healthcare processes as a new digital healthcare tool.

Trial 1: Assessing attitudes to robots and willingness of participants to engage with and discuss sensitive topics with a robot versus a human health practitioner

This 500-participant online study will identify specific areas in which the use of a social robot may have the greatest impact based on individual preference and perceived utility of the robot to assist people. It will also provide better insight into areas where people strongly prefer human intervention and guide future efforts to avoid pressing social robots into contexts that are not yet acceptable and where people have a low willingness to discuss issues with a robot.

Trial 2: Assessing the impact of a social robot delivering advice on wellbeing in a one-on-one session

This trial tests whether social robots can be a viable option for helping people with wellbeing exercises. The social robotics team recruited more than 200 participants with each session taking approximately 10 minutes. Sessions were conducted at QUT (Gardens Point campus), either in a semi-private room in The Cube study space or in Hi-Q in the Library.

Participants were a mixture of University students, staff and visitors. They underwent a live face-to-face session with Pepper to investigate the use of a robot to deliver a wellbeing technique and measure the individual's rating of the robot providing the instruction.

This was a two-condition study. One involved the brief wellbeing exercise, and the other involved a casual conversation, with participants asked general questions about their interests. The results will also help to inform the important elements in a social robot interaction when designing longer session interactions or health interventions, such as mental health interventions or wellbeing assessments.

The use of a control condition (where the participant discusses their interests) will help to identify whether the perceived incentives and utility of a social robot is based on novelty alone, rather than the context or session interaction.

Trial 3: This study, starting early 2019, will take Pepper into an operating health clinic at QUT's Kelvin Grove Campus.

It will quantitatively assess the performance and acceptability of using a social robot in a health clinic to deliver an automated health assessment and then provide customised feedback to take to a professional healthcare appointment.

We prefer our robots more 'human' thanks

Centre Research Fellow, Nicole Robinson studied people's responses to social robots during her PhD. Her findings inform our research in this space. Dr Robinson's topic was motivated by her background in psychology and health, as well as her interest in robotics.

Mental illness rates and non-communicable diseases are on the rise. There is an urgent need to find new ways to tackle the problem to help improve health and wellbeing in people.

Digital healthcare could include social robots to be part of an overall approach to reducing the treatment gap and helping more people with their health and wellbeing. This type of technology could support people in areas where it is difficult to access health services or people feel too stigmatised to reach out for help.

Published in *Science Robotics* (September 2018), Dr Robinson's results suggest people are more willing to respond to a social robot with 'humanoid' qualities rather than obviously mechanical ones.

Additionally, their intention to use a social robot can be predicted by their emotional response, perceived utility and social/relational connection to the robot, which can be measured via a simple questionnaire. This can help identify key areas where robots would be well-liked and acceptable to use, or areas where people are not ready to interact with a social robot.

To encourage interaction with a social robot, researchers may need to find new ways for people to have a positive emotional response to them. The challenge is to make them better companions and increase the perceived utility of a social robot, such as being able to help a person solve a problem.

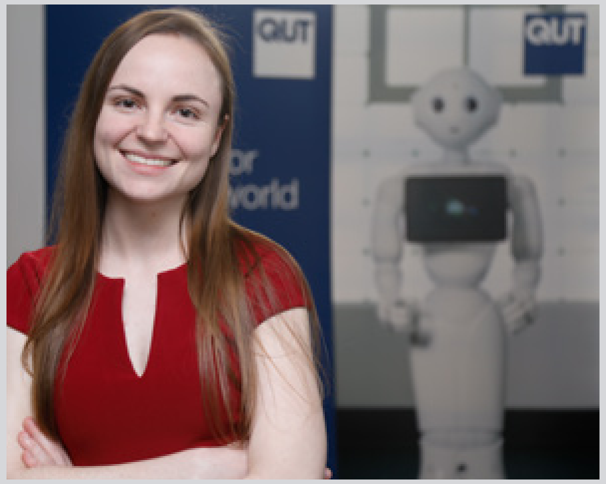

Pepper at Townsville Hospital

To be effective, applications should be developed in conjunction with healthcare professionals, and with an understanding of the factors that determine the acceptability and efficacy of a social robot in the relevant context. It's also important to manage expectations so that spending decisions are well informed. Providing a valuable service involves more than simply purchasing or leasing a robot.

The Centre's social robotics team supported a trial of the Pepper robot at Townsville Hospital and Health Service conducted by James Cook University and nursing staff between July and December 2018.

The team developed two applications for this study. One involved Pepper acting as a concierge in the Emergency Short Stay Unit providing answers to frequently asked questions. The other involved Pepper standing in the main foyer area asking people about their attitudes to flu vaccinations, and providing public health information about the virus, vaccination, and the importance of practices like hand hygiene in preventing the spread of disease.

Despite some initial scepticism around the presence of a robot in the hospital, Pepper was very well received.

PHOTO by Ian Hitchcock: Pepper with staff at Townsville Hospital, from left: Anne Elvin (Enrolled Nurse), Professor Cate Nagle (Professor of Nursing & Midwifery, James Cook University), Dr Wendy Smyth (Nurse Manager - Research, Townsville Hospital & Health Service), Chris McIntosh (Nurse Unit Manager – Emergency Short Stay Unit, Townsville Hospital)

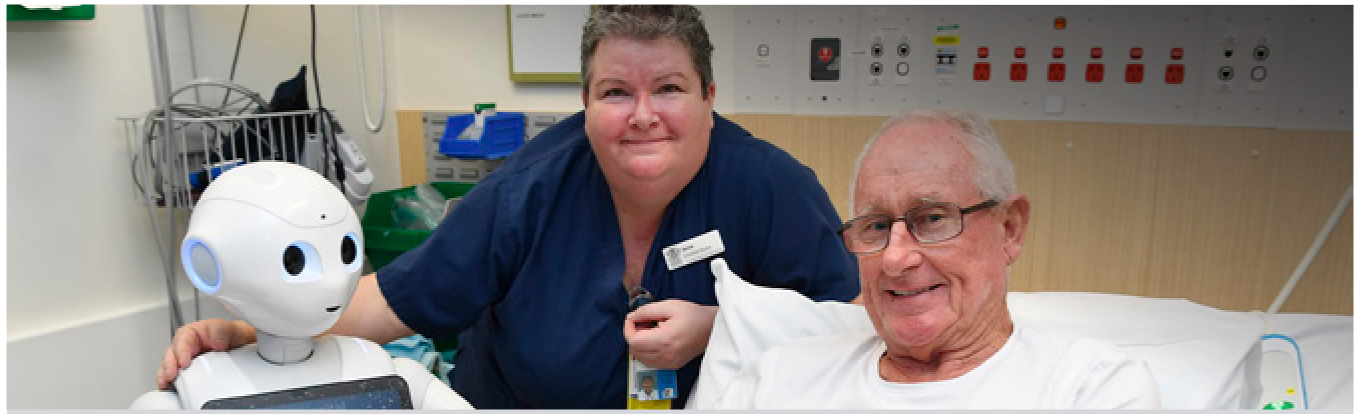

PHOTO by Ian Hitchcock

OBSERVATIONS FROM THE FIELD

It's critical to have a champion on the ground when new technology is launched, and it was the idea of enrolled nurse Anne Elvin to introduce a robot into the midst of a busy working hospital. Together with nurse unit manager Chris McIntosh, Ms Elvin's enthusiasm and persistence were key to overcoming privacy, ethical, technical, and procedural issues involved in introducing a new technology into a public hospital.

Ms Elvin, who acted as Pepper's chaperone during the hospital trial, said: "We've found it a really positive experience for patients, visitors and any of the staff who engage."

She observed that while some people were initially unsure of the robot, they tended to warm to it very quickly, particularly when they noticed small hand movements, face tracking and simulated blinking. These features contribute to the behaviour Pepper's creators at SoftBank term 'autonomous life'.

People who engaged with Pepper also seemed amused and intrigued to be corrected by the robot if they responded incorrectly to a question (for example, about hand hygiene or the flu virus), with Pepper saying "that's not quite right" before delivering the correct information.

Even a mild scolding appeared acceptable from the robot. Patient Rodney Whiting said: "Pepper is not backward about coming forward. I asked her a smoking question and I got the lecture about smoking." That lecture was apparently even more effective on another patient who was overheard asking nursing staff for nicotine patches because "Pepper told me smoking is bad for my health, and you can help me quit."

People were also quite tolerant when Pepper misheard their questions or responses. Ms Elvin recalled: "An emergency department consultant once asked Pepper where to find a doctor, which was not a line we had it programmed to answer, and Pepper replied with information about where the shops and cafes were. Everybody had a big laugh because it was kind of accurate."

"Pepper is not backward about coming forward. I asked her a smoking question and I got the lecture about smoking."

On a more serious note, the trial of Pepper in Townsville Hospital has inspired many ideas from healthcare professionals about how social robots could assist in their field. It also helped allay fears about the future role of robots in healthcare. Staff, patients and visitors were reassured that while Pepper might be informative, entertaining and indefatigable, robots will never be a substitute for humans in the most important roles.

In the words of trial partner, head of nursing and midwifery at James Cook University, Professor Melanie Birks: "As much as robots might have 'personalities' they can't provide the level of caring and therapeutic engagement that nurses and other health professionals provide."

Something to chew over

In November 2018, Pepper was put to work to advocate for healthy habits starring alongside nine-year-old Juiced TV presenter, Andie Rose, in an important educational video for Children's Health Queensland.

Filmed at the Centre, Pepper and Andie's conversation about healthy eating will be the first video as part of a public outreach program under Children's Health Queensland and Queensland Health Queensland Child and Youth Clinical Network - Growing Good Habits program. Robots just like Pepper have huge potential to promote and encourage kids to learn more about healthy eating and exercise habits in a way that's fun and entertaining for young people.

As a follow-on project, the Centre is working with clinical dietitians to scope an application to enable the use of Pepper as a healthcare advocate for encouraging children to develop and maintain healthy eating habits.

Pepper gets moving

So far, deployments of SoftBank's Pepper robot outside laboratory settings have typically involved the robot standing in place or moving within a constrained area or along a predetermined path. Indoor robots with more sophisticated navigational capabilities usually rely on expensive sensors, such as scanning laser range finders, which Pepper does not have.

The ability to move in the world is important for social interaction so that, for example, Pepper could visit patients in a hospital or residential aged care facility and engage in meaningful dialogue and interaction based on the situation. This would take into account where the patients or residents are, what they're doing, and what's happening around them. To do this, social robots like Pepper will first need to be able to competently and safely navigate, and then detect people who are standing, walking or engaged in a range of activities.

As a first step, the social robotics team has developed a system to enable Pepper to move safely and navigate reliably in a known indoor environment. The team has mapped this environment using CSIRO's handheld mobile mapping technology (commercially available as ZEB-REVO from GeoSLAM Ltd).

Pepper then performs localisation, path planning, and mobile navigation using information from an in-built depth sensor and point lasers in lieu of a scanning laser range finder.

The navigation system has been demonstrated with Pepper giving tours of the open plan office and laboratory spaces at QUT to numerous visitors to the Centre. The technical components of this work were presented in a paper at the Australasian Conference on Robotics and Automation (ACRA) 2018, in Canterbury, New Zealand. The software resources developed – which include a virtual machine, ROS core, and improved motion controllers for Pepper – have been made available as open source in a bitbucket repository. This has so far been accessed more than 250 times by researchers in Australia, New Zealand, Japan, Europe and the US.

BEYOND THE METRIC MAP

One disadvantage of the approach to indoor robotic navigation is that it first involves generating a comprehensive map of the environment in which the robot will operate. This requires considerable technical expertise. While that might be feasible for a large institution like a university or major hospital, any approach where a technical expert is required to configure the system doesn't readily support introducing a robot like Pepper into a smaller health clinic, temporary venue or private home.

Another challenge arising in traditional map-based robotic navigation is that any unmapped objects are treated as 'obstacles' to be avoided – and that includes humans. A social robot needs to not only recognise humans, but interact in a socially meaningful way.

To address these challenges, the team is developing a system for visual navigation so that Pepper can be shown around an indoor environment and thereafter use visual cues to find its way around. The aim is to ultimately enable a person who does not have any robotics expertise to configure the system by physically guiding the robot around and pointing out important locations; much as you might expect to do with a visiting human.

In 2018, the social robotics team successfully developed and tested appearance-based navigation capabilities for Pepper based on features in a set of reference images, and Image-Based Visual Servoing (IBVS). Using this technique, the robot can navigate within the Centre's open plan office without a metric map, using only visual features such as lines and corners.

A social robot needs to not only recognise humans, but interact in a socially meaningful way.

The first step involves 'showing' or leading Pepper around the space by holding the robot's hand. The system then automatically selects and stores a set of reference images which it subsequently uses to navigate. The navigation environment is represented internally by a set of reference images with overlapping landmarks organised in a topological graph defining the path the robot should take.

The robot can subsequently find its way around by comparing what it can see currently with its set of nearby reference images. The advantage of this technique is that it doesn't require a global metric map or expensive sensors. It could be used to intuitively teach a robot to navigate when installed in a real-world healthcare setting without input from a robotics expert.

Importantly, the team successfully implemented code onboard the robot so that it is not dependent on a reliable Wi-Fi connection to a controlling computer to be able to navigate. This means the entire process of image acquisition, processing and motion control is done inside the Pepper platform.

Work is underway on the task of detecting and moving around obstacles which were not present during the initial tour. That way Pepper will not be bamboozled when someone moves the furniture.

Medical and Healthcare Robotics

Escalating use of robots and robotic technology in the healthcare space is a global trend. Its adoption is seen as a way of increasing the quality of healthcare, improving recovery time and reducing the need for further medical intervention.

The actual deployment of robotics in healthcare, however, is challenging on a number of fronts: technical requirements, ethical and governance protocols, acceptance (by medical practitioners and patients) and testing (regulations).

Development of robotics in medicine also relies on the fusion of teams from different science and engineering backgrounds. The **Medical and Healthcare Robotics Group at the Centre** is comprised of engineers, scientists, biologists and industrial designers – all working together to solve very complex problems.

The group is co-led by Chief Investigator Jonathan Roberts and Associate Investigator Ross Crawford, who is also an orthopaedic surgeon.

MINIMALLY INVASIVE ORTHOPAEDIC SURGERY

Unlike 'open' surgery, which involves cutting multiple tissue layers in order to access the surgical area of interest, minimally invasive surgery is conducted via small incisions. This aims to reduce surgical trauma and post-operative recovery time. Despite increasing demand for minimally invasive surgery, there are some common drawbacks. They include limited access to the surgical site and a reduced field of view; lack of haptic (feeling) feedback for the surgeon; loss of stereo vision and depth perception; diminished hand-eye coordination; and prolonged learning curves and training periods. This leads to increased operation time and cost to the health system.

Detailed and accurate 3D reconstruction of a surgical site is important to the medical robotics research community. An advantage of seeing the surgical site in 3D is that this data can be used by a surgeon or robot to aid planning before the surgery takes place. This information can also be used with pre-operative MRI and CT scans to further improve planning.

The Medical and Healthcare Robotics Group is undertaking a series of projects to tackle many of these challenges.

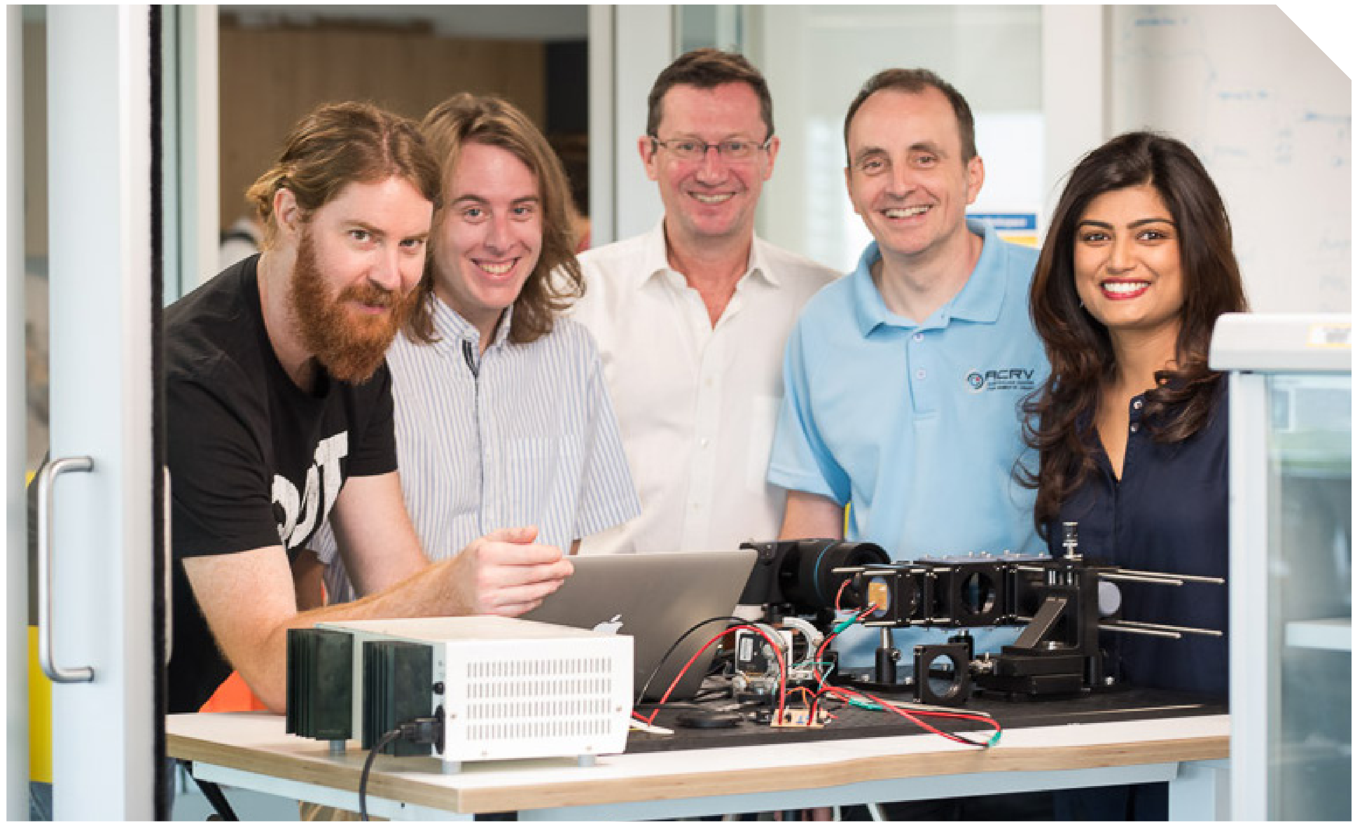

AUSTRALIA-INDIA STRATEGIC RESEARCH FUND PROJECT (AISRF) ON INTELLIGENT ROBOTIC IMAGING SYSTEM FOR KEY HOLE SURGERIES

The AISRF53820 project started in 2017 as a collaboration between Centre partners, QUT and University of Adelaide, and the Healthcare Technology Innovation Centre at Indian Institute of Technology (IIT) in Madras. Its goal is to provide real-time 3D reconstruction to circumvent the apparent lack of perception among surgeons training to become minimally invasive surgical specialists.

The project is using robotic vision techniques applied to ultrasound video streams from 4D ultrasound – think of a video made using an ultrasound machine instead of a video camera. The images from the system will be automatically segmented and labelled enabling tissue visible in an endoscopic camera view to be annotated in real time. Centre Research Affiliate **Dr Ajay Pandey** is the project leader and, within the project, is the 3D vision and intelligent imaging system lead.

"Work on the Australia-India project extends to exploration of the application of artificial intelligence (AI) in medical endoscopic image augmentation and segmentation"

As part of the AISRF53820 project, **Centre PhD Researcher Shahnewaz Ali** is developing an Intelligent Imaging System for surgeons performing minimally invasive surgery using an endoscope. Ideally, the vision system will infer potential information such as depth through the captured images and label what becomes visible during surgery, such as tissue type.

This work started in 2018 and has explored how to control illumination in underwater environments when a camera is moving to ensure quality of images in terms of image context. Work in 2019 will involve development of techniques to estimate depth from endoscopic images.

One of the biggest challenges in knee arthroscopy is the 2D limited field of view provided by the arthroscope. This can lead to complex surgical manoeuvres and, potentially, unintended damage to sensitive knee structures. At present, ultrasound imaging is the only real-time volumetric imaging modality clinically available in operating theatres. As part of the project, QUT **PhD student Maria Antico** is working with **Centre Chief Investigator Gustavo Carneiro** (University of Adelaide node) and QUT Faculty of Health **Senior Lecturer Davide Fontanarosa** to develop techniques that use an ultrasound imaging system for guidance in knee arthroscopy. Their work will create, intra-operatively, a full real-time map of the knee to safely and precisely guide surgical tools in the knee joint and possibly to improve the surgical outcome.

Although there is currently no clinically available system using ultrasound guidance for arthroscopic procedures in the knee, several studies in literature have investigated procedures that could be adapted to (autonomous) arthroscopic applications. The team is currently working on using 3D/4D ultrasound imaging for identification and time-resolved localisation of knee structures of interest to make 3D/4D ultrasound guidance in robotic knee arthroscopy possible.

The project extends to exploration of the application of artificial intelligence (AI) in medical endoscopic image augmentation and segmentation. This work is being undertaken by QUT **Postdoctoral Research Fellow, Yaqub Jonmohamadi**, who joined the project in 2018. Application of augmented reality has become a common practice in surgery, particularly in minimally invasive surgery. During the past few years, automatic image segmentation techniques based on AI have proven to outperform humans in diagnosis of medical images. Besides diagnosis applications, currently, AI is finding its way in the real-time surgical operations by way of augmented reality.

Work so far has shown automatic segmentation of the femoral cartilage in images, achieved by training a deep neural network on 1,100 manually-contoured images. Recently, the group's embedded clinicians have contoured a further 3,500 images. In 2019, using the newly labelled data, it is expected to expand the auto segmentation process for other structures such as ligaments and meniscus.

AISRF53820 project progress meeting at IIT-Madras. From Left: Dr Keerthi Ram, Prof. Mohansankar Sivaprakasam, Centre Research Affiliate Dr Ajay Pandey and Mr Ramdayalan Kumarasami).

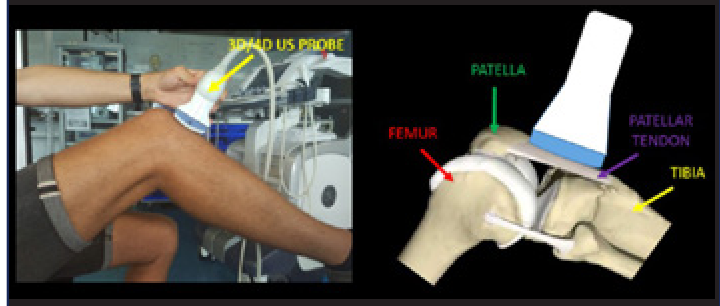

Photo: [Left] Lateral view of the knee joint with an ultrasound probe placed on the patellar tendon. **[Right]** Schematic ultrasound probe positioning representation, showing the positions of reference structures relative to the probe. In the final set-up, the probe will be held in position by a robotic arm.

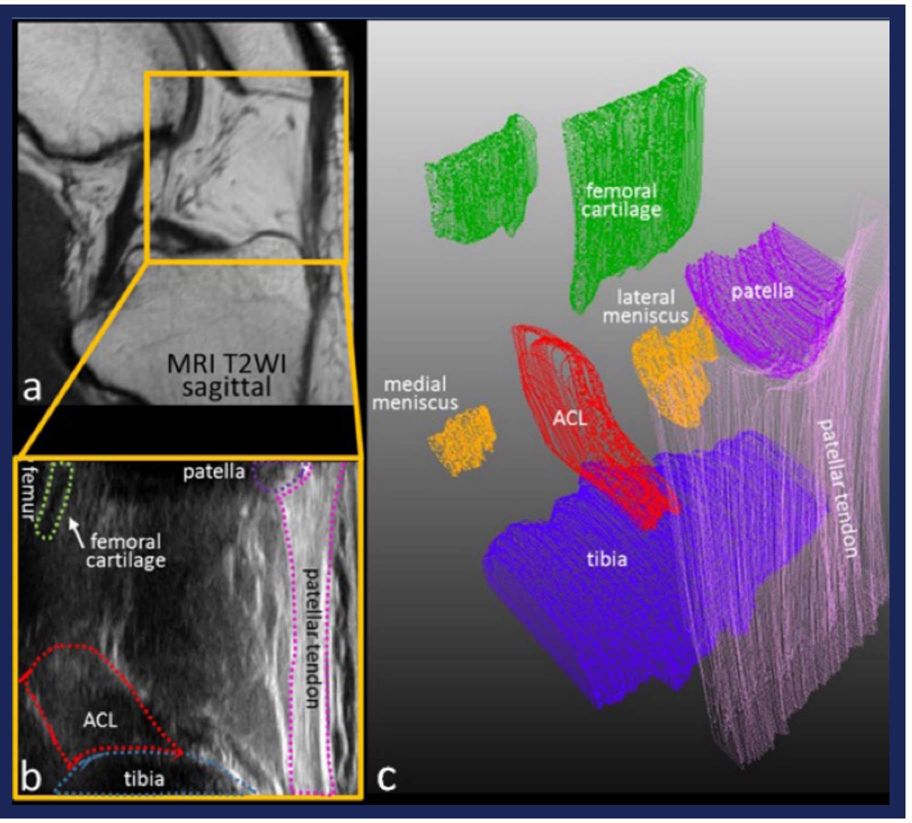

MRI (a) and ultrasound (b) of a volunteer's knee, flexed at 60 degrees. The yellow boxes in (a) and (b) highlight the corresponding section in the two modalities, as 3D/4D US acquires a smaller field of view compared to MRI. 3D segmentation of the relevant structures in (c): femoral cartilage, ACL, menisci, patellar tendon, patella and tibia. The anatomy structures were manually contoured on the ultrasound volume by an experienced surgeon. The location of the structures were checked by visually inspecting and comparing the MRI and ultrasound images. The contours are highlighted by the same colours and labels on the ultrasound image (b).

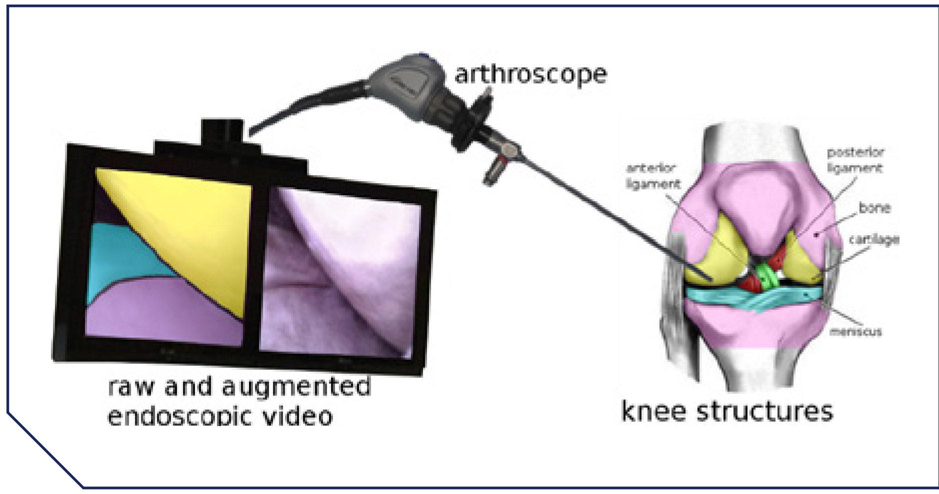

Augmented reality in knee arthroscopy: a sample raw arthroscopic image on the right side of the surgical screen and its augmentation on the left. The colour coding represents different structures of the knee anatomy (as shown in right-hand figure).

The Australia-India project is not the only project focused on the development of robotic assistance for minimally invasive surgery. A number of our other researchers are also working on associated projects.

Centre Research Affiliate Dr Liao Wu is developing the second iteration of a hand-held snaking hip arthroscopy tool. It involves novel application of a variable neutral-line arm constructed using printed discs and actuated by tendons. His work also aims to improve the current model of this arm type to allow for controlled rotation at the tip and collision detection using tendon displacement and tension measurements.The entire robot can be held in the hand by a surgeon.

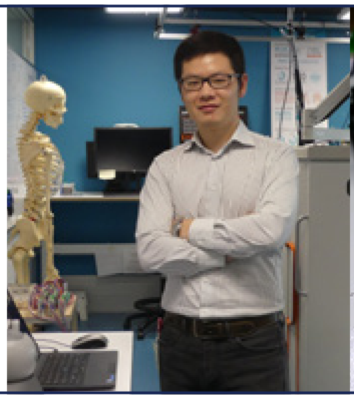

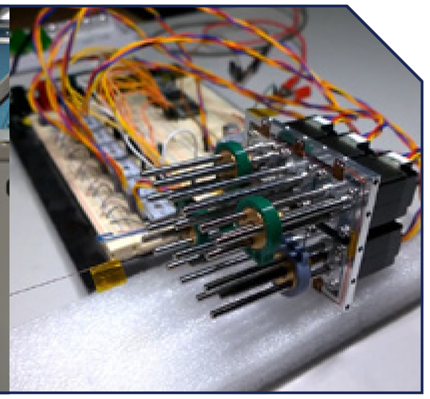

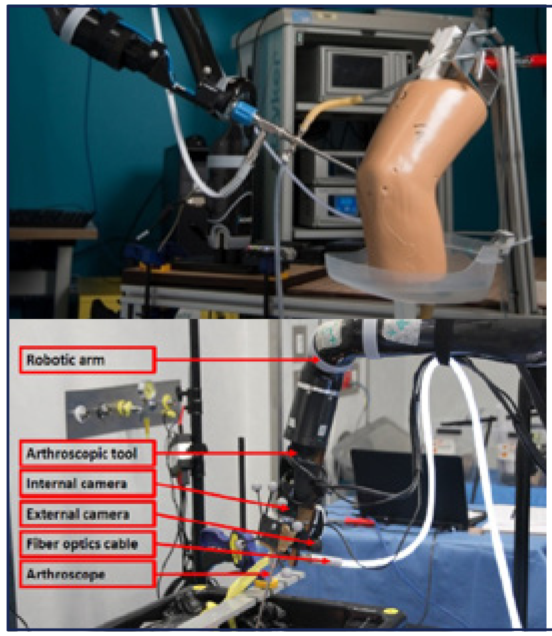

Centre PhD Researcher Andres Felipe Marmol Velez is working on the development of a vision-based robotic assistant for arthroscopy. This system focuses on assisting surgeons during complex and strenuous manipulation of instruments to reduce risk of unintended injury to the patient and to the surgeon. Its hardware is composed of two visual sensors rigidly attached to the end-effector of a robotic arm; exploited within a hybrid Simultaneous Localisation and Mapping (SLAM) framework designed to overcome common challenges in the arthroscopic context.

The system's performance has been demonstrated with extensive experimentation in both knee models and cadaveric tissue. Experiments have included various challenging manoeuvres, image settings and operative conditions common to arthroscopy. Results show the system is capable of localising the scope robustly and reliably inside the body. Furthermore, the internal anatomy is mapped with great detail and sub-millimetre accuracy. Current intra-articular maps allow surgeons to readily recognise anatomical areas of interest and quantify their morphology. Future work seeks to exploit the localisation and mapping estimates to actively navigate the inner anatomy, allowing for automatic tissue/tool tracking capabilities.

Centre PhD Researcher Artur Banach is developing a robotic-imaging system based on machine learning, computer vision and robotic control techniques. It will assess distance between tools and anatomy to assist a surgeon or robot move surgical tools inside the knee cavity without unintentional contact.

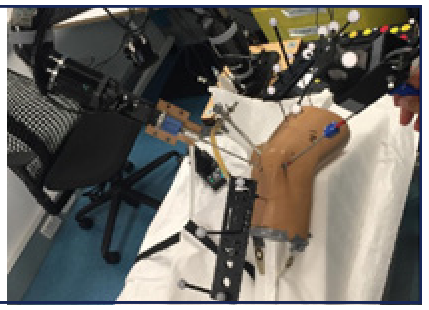

Snake-like Tools for Knee Arthroscopy

Centre PhD Researcher Andrew Razjigaev is developing snake-like robots for knee arthroscopy, as demonstrated to the Royal Family's STEM advocate, His Royal Highness The Duke of York, on a visit to QUT in November 2018.

Many of the current surgical tools for knee arthroscopy are straight and rigid, making them difficult to manoeuvre during surgery.

Snake-like robotic arms for minimally invasive surgery can be very flexible and dexterous manipulators. They are ideal in teleoperation because the surgeon can handle organic tissue with ease. The challenge is to design robotic manipulators for this task and to control them with computer vision in surgery by a surgeon.

In 2018, work was undertaken towards designing tendon-driven snake-like tools for the procedure. A new process for designing a robot was created in simulation. Designs for a snake-like robot were assessed to see how dexterous the arm is at accessing target regions in an anatomical structure that is provided by a patient's knee scan.

The simulation took account of obstacles inside the knee and ensured that the robot arm was capable of bypassing them to reach the target. Using an evolution algorithm, the robot joints of the snake-like robot were optimised for better dexterity by letting them 'evolve' for the simulated tasks. The evolution resulted in a snake-like tool that scored the best dexterity for knee surgery.

There's still further work to be done, with the focus in 2019 to integrate a flexible snake-like robotic arm to the Raven II surgical system in a teleoperation control scheme.

Image 1

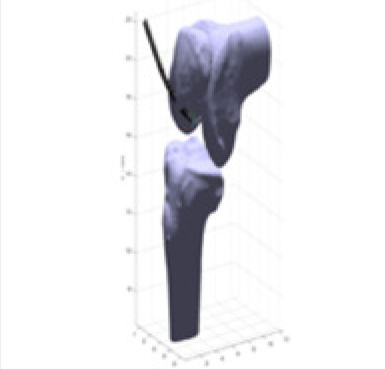

Image 2

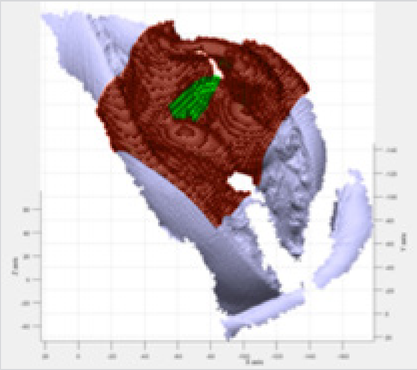

Image 3

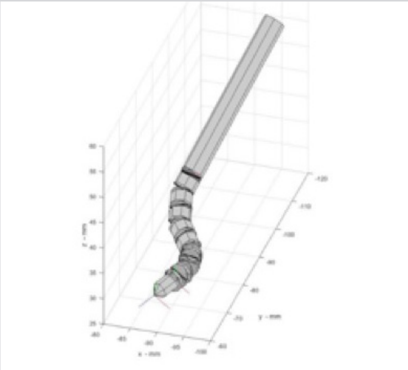

Image 1 and 2 show the conversion of knee 3D points into voxels for the obstacles (red cubes) and target region (green cubes) for the simulation. **Image 3** shows a plot of a snake-like robot

Every eye tells a story!

Centre Research Affiliate Dr Anjali Jaiprakash – an Advance QLD Research Fellow working in the Medical and Healthcare Robotics Group – is leading a team to develop an innovative retinal diagnostic platform. It aims to improve rates and accuracy of retinal screening, enable personalised early diagnosis of eye disease and reduce preventable blindness.

At the heart of this platform is a revolutionary retinal imaging device, the **Retinal Plenoptoscope.**

This non-invasive, portable, user-friendly product employs advanced light field technology to create patient-specific 3D imaging data. Sophisticated machine learning algorithms will enable robust, cost-effective analysis and accurate diagnosis.

The Retinal Plenoptoscope (light field fundus camera) captures both information about the intensity of light in a scene and the direction light rays are travelling in space.

This contrasts with a conventional camera, which records only light intensity. With this design, images can be refocused with a much higher spatial resolution than with images from the standard plenoptic camera.

The Retinal Plenoptoscope will significantly improve the affordability, reliability and portability of retinal imaging without dilating the eye while also delivering new diagnostic features, the ability to create 3D views, metric measurement of retinal features that is patient-specific and post-acquisition refocusing.

DID YOU KNOW?

- 80 per cent of visual impairment is preventable or curable with regular monitoring and timely treatment. Despite this, an estimated 285 million people world-wide are living with preventable eye disease.
- Over 300,000 Australians have visual impairment, with rates of blindness six time higher in Aboriginal and Torres Strait Islander communities.
- Diabetes-related blindness is increasing in line with the global obesity epidemic.
- In 2010 there were about 422 million patients with diabetes (8.5 per cent of the global population). In Australia, there are at least 1.5 million people with diabetes. Alarmingly, half do not get their recommended eye checks. It is estimated that in 2002 diabetic retinopathy accounted for about five per cent of world blindness.

Dr Anjali Jaiprakash, winner of the 2018 MIT Technology Review Innovator Under 35 (Asia Pacific), shares the team's vision to reduce preventable blindness on a TedX platform before a 5,000-strong audience in Sydney.

This new class of retinal imaging system brings together human-centred service design to ensure services and technologies meet user and societal requirements. Of note, the team is working with Aboriginal and Torres Strait Islander people living with diabetes and their associated healthcare professionals. This involves mapping current diabetic retinopathy eye screening systems and experiences in primary healthcare.

With proof of concept and protectable IP in place, the Retinal Plenoptoscope is being developed into a product viable for commercialisation, targeting improved health outcomes in areas of greatest need, particularly remote communities.

Additionally, the team is investigating how computer vision can be incorporated into the RAVEN-II surgical robotic platform to improve the speed and safety of eye surgery procedures such as capsulorhexis, phacoemulsification and retinal vein cannulation.

Revolutionising cancer detection

Centre Chief Investigator Gustavo Carneiro, an expert in deep learning based at the University of Adelaide, has developed an algorithm that can rapidly detect breast tumours 1.78 times faster than existing cancer detection methods.

"The interesting part is that our algorithm is based on reinforcement learning," said Professor Carneiro.

"This is the same method commonly used to train robots by providing rewards or penalties as they interact with their working environments in a positive or negative manner."

Reinforcement learning is a form of AI that enables a computer – or robot – to learn how to perform complex tasks without being programmed by humans.

Professor Carneiro and his previous PhD student, Gabriel Maicas Suso (currently a research associate at the University of Adelaide), applied reinforcement learning to train a computer program to independently analyse breast tissue with a relatively small amount of data.

In conjunction with an MRI scan, the autonomous program operates similar to the video game, Tetris. It uses a green square to navigate and search an MRI image for lesions. The square turns red if a lesion is detected.

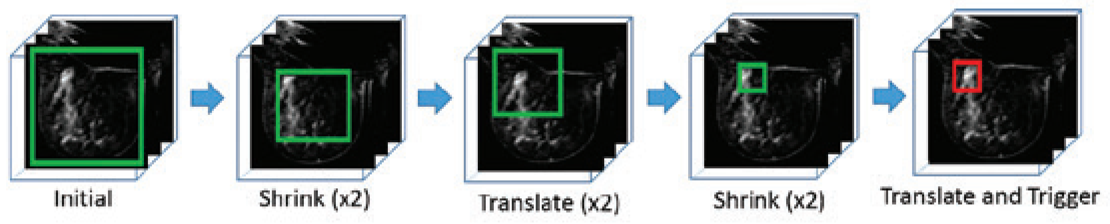

"By incorporating deep learning into medical imaging analysis, we have developed a program that intuitively locates lesions quickly and accurately," Professor Carneiro said.

"More research is needed before the program can be used clinically. Our ultimate aim is for this detection method to be used by radiologists to complement, support and assist their important work in making a precise and quick diagnosis.

"The potential to use AI in the imaging medical field is boundless."

In other work, Professor Carneiro has progressed a deep learning algorithm that uses more than 500 images to help detect and classify masses in mammograms.

The new method enables a whole breast MRI volume, as opposed to individual masses, to be classified into two classes: malignant or non-malignant.

"The interesting aspect of this research lies in the way we model the classifier that mimics how radiologists are trained," Professor Carneiro said.

"It's trained to solve several classification problems of increasing difficulty instead of solving only the question of malignant versus non-malignant findings."

In 2018, his research interest in medical image analysis and deep learning also led to the development of a classifier that can estimate the quality of an ultrasound image scanned for prostate cancer radiotherapy. This work was conducted in collaboration with QUT researchers.

"This is important because these images assist in the radiotherapy treatment, so they must have high quality in order to maximise the efficacy of the treatment."

In 2019, Professor Carneiro will work on methods that automatically find new medical imaging biomarkers for the classification of not only cancer, but other diseases. He foresees the exciting part of this new direction is the possibility of discovering biomarkers that are currently unknown for the medical community.

Flying robots to the rescue

CASE STUDY

The world's biggest outdoor airborne robotics challenge, the UAV Challenge

It's renowned as the toughest rescue mission on the planet. For good reason; pushing the real-world capabilities of flying robots on the frontline towards the ultimate goal of saving lives.

Organised by QUT, with the support of the Australian Centre for Robotic Vision, and CSIRO's Data61, the UAV Challenge stakes its claim as the world's biggest outdoor UAV competition (number of global teams involved). It transforms the rural Queensland community of Dalby into a 'global robotics capital'.

In 2018, 11 teams from across Australia, Poland, Thailand, the Netherlands and India descended on Dalby, qualifying to contest the biennial Medical Express challenge from a pool of 55 entries. The Mission Impossible-esque competition requires unmanned aerial vehicles (UAVs) to not only autonomously land in difficult, unseen conditions (having flown up to 30km from take-off), but return swiftly and safely with precious cargo. Namely, a blood sample from a stranded, fair dinkum dummy, Outback Joe.

Outback Joe's perilous life-or-death situation is further amplified by the addition of simulated flash flooding. Meanwhile, in the air, UAVs must contend with potential air strike from magpies during swooping season, not to mention simulated obstacles like commercial planes.

Over the past 12 years, the visionary event has drawn an eye-boggling array of agile, low-cost UAV lifesavers. Flying robots of all shapes and sizes, weighing-in anywhere from 3kg (the weight of a brick) to around 20kg (think three bowling balls or checked-in luggage) and resembling everything from mini planes to UFO-like contraptions.

All of them tasked with the white-knuckle challenge of navigating remote terrain, bad weather and an onslaught of obstacles (terrestrial and airborne), to aid a person in need of emergency medical assistance.

"There are always quite a few white-knuckle moments, as the teams work to reach poor Joe," said Australian Centre for Robotic Vision Chief Investigator Jonathan Roberts; one of the brains behind the UAV Challenge.

"But it's always a lot of fun and very rewarding. Importantly, the challenge is all about developing technology to help save lives in the real world.

"Over the years, the UAV Challenge has inspired advances in drone design as well as software and communications systems, not least being enhancements to the functionality and codebase of open source autopilot software, ArduPilot, which is now embraced by major players such as Microsoft and Boeing."

While the biennial Medical Express won't be held again until 2020, the UAV Challenge also hosts an annual high school Airborne Delivery Challenge. This requires teams to use a UAV to drop an EpiPen to a stranded individual, desperate and suffering an allergic reaction.

In 2018, a group of Californian high school students were crowned champions of the Airborne Delivery Challenge in a field of 12 teams from across Australia, the United States and South Korea.

The biennial Medical Express challenge, meanwhile, lived up to its 'Mission Impossible' moniker, with the winner's trophy remaining unclaimed in 2018.

"They came, they almost conquered and I can guarantee you they'll be back," said Professor Roberts.

"Trying to overcome the seemingly impossible is all part of the challenge. Cracking the Medical Express challenge really does involve the superhuman. And that's exactly where we want to take the technology.

"The whole point of the competition is to advance UAV design, software and communications systems so that these flying robots can go where it is too dangerous or remote for human-led rescue efforts alone.

"Two of the teams came very close in 2018 and that's a huge achievement in itself."

The 2018 UAV Challenge was sponsored by: Queensland Government; Insitu Pacific; Boeing Research & Technology – Australia; Northrop Grumman; Lockheed Martin Australia; Defence Science and Technology Group (part of the Australian Government Department of Defence); and MathWorks.

The 2019 Airborne Delivery Challenge moves to Calvert in the City of Ipswich, 55km south-west of Brisbane. The competition for high school participants takes place October 12-13. For more details: **www.uavchallenge.org**

Dialling-in AI in the fight against skin cancer

CASE STUDY

Self-diagnosing skin cancer via a personal full-body scan using a smart phone at home

Artificial intelligence is being enlisted to develop a smart phone app destined to help Australians self-diagnose skin cancer at the earliest possible stage.

The project led by Australian Centre for Robotic Vision Associate Investigator Anders Eriksson – an ARC Future Fellow at QUT in 2018 (moving to University of Queensland in 2019) – promises the ultimate ‘selfie’ thanks to advances in computer vision and machine learning.

Namely, personal full-body skin mapping of moles and lesions. All possible on a smart phone in the privacy of your own home.

“The most critical cue for early diagnosis of skin cancer is changes in the appearance of moles and lesions over time,” Dr Eriksson said.

“The technology we’re developing will allow a person to simply scan their body with a smart phone by taking a large number of high-resolution images. The images will be sent to a remote server for processing and within a matter of minutes, a complete 3D reconstruction of the individual’s body is returned, with every mole and lesion identified, analysed, assessed and compared with previous scans.”

The Australian Centre for Robotic Vision hosts a demo on its cloud computer vision platform, CloudVis (**www.cloudvis.qut.edu.au**)

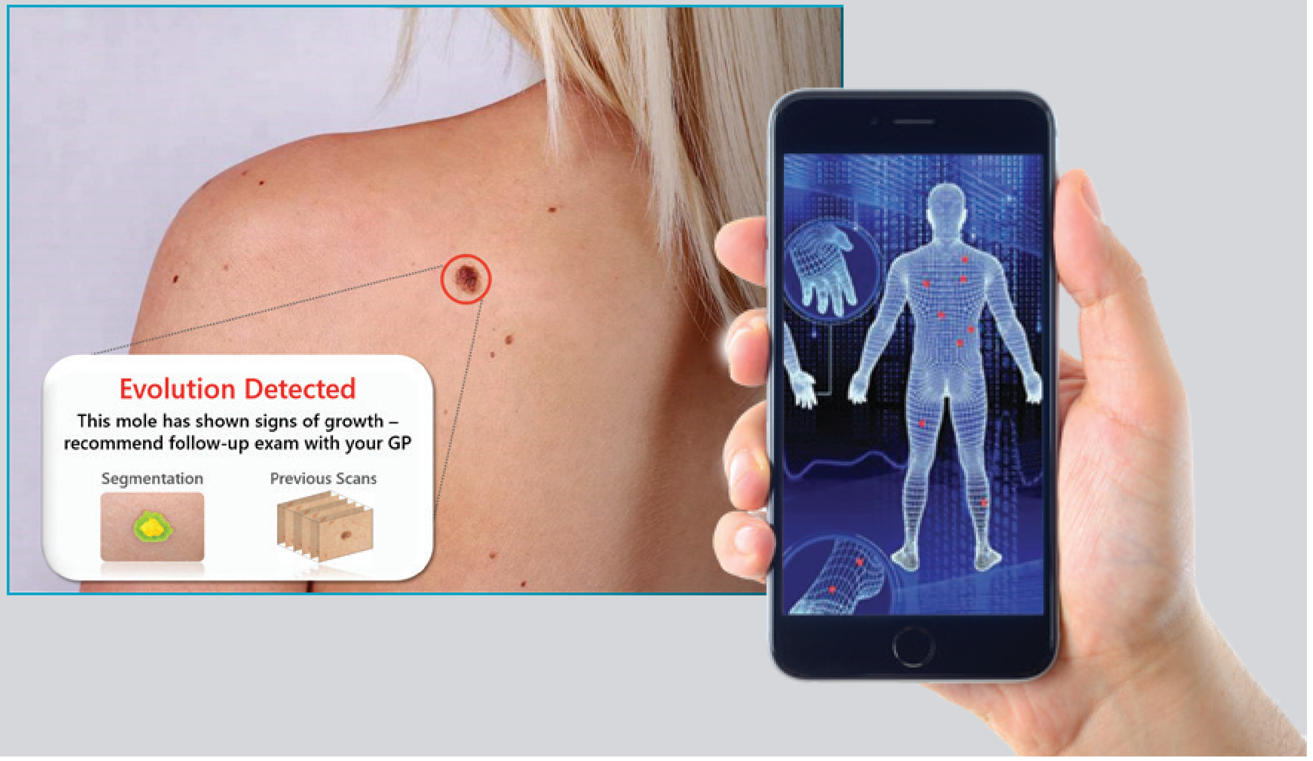

Where we live

The environment in which we live is critical to our wellbeing, comprising both our natural and built environments.

THE NATURAL ENVIRONMENT sustains us with fundamental necessities such as drinkable water and breathable air, but also provides beauty and wonder that nourishes the human spirit.

The natural environment is vast. Today, we can only afford to infrequently inspect a tiny fraction of it, and rely on extrapolation to understand the full condition. Local variations in time and space can be overlooked.

Robotic vision technology has the potential to reduce associated costs, allowing more of the world to be inspected, more frequently. Ultimately, intelligent, vision-enabled robots would be capable of environmental remediation as well as inspection. This is already happening in reef monitoring and management trials on the Great Barrier Reef involving RangerBot – the world’s first autonomous robot designed specifically for coral reef environments.

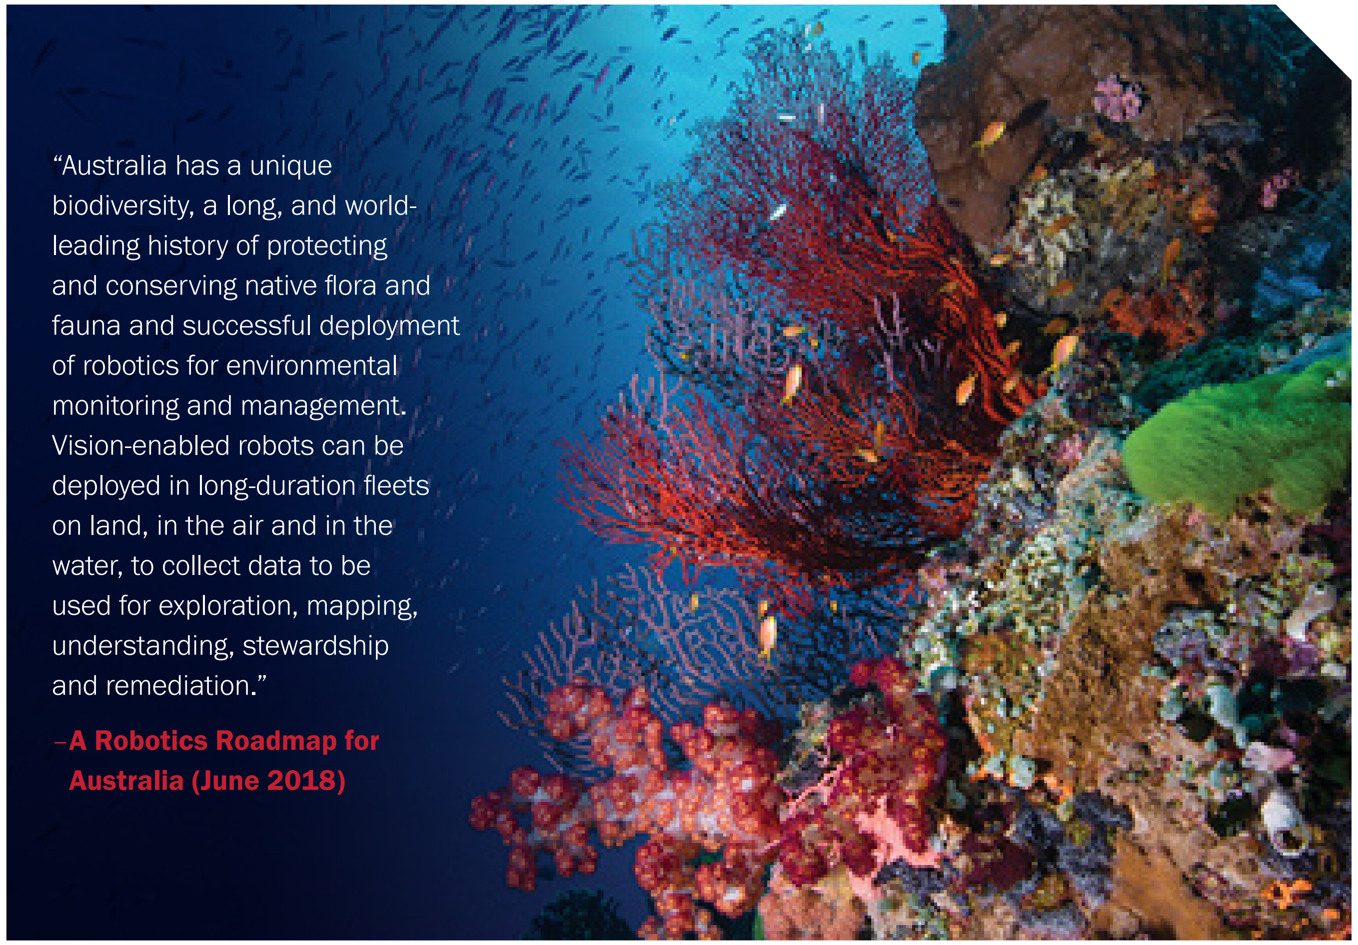

“Australia has a unique biodiversity, a long, and world-leading history of protecting and conserving native flora and fauna and successful deployment of robotics for environmental monitoring and management. Vision-enabled robots can be deployed in long-duration fleets on land, in the air and in the water, to collect data to be used for exploration, mapping, understanding, stewardship and remediation.”

–A Robotics Roadmap for Australia (June 2018)

PHOTO by Tourism and Events Queensland

Romancing the Reef

CASE STUDY

RangerBot the robo reef protector

The world's first robotic-vision empowered coral reef protector, RangerBot, literally helped share the love in a coral seeding program on the Great Barrier Reef.

In November 2018, RangerBot-turned LarvalBot creator, Centre Chief Investigator Matthew Dunbabin, ventured on a field trip to the Great Barrier Reef. The trip was timed during Mother Nature's annual splendour of coral spawning; a phenomenon best described as the Mount Everest of reproduction in nature.

His mission? To put two LarvalBots to work in a world-first coral larvae seeding program, spearheaded by Southern Cross University (SCU), delivering baby coral (or larvae) following incubation in a floating nursery of purpose-built rearing ponds on the reef.

The overall project is led by SCU marine biologist Professor Peter Harrison – an expert in coral reproduction ecology who, in 1981, was part of a research team that discovered the mass coral spawning phenomenon on the Great Barrier Reef. It aims to restore damaged parts of the World Heritage-listed wonder, ultimately speeding up recovery of precious ecosystems affected by coral bleaching.

"Everything went to plan despite challenging weather," said Professor Dunbabin of his part in the project funded through a $300,000 award by the Great Barrier Reef Foundation's Out of the Blue Box Reef Innovation Challenge supported by Tiffany & Co. Foundation.

"For me, it's a dream come true, to be able to help protect and rebuild the world's greatest natural wonder."

In the weeks leading up to the field trip, the roboticist and ardent conservationist worked tirelessly with his team to transform RangerBot into LarvalBot, building innovative and 'life-giving' attachments in his Lab at QUT.

PHOTO by Gary Cranitch, Queensland Museum

PHOTO by Tourism and Events Queensland

PHOTO by Gary Cranitch, Queensland Museum

Ticked off his 'To Do' list: a 'coral baby' sling of sorts, complete with bladder, rigging and a triggered release mechanism allowing larvae to be dispersed across the widest possible area, with the potential to cover 1,500 sq.m/hour per robot.

Fittingly, the larval delivery system for RangerBot, which won the 2016 Google Impact Challenge People's Choice prize as a high-tech, low-cost autonomous underwater robot, largely utilised everyday items from local camping and hardware stores.

"RangerBot is the world's first underwater robotic system designed specifically for coral reef environments, using only robotic vision for real-time navigation, obstacle avoidance and complex science missions," Professor Dunbabin said.

In the not too distant future Professor Dunbabin believes RangerBot will be affordable enough for use in global citizen science projects to help monitor and protect the world's reef environments, similar to the accessibility provided by drones.

"Its transformation into LarvalBot has been a successful test of functionality. What's also exciting is RangerBot can map expansive underwater areas at scales not previously possible, making it a valuable tool for reef research and management."

RangerBot and its predecessor, COTSbot (created in 2015 and displayed at the Queensland Museum in 2018), are nothing short of a labour of love for Professor Dunbabin, who has spent more than a decade devising and refining the technology.

A roboticist who describes himself as akin to the Babel fish from *The Hitchhiker's Guide to the Galaxy* – the fictitious alien fish that performs instant translations – Professor Dunbabin reckons his greatest talent lies in his ability to help bridge the gap between the technical (robotics) and the natural sciences and how each can help the other.

Did you know?

RB2 (RangerBot Number 2) – one of the two robots that moonlighted as LarvalBots on the Great Barrier Reef – went on to help a Centre-supported student team take second place in the coveted biennial Maritime RobotX Challenge in Hawaii (December 2018).

In this competition, RangerBot was modified to step up as ‘Thunder4’, providing underwater support to an Autonomous Surface Vehicle or ‘robot-boat’ nicknamed ‘Bruce’. RangerBot’s support role was possible after Maritime RobotX Challenge organisers introduced a new requirement in 2018. This called on competitors to construct a ‘System of Systems’ with robots operating in multiple domains, including the ability to sense and act underwater.

As a result, RangerBot was put to work as the team’s eyes beneath the water, while *Bruce* played the lead role as ‘robo mother ship’. That is, a 16-foot WAM-V (Wave Adaptive Modular Vessel) provided as a shell to each team (ahead of the competition) to fit with a system of sensors, software and hardware that work together to enable its autonomous capabilities.

“In addition to fully autonomous movement, *Bruce* must be smart enough to make decisions regarding navigation and mapping,” Professor Dunbabin said. “On-board intelligence developed by the team, enabled it to successfully perform tasks like obstacle avoidance; signal recognition, including reading light sequences to determine whether to dock or circle buoys in a certain direction; obstacle detection; and recovery of objects located above and below the water’s surface. All without human intervention.”

Professor Dunbabin said the Maritime RobotX Challenge, held biennially since 2014, provided a ‘fantastic technical opportunity for students to network and gain exposure to an incredibly important industry that is all about saving lives, moving freight and protecting natural environments’.

“Ongoing advances in autonomous maritime vehicles or ‘robo-boats’ will be important in real-world emergencies, enabling search-and-rescue missions in all conditions, particularly when human-led efforts are thwarted by dangerous or risky scenarios,” he said.

And finally, as for the inspiration behind RangerBot’s alter ego, *Thunder4*, playing cameo to *Bruce*?

“Thunderbird 4, of course! The underwater rescue vehicle in *Thunderbirds*,” said Professor Dunbabin of the cult British sci-fi classic, set in the 2060s, following the adventures of International Rescue, a life-saving organisation equipped with technologically advanced land, sea, air and space rescue craft.

“RangerBot, like Thunderbird 4, is the same superhero yellow and has its own superpower of robotic vision. What’s worrying, however, is some of my young PhD students haven’t heard of *Thunderbirds*. So, clearly they’ve got a bit of retro TV to binge watch.”

THE BUILT ENVIRONMENT covers our cities, roads, dams, tunnels, pipelines and electrical distribution networks. It represents a massive investment over generations that needs to be maintained and enhanced and passed on to future generations.

Maintaining these assets in good condition requires continuous inspection. Traditional labour-intensive methods are expensive. Intelligent vision-capable robots have the potential to reduce the cost of inspection, perhaps even allowing more frequent inspection than could be imagined today. Ultimately, robotic vision technology would be capable of repair as well as inspection.

AI – the ultimate back-seat driver!

CASE STUDY

Connected and Highly Automated Driving (CHAD) Pilot

Artificial Intelligence (AI) is about to buckle up as the ultimate back-seat driver in a novel, tech-savvy research project aimed at ensuring automated cars of the future will be smart enough to handle tough Australian road conditions.

The Queensland Department of Transport and Main Roads' (TMR) Connected and Highly Automated Driving (CHAD) Pilot will see QUT researchers test high-tech sensors, cameras and computer on a 1,200km road trip covering a wide range of road and driving conditions.

This research is headed by Centre Chief Investigator, Professor Michael Milford in partnership with TMR and iMOVE Cooperative Research Centre. It will help shape the future of road infrastructure by researching how automated vehicle technology responds to Australian road conditions.

Professor Milford and his team of Centre researchers will take an electric Renault Zoe, fitted with high-tech sensors and computers, on the road trip in early 2019.

"During this trip, you could say AI will become our ultimate back-seat driver," Professor Milford said.

"Engineers at QUT's Research Engineering Facility have developed a research car platform equipped with a range of state-of-the-art camera and LIDAR sensors used in highly automated vehicles.

"So, as we drive, AI will watch and determine if it could perform the same as a human driver in all conditions.

"The big problem that faces highly automated vehicles right now is that they don't drive as well as humans in all possible conditions.

"We're targeting how the car might use infrastructure to help it to drive well."

Professor Milford said current automated car systems, when faced with some of the road conditions Australian drivers deal with daily, either refused to go into automated mode or handed control back to a human driver.

This research project will look at how a highly automated vehicle's artificial intelligence systems copes with Australian road conditions in four main areas: lane markings, traffic lights, street signs and how to determine a vehicle's exact position despite errors that occur with GPS systems in highly built-up urban areas or poor reception areas such as tunnels.

PHOTOS by Anthony Weate

PHOTO by Anthony Weate

Professor Milford said past studies and his team's initial experiments show that automated cars could have difficulties on rural roads which often lacked lane markings on the side or even a centre line.

"A human driving down a rural road knows to stick on the left and they infer or imagine that there is a line in the middle of the road," Professor Milford said.

"But they will also cross that imaginary line to go around obstacles quite freely. That's very hard for an automated car."

Professor Milford said early testing of the system had already revealed how a paint spill on the road from the back of a truck could confuse a self-driving AI system into wrongly identifying it as a lane marking.

During 1,200km of testing, spread over three months, the research team will assess the car's AI data each day to analyse how it responds to the road conditions, lane markings and road signs.

"Robotics and AI are ultimately about enhancing human life in some way," Professor Milford said.

"The primary goal of our research is to determine how current advances in robotic vision and machine learning – the backbone of AI – enable our research car platform to see and make sense of everyday road signage and markings that we, as humans, take for granted.

"So, safety is an obvious off-shoot, but not the focus of this particular study. What's important is understanding how AI performs and potential improvements to both the technology and physical infrastructure as the automated car revolution unfolds."

"…as we drive, AI will watch and determine if it could perform the same as a human driver in all conditions."

The pilot project is part of the Queensland Government's wider Cooperative and Automated Vehicle Initiative (CAVI).

Expected to run for 12 months, this study will assess the degree to which modern sensors and AI techniques used on highly automated vehicles can interpret and understand signage and road markings to inform future development and investment in infrastructure.

"We'll be out on the roads day and night and in all weather conditions to be sure AI is put to the 'real-world' test," Professor Milford said.

Flying high

CASE STUDY

FrontierSI Project with Ergon Energy

Maintaining a network of hundreds of thousands of kilometres of power lines and supporting poles is a daunting, but critical task in the face of annual electric network outages that drain $160 billion globally.

The energy industry has been deploying Remotely Piloted Aircraft Systems (RPAS or drones) for power network inspection within preventative maintenance practices. Improvements in autonomous drone technology will amplify the benefits.

Coupling the Centre's expertise in sensor and perception technology and control with QUT's Research Engineering Facility flight test team, this project succeeded in inventing new autonomous drone technology to enable easier, faster and more reliable power pole inspection.

The project team successfully developed lightweight pilot assist technology using computer/machine vision to automatically detect a drone's distance from and relative location to power poles and wires.

Key to this technology is the autonomous collision avoidance capability. This allows operators to safely inspect distant infrastructure either manually or using the automated functions.

Two automated functions were developed during the project:

/ Transition: the drone navigates between poles autonomously; it knows when it gets to the next pole and notifies the operator.

/ Orbit: the drone can circumnavigate while simultaneously inspecting a pole without the operator guiding it.

End result: an innovative system that reduces pilot cognitive workload by automating the precision flight required for pole inspection and automatically sensing and avoiding any potential collisions with infrastructure.

The project ended in August 2018 with a successful live demonstration of pilot assisted power pole inspection in front of energy sector professionals and project sponsors, including FrontierSI and Ergon Energy.

Next steps include exploration of commercialisation opportunities and payload miniaturisation.

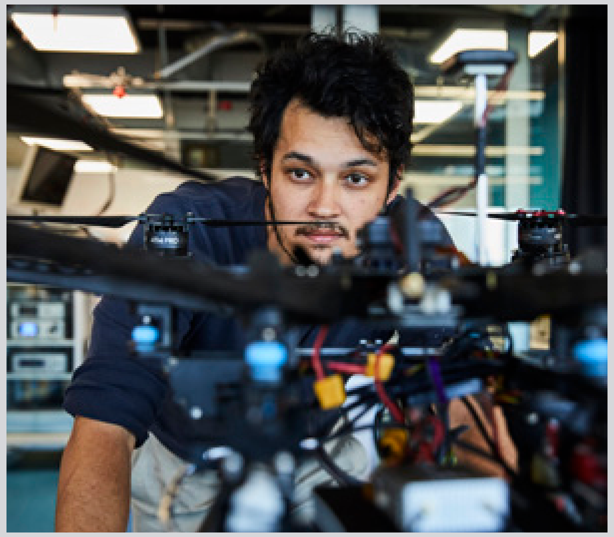

Means to live

Our modern lives are resource intensive, demanding reliable and low-cost supplies of food, energy and minerals. Australia is an abundant producer and exporter of these resources.

Agriculture/farming, with its extensive grazing and feedlot practices, is not highly labour intensive but the work is physically hard and often located in remote areas. Horticulture's production of fruit, vegetables, nuts and flowers vary from highly mechanised to labour-intensive, and is often physically hard work in remote locations. It is increasingly difficult to source workers to carry out this work, which can cause incomplete harvests, wasted produce and an inability to expand production to meet growing global demand. Robotic vision technology has the potential to identify and control weeds, pick fruit, pollinate flowers and even muster cattle, providing a source of support to farmers and their workers by improving production yields and reducing physical risks.

> The Centre is focused on giving the next generation of robots the vision and understanding to help solve real global challenges including sustainable food production

With the global population projected to reach over nine billion by 2050, the Centre is focused on giving the next generation of robots the vision and understanding to help solve real global challenges, including sustainable food production.

Mineral production, in particular mining, is highly mechanised and automation is increasing. The business drivers for this are to increase productivity and maximise the use of capital, as well as reduce machine damage and take workers out of hazardous environments. In this industry, robotic vision technology has the potential to drive vehicles, guide excavation and truck loading, and perform drilling, blasting and survey work in surface, underground, undersea and space mining.

CASE STUDY

Farm-hand Robots

Centre researchers based at The Australian National University in Canberra are developing a novel robotic harvesting platform for multiple crops, starting with asparagus, capsicum and chilli.

“Ultimately, we want to deliver an intelligent and practical robotic harvesting solution that will be ‘truly useful’ for Australian famers and the wider horticultural community,” said Centre Chief Investigator Robert Mahony. “Our platform is at an early stage of development, but the objective is to have commercially viable units available in 2023.”

Professor Mahony said the planned robotic harvester consisted of three main systems: perception, manipulation and transportation. In 2018, Centre researchers focused on developing and refining a robotic vision perception system capable of locating crops among surrounding plants in real-world outdoor scenarios.

“Once complete, this information will be used to optimise the design of the manipulation and transportation systems,” said Professor Mahony. “We aim to run the first on-farm harvesting trials of a single crop – asparagus – in 2019, extending to multiple crop types from 2020 to 2022.”

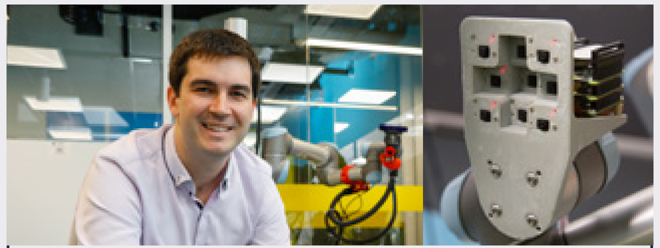

Centre Reseach Affiliate (QUT) Chris Lehnert has turned an exciting new leaf, adding to ground-breaking work on ‘Harvey’ (aka the capsicum-picking robot prototype), developing a 3D-printed camera system with more eyes than a spider. That’s nine eyes or cameras operating at different depths, enabling a robot like Harvey to look around obstructing leaves much like a human.

CASE STUDY

Robotic Vision and Mining Automation

In collaboration with Caterpillar, Mining3 and the Queensland Government, Centre researchers have developed new technology to equip underground mining vehicles to navigate autonomously through dust, camera blur and bad lighting.

The team led by Chief Investigator Michael Milford has completed five field trips to Australian mine sites (including three trips in 2018). Other team members are: Centre Associate Investigator Thierry Peynot, Centre Associated Research Fellows Adam Jacobson and Fan Zeng, and Mining3 industry fellows Tim Hojnik, Umesh Mutubandara and Maciej Matuszak.

Work covers finalising development of a cheap approximate positioning system for tracking mining vehicle inventory underground, and continued development of a more accurate system to improve automation capability of vehicles including Load-Haul-Dump (LHD) loaders.

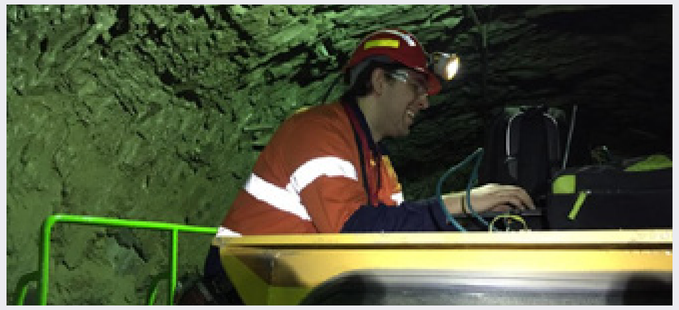

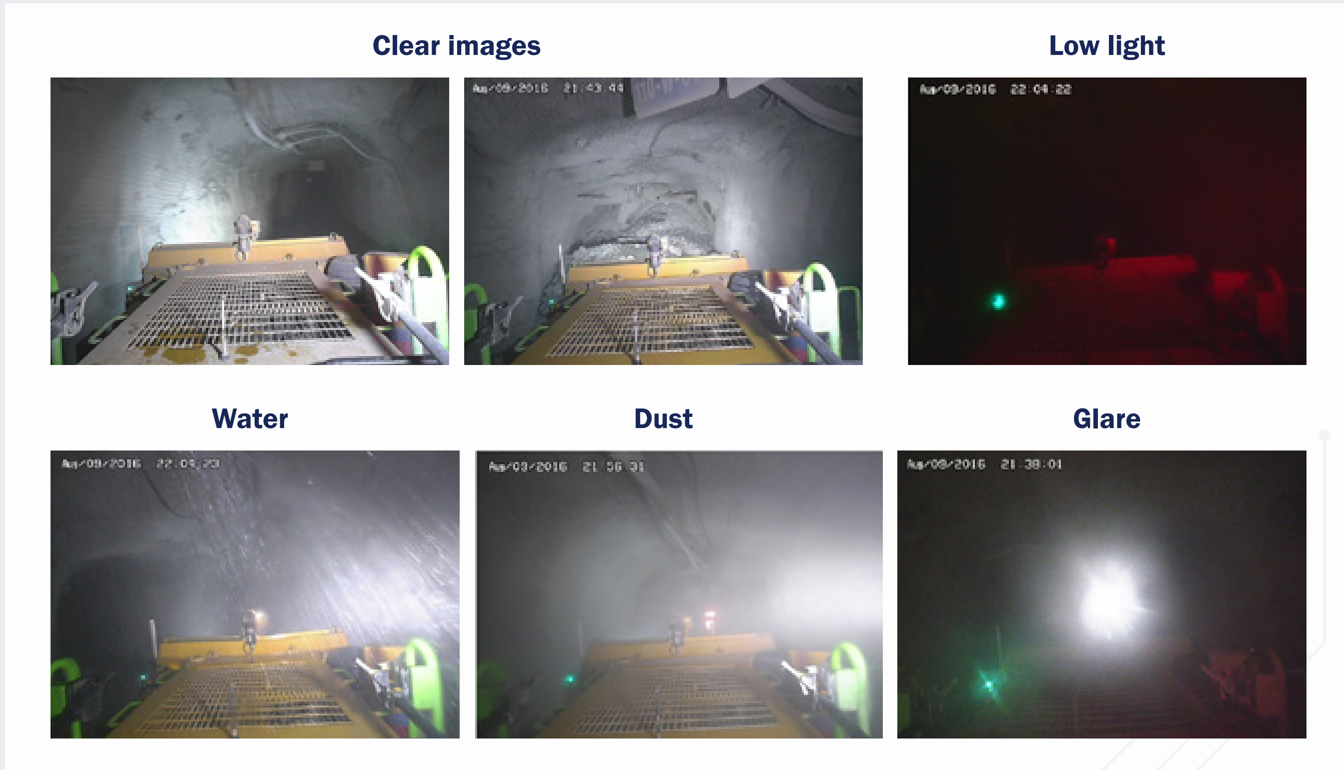

Section 4 **National Benefit**

Mapping Australia's **new robot economy**

The Centre partnered with industry, researchers and government agencies to develop Australia's first Robotics Roadmap, released June 2018.

The Centre believes Australia has the opportunity to take a collaborative, multi-sector approach to education, funding and legislation to build a leading role in developing robotic technologies and the tech sector more generally.

Robotic technologies are central to the fourth industrial revolution (Industry 4.0), where the physical and digital worlds collide.

Our continued standard of living depends on us improving productivity 2.5 per cent every year. This cannot come from labour productivity alone, and will rely on automation, which is forecast to boost Australia's productivity and national income by up to $2.2 trillion by 2030.

Australia's robotics industry has a significant store of talent in research communities, small businesses and large corporations. Combined with the right investment, this ecosystem can build, and feed, an innovation pipeline that will realise new robotic vision products, services and businesses and the skills needed to service them.

Over the next decade, a range of new technologies, importantly including robotic vision, will see robots that are more tactile, more capable of interacting with their physical environments, more closely working with humans and more self-sufficient.

A vibrant robot economy in Australia will help maintain living standards; safeguard the environment; provide services to remote communities; reduce healthcare costs; create safer and more fulfilling jobs; encourage investment; and reshore jobs.

ROADMAP RECOMMENDATIONS

INDUSTRY

Ensure Australia's ongoing prosperity by stimulating formation of new hi-tech firms, encouraging global tech giants to invest in Australia, and reskilling Australian workers.

EDUCATION

Equip all Australians with Industry 4.0 relevant skills.

GOVERNMENT

Lead the region in catalysing robotics activity by setting ethical, legal and standards frameworks and adopting robotics in government services.

R&D

Develop clusters of robotics activity, encourage VC investment, develop aspirational research challenges and encourage multidisciplinary problem-solving.

CULTURE

Support an entrepreneurial culture around Australia's niche robotics capability and harness the nation's imagination through aspirational challenges that solve pressing needs for our nation.

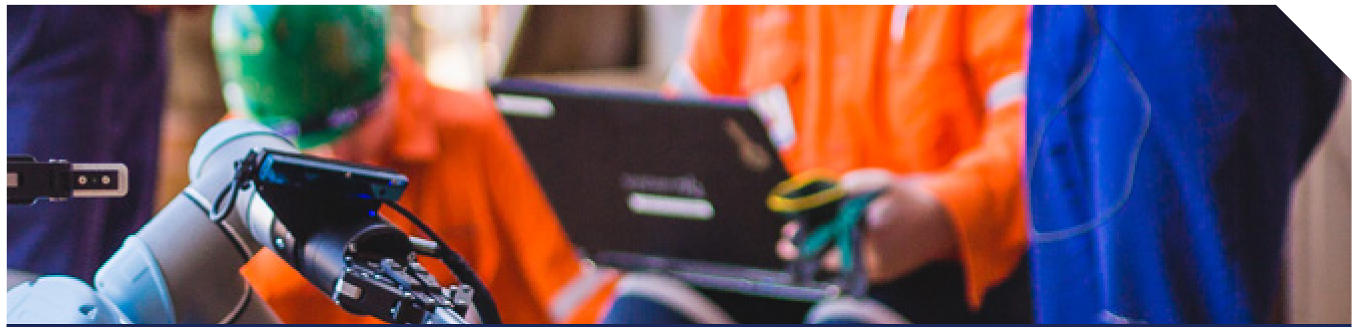

5 KEY PRINCIPLES UNDERPIN A ROBOTICS ROADMAP FOR AUSTRALIA

JOBS MATTER

A robotics industry will enhance economic competitiveness to create meaningful jobs and, with the right policy settings, help adapt existing ones

SAFETY IS IMPERATIVE

Robotics reduces the risk of workers being placed in hazardous situations

CERTAINTY COUNTS

Investment decisions need clarity of understanding. The roadmap offers the means to find that clarity as well as demonstrating the current and likely future state of robotics in Australia

TIME MATTERS

The right use of robotics eliminates workplace routine, improves efficiency and allows workers to dedicate time to interesting and more fulfilling tasks

REMOTE COMMUNITIES NEED TO BE SERVED

Automation could help provide better and more consistent services to remote areas difficult to serve

Why Australia's first Robotics Roadmap is important?

Centre Chief Operating Officer Sue Keay was the key driver behind the development of *A Robotics Roadmap for Australia*. She reflects on its importance as a 'living document' and Centre legacy.

The original idea for a roadmap was to help identify companies operating in the space of robotics and computer vision here in Australia and to get a better sense of their operations and value to the Australian economy.

"Thanks to the roadmap we now know that Australia has more than 1,100 companies operating in the robotics space, employing more than 50,000 people and generating more than $12 billion in revenue to the economy."

We also know that the development of national robotics roadmaps and strategies has had a significant impact on the level of investment in both the robotics industry and robotics R&D in the countries concerned.

It is hard to encourage investment in an industry that is poorly defined. So the roadmap journey commenced with an exercise looking to define Australia's capability in this space and the results were impressive.

Thanks to the roadmap we now know that Australia has more than 1,100 companies operating in the robotics space, employing more than 50,000 people and generating more than $12 billion in revenue to the economy.

Those numbers are important because, until recently, Australia's history of investment in new robotic start-ups has been very low compared to our peer nations.

Until we can encourage more investment there is little incentive for the talent and technologies we are currently developing in Australia to remain in Australia.

The roadmap was also a great opportunity to define Australia's strengths in robotics and some of the many world firsts we have achieved.

As a nation we create great technology to overcome the barriers of being a large country with a small population. Field robotics and remote sensing is one of our strengths, as well as the application of robotics in the resources and agricultural industries, in logistics and infrastructure inspection.

Importantly the roadmap also considered the impact that robotics would have in areas of the Australian economy that we often don't associate with robotics. For example, the Services sector, which employs more than 80 per cent of Australia's workforce.

The roadmap considers what the future might look like across all Australian industries if we are to develop suitable robotics technologies to transform them.

Australia can continue to build niche solutions to meet the unique challenges our country faces while also developing a competitive robotics industry to create new export markets.

It's a great opportunity and I look forward to seeing it become a reality.

Living Legacy: to ensure the Robotics Roadmap remains a 'living document', the Centre is launching a 'Robotics in Australia' website platform to celebrate ongoing advances in the development of a vibrant robot economy: **www.roboticsinaustralia.com.au**

Science and Research Priorities

In 2015, the Australian Government defined a set of nine national Science and Research Priorities that we have mapped our Centre's research against. The priorities that overlap with the Centre's research are highlighted in bold text.

Food

- Demand, supply chains and the identification of country specific preferences for food Australia can produce.
- Social, economic and other barriers to achieving access to healthy Australian foods.
- Enhanced food production through:
  - **novel technologies, such as sensors, robotics, and real-time data systems and traceability, all integrated into the full production chain;**
  - better management and use of waste and water; increased food quality, safety, stability and shelf life;
  - protection of food sources through enhanced biosecurity;
  - genetic composition of food sources appropriate for present and emerging Australian condition.

Soil and water

- **New and integrated national observing systems, technologies and modelling frameworks.**
- Understanding sustainable limits for productive use of soil, water, terrestrial and marine ecosystem.
- Restoration and remediation of soil, fresh and potable water, urban catchments and marine systems.

Transport

- **Low emission fuels and technologies for domestic and global markets.**
- **Urban design, autonomous vehicles, electrified transport, sensor technologies, real time data and spatial analysis.**
- Effective pricing, operation and resource allocation.

Cyber security

- Highly secure and resilient communications and data acquisition, storage, retention and analysis.
- Secure, trustworthy and fault-tolerant technologies.
- New technologies and approaches to support the nation's cyber security.
- Understanding the scale of the cyber security challenge for Australia.

Energy

- Low emission energy production from fossil fuels and other sources.
- New clean energy sources and storage technologies.
- Australian electricity grids that can readily integrate and more efficiently transmit energy.

Resources

- A fundamental understanding of the physical state of the Australian crust, its resource endowment and recovery.
- Knowledge of environmental issues associated with resource extraction.
- Lowering the risk to sedimentary basins and marine environments due to resource extraction.
- **Technologies to optimise yield through effective and efficient resource extraction, processing and waste management.**

Advanced manufacturing

- Knowledge of Australia's comparative advantages, constraints and capacity to meet demand.
- **Cross-cutting technologies that will reduce risk, scale up and add value to Australian manufactured products.**
- Specialised, high value-add areas such as high-performance materials, composites, alloys and polymers.

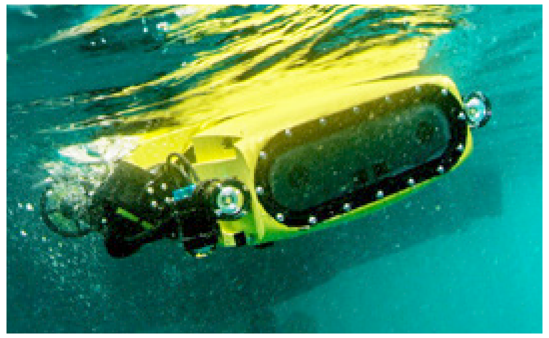

Environmental change

- **Predicting and measuring the impact of environmental changes caused by climate and local factors.**
- **Resilient urban, rural and regional infrastructure.**
- **Options for responding and adapting to the impacts of environmental change on biological systems, urban and rural communities and industry.**

Health

- **Better models of healthcare and services that improve outcomes, reduce disparities for disadvantaged and vulnerable groups, increase efficiency and provide greater value for a given expenditure.**
- Improved prediction, identification, tracking, prevention and management of emerging local and regional health threats.
- Better health outcomes for Indigenous people, with strategies for both urban and regional communities.
- **Effective technologies for individuals to manage their own healthcare.**

Section 5 **Engagement**

PHOTO by Anthony Weate

Our focus on end users

The Centre is focused on generating internationally impactful science and new technologies that will transform vital industries and provide solutions to some of the hard economic and social challenges facing Australia and the world.

We are tackling the kind of challenges that threaten our quality of life, including:

/ labour shortages and low productivity growth in key industries;
/ diminishing competitiveness;
/ increasing occupational health and safety compliance costs;
/ ageing infrastructure;
/ rising healthcare costs; and
/ the balance between environmental sustainability and growing demands for minerals, energy and food.

It is not enough for the Centre to do fantastic, internationally impactful science if it fails to translate research into tangible benefits to end users.

Our role is not to compete with existing businesses in the Australian marketplace, but to identify companies we can partner with to build and deliver technological solutions that will serve across a range of industries.

Our key focus is on developing robotic vision technologies in those areas where their application is identified as having the highest potential economic benefit to Australia, as outlined in Australia's first Robotics Roadmap.

PHOTO by Anthony Weate

Market size and maturity of sectors of the Australian economy that will be impacted by the adoption of new robotic, automation and computer vision technologies.

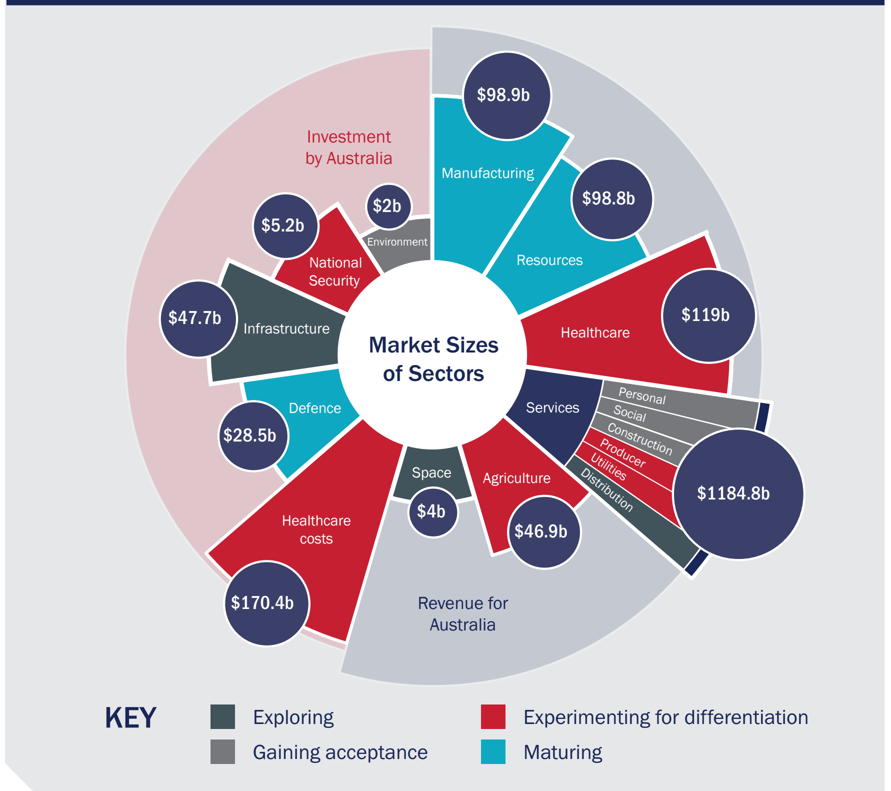

Graph from 'A Robotics Roadmap for Australia'

How we **engage**

Robotic vision has the potential to impact all industries. As a result, the Centre's End User Engagement Strategy must incorporate a diversity of approaches.

We achieve this through:

- local demonstrations where existing and potential end users can see our technologies in operation;
- sector-specific industry workshops, engaging directly with businesses to understand what challenges they are facing and which of these challenges can be addressed through robotic vision solutions;
- traditional media engagement, social media and public events;
- one-on-one consultations where our members host meetings with representatives from various industry segments related to our targeted impact areas;
- bringing high profile international conferences to Australia;
- VIP tours (the Centre regularly invites high-profile visitors to our various nodes to showcase our evolving work); and
- publication of research papers and articles in respected trade and peer-reviewed journals.

We track our engagement activities using a CRM system which records contacts, meetings and other relevant information.

As demonstrated in Australia's first Robotics Roadmap, we lobby governments to work with our national robotic vision community to build a strong robotic vision industry in Australia.

CENTRE STAKEHOLDERS

Our end users provide us with the real-world challenges we must address. Their businesses and the Australian economy will benefit by implementing the technologies we develop. Our success will come from engaging with a wide range of stakeholders, from our funding agencies and business partners, through to hobbyists and community groups that inspire the next generation of end users.

CENTRE STAKEHOLDER MAP

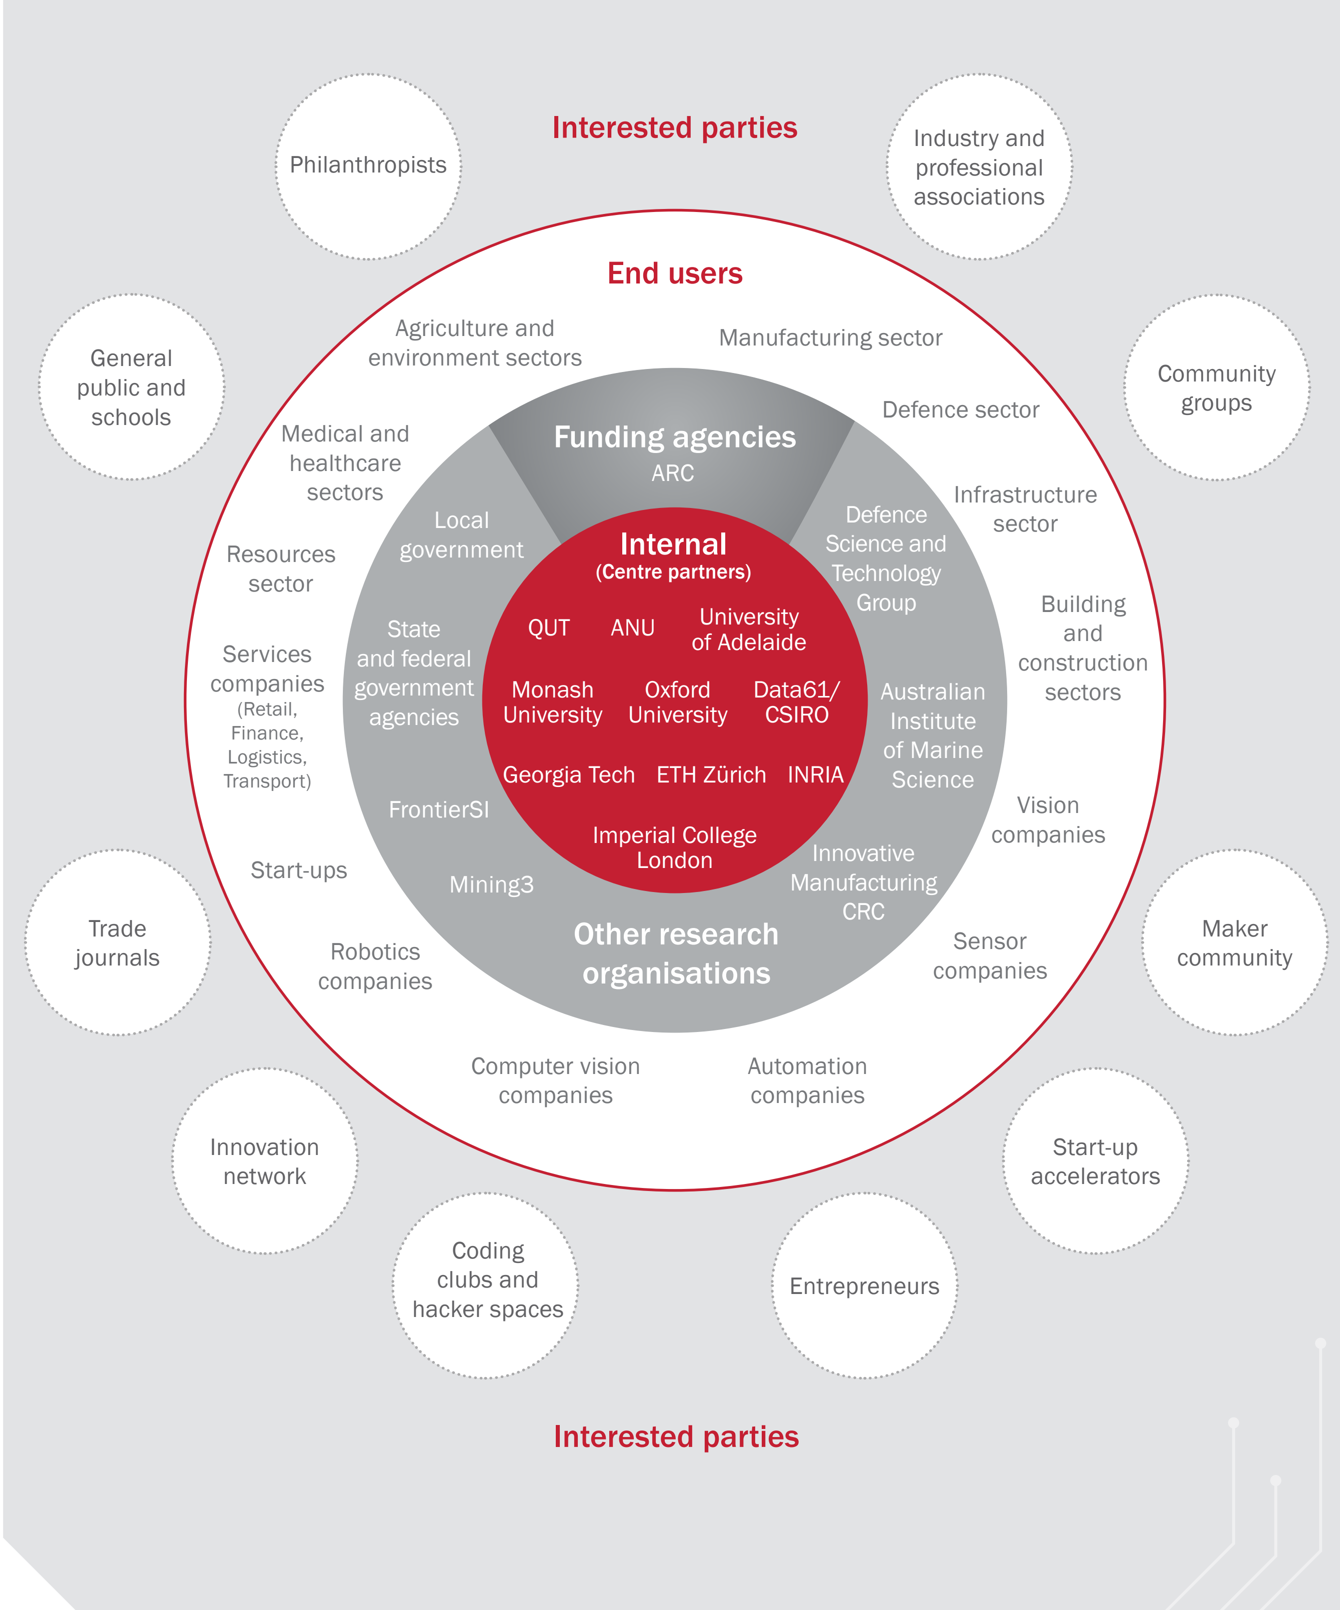

COMMUNICATION AND MEDIA ENGAGEMENT

As part of the Centre's strategic plan to increase visibility and public/stakeholder impact, a full-time Communications Specialist commenced in July 2018.

TRADITIONAL MEDIA

2018 has been something of a watershed year for national and international media coverage about the Centre's work. A year to celebrate, punctuated by the release of **Australia's first Robotics Roadmap**; Australia's first staging of the world's premier robotics research forum, **International Conference on Robotics and Automation (ICRA)**; and **tangible, real-world applications of robotic vision** once only dreamed possible. Not least being, deployment of the world's first robotic vision-empowered reef protector, RangerBot, helping a coral seeding program on the Great Barrier Reef.

The Centre and our researchers were the focus of more than 1,500 stories across print, online and broadcast media (radio and TV). This is valued at more than $12.2 million (according to the AVE measurement based on equivalent advertising value). Media coverage reached more than 1.3 billion people globally across Australia and more than 40 countries.

SOCIAL MEDIA

The Centre maintains a range of social media channels – Facebook, Instagram, Twitter and YouTube – to inform, inspire and better understand the widest possible audience. In the second half of 2018, we focused on increasing meaningful social media engagement via Facebook and Instagram, resulting in a 25 per cent increase in followers. Of note, the Centre's release of Australia's first Robotics Roadmap drew seven million social media impressions with positive sentiment.

More than just a pinch of Pepper!

SoftBank's social robot, Pepper, again captured hearts and minds as our most engaging robot.

Stories about the work of the Centre's social robotics team appeared across broadcast, print and online media, including Fairfax and News Corp titles; Xinhua News Agency; national news programs (9 News, 7 News, ABC, 10 News, ABC); Totally Wild and Scope (Network Ten); BrainBuzz (Nine Network); and Juiced TV (made by the kids in hospital for kids in hospital). Pepper even starred in a tech-related fashion shoot for *Lita Magazine* with the Centre as a backdrop.

In May, the Centre's social robotics team were kept busy running continuous live demonstrations with Pepper over four days at ICRA 2018 in Brisbane – an event which attracted over 3,000 delegates and significant media coverage. See page 98 for full details on this conference.

During National Science Week, a heart-warming new friendship was born at the Centre's QUT headquarters. Ella Wrigley, 7, the first Queenslander diagnosed with a life-threatening metabolic condition – propionic acidemia – became 'besties' with Pepper. The meeting was organised following a chance encounter at ABC Radio Brisbane studios between Ella's mum, Amanda, and Centre Chief Operating Officer Sue Keay. The inspiration? A deep *Star Wars* connection. Ella needed a liver transplant for her condition. She named her new liver R2-D2, after the fictional droid that also inspired Dr Keay's love of (and career with) robots. Building on this connection, Ella visited the Centre with her Mum and older sister, Ava, 11. Their visit drew national media coverage.

A key focus of National Science Week is to encourage more children – in particular, girls – to consider STEM as a career. This message was reinforced by Queensland's Science Minister, Leeanne Enoch, who filmed an educational video at the Centre (also with Pepper as her guide).


PHOTO by Heidi Huttunen/LITA Magazine

Pitch Perfect

His Royal Highness The Duke of York witnessed first-hand the Centre's life-enhancing research, creating robots able to see for the good of all people and our planet. He visited QUT on 29 November 2018.

The Royal Family's STEM advocate was in Australia for Pitch@Palace, a unique global initiative he founded in 2014 to give a platform to entrepreneurs.

Photos from The Duke's meeting with Centre researchers were 'tweeted' on the Royal Family's official Twitter account. The Duke met Centre Director Peter Corke, Chief Operating Officer Sue Keay and researchers including PhD Researcher Andrew Razjigaev and Research Fellow Feras Dayoub. He tried his hand at operating Snakebot (destined to revolutionise keyhole surgery) and was briefed on the work of RangerBot on the Great Barrier Reef.

PHOTO by Anthony Weate

WEBSITE

roboticvision.org is at the core of our online presence. Our website provides comprehensive information about the Centre including our purpose; detailed profiles of our research programs; people; field testing; news and videos; access to our publications; opportunities for industry collaboration; and workshops and events. It is also a key point of contact for our researchers around the world, enabling access, through a secure portal, to Centre operational information, collaborative research forums, and policies and procedures including the Centre's treatment and protection of intellectual property.

Website traffic doubled in 2018, with 32,405 new users attracted to the site (up from 16,042 in 2017). The site had 109,587 page views during the year (up from 55,495 in 2017) with key interest in our researcher profiles, job opportunities, the Robotics Roadmap and Robotic Vision Summer School (RVSS).

DISTINGUISHED VISITORS AND EXCHANGE OPPORTUNITIES

The Centre hosted 50 international visitors from 41 organisations and 18 countries in 2018. They included Professor Riaan Stopforth (University of KwaZulu-Natal); Assistant Professor Jens Kober (TU Delft); Professor El Mustapha Mouaddib (University of Picardie Jules Verne); Professor David Lane (Edinburgh Robotics); Associate Professor Veronique Auberge (LIG-Grenoble Computer Science Research Lab); and Raia Hadsell (Google DeepMind).

International guest speakers at our Robotic Vision Summer School included: Professor Margarita Chli (ETH Zürich); Professor Vincent Lepetit (University of Bordeaux); Associate Professor Yarin Gal (University of Oxford); and Associate Professor Andrea Cherubini (Universite de Montpellier).

Fifty students (16 international, 34 domestic) including six partner organisation students from INRIA, Oxford University and Imperial College London attended RVSS.

Nine of these international students visited our QUT, Adelaide and ANU Centre nodes following RVSS.

Steven Parkison, Lu Gan and Arash Ushani (University of Michigan); Axel Lopez Gandia (IRISA); and Arnab Ghosh (Oxford University) spent a week at QUT. Agniva Sengupta (INRIA) visited ANU. Bibin Xu and Zoe Landgraf (Imperial College); and Viveka Kulharia (Oxford University) visited the University of Adelaide.

Other centre-hosted international exchanges included: PhD Researchers Horia Porav (Oxford University); Weihao Li (Heidelberg University); Caner Hazirbas (TU Munich); Anh Nguyen (Italian Institute of Technology); and Adriene Spaenlehauer (Université de Technologie de Compiègne). Visiting Masters Researchers were: Francesca Cini (Biorobotics Institute at Scuola Superiore Sant'Anna); Lisa Teunissen (Eindhoven University of Technology); Attila Lengyel (TU Delft); Thomas Schopping; Timo Korthals's (Bielefeld University); and Masters graduate Antoine Dubois (University of Liège, Belgium).

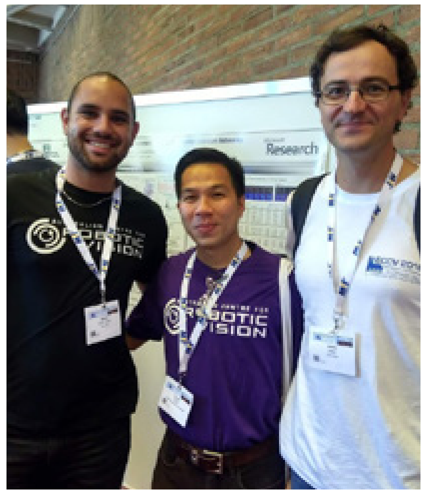

PUBLICATIONS AND PRESENTATIONS

Our researchers published a total of 180 papers, including 70 conference papers and 40 journal articles. See page 134 for a full list of publications from 2018.

Key publications included:

- Research Fellow Nicole Robinson's work published in Science Robotics: *"Measures of Incentives and confidence in using a social robot".*
- Research Fellow Dylan Campbell received first prize in the IEEE Australia Council Postgraduate Student Paper Competition for the paper *"Globally-Optimal Inlier Set Maximisation for Camera Pose and Correspondence Estimation".*
- Chief Investigator Dr Niko Sünderhauf and colleagues had their proposal for a special issue on Deep Learning for Robotic Vision accepted in the International Journal of Computer Vision (IJCV). This is a follow-up from the workshop on "Deep Learning for Robotic Vision" held at the Computer Vision and Pattern Recognition (CVPR) conference in Honolulu, Hawaii. The guest editors are: Anelia Angelova, Gustavo Carneiro, Niko Sünderhauf and Juxi Leitner.
- PhD Researcher Lachlan Nicholson and Chief Investigators Dr Niko Sünderhauf and Professor Michael Milford won a best workshop paper award at the ICRA 2018 workshop 'Representing a Complex World: Perception, Inference, and Learning for Joint Semantic, Geometric, and Physical Understanding' for their paper *"QuadricSLAM: Constrained Dual Quadrics from Object Detections as Landmarks in Semantic SLAM".*
- Centre researchers were finalists in the 2018 Amazon Robotics Best Paper Awards in Manipulation for their paper *"Cartman: The low-cost Cartesian manipulator that won the Amazon Robotics Challenge".*
- Best Science Paper Award at the British Machine Vision Conference (BMCV) for the paper *"Non-smooth M-estimator for Maximum Consensus Estimation"* by Huu Le, Anders Eriksson, Michael Milford, Thanh-Toan Do, Tat-Jun Chin, and David Suter.
- Chief Investigator Professor Tom Drummond and colleagues won the ISMAR Impact Paper Award 2018 for their paper *"Pose tracking from natural features on mobile phones".* The IEEE International Symposium on Mixed and Augmented Reality (ISMAR) was held in Germany.

We led over 200 government, industry and business community meetings and presentations, bringing the complex world of robotic vision to a wide audience.

CONFERENCES AND WORKSHOPS

International conferences play a key role in growing our networks both at home and abroad. The Centre supports our researchers' attendance at these events to share our findings, learn about important work other researchers are doing, and to recruit new researchers and students. These activities lift our international reputation as a research centre. This helps us attract world-leading researchers to Australia, further enhancing our international linkages and collaborative partnerships. Centre researchers presented papers at over 10 conferences in eight countries, including here in Australia.

The largest conference in 2018 was the International Conference on Robotics and Automation (ICRA), staged in Brisbane in May. ICRA is the IEEE Robotics and Automation Society's flagship conference and is a premier international forum for robotics researchers to present their work. ICRA 2018 committee members from the Centre included: Alex Zelinsky, General Chair; Peter Corke, Program Chair; Rob Mahony, Tutorial Chair; Henrik Christensen, Government Forum Chair; Luis Mejias, Publication Chair; Jon Roberts, Competitions Chair; and Thierry Peynot, Local Arrangement Chair.

The following outlines the breadth of the Centre's involvement at ICRA 2018:

- Partner Investigator Professor Francois Chaumette (INRIA) and Chief Investigator Professor Robert Mahony (ANU) presented lectures as part of the "Tutorial On Vision-Based Robot Control" alongside Professor Pieter Abbeel (UC Berkeley) and Fabien Spindler (INRIA).
- Chief Operating Officer Dr Sue Keay co-organised the Social Robotics Forum with Centre Associate Investigator Professor Elizabeth Croft (Monash University) and Distinguished Professor Mary-Anne Williams (University of Technology, Sydney). The forum had 12 speakers; eight female and four male speakers. They included Centre Research Fellow Dr Nicole Robinson; Assistant Professor Angelica Lim (Simon Fraser University); Dr Amit Kumar Pandey (SoftBank Robotics Europe); and Assistant Professor Maya Cakmak (University of Washington).
- Partner Investigator Professor Paul Newman and Centre Director Distinguished Professor Peter Corke spoke at the PhD Forum which provided an opportunity for a group of PhD researchers to discuss and explore their research interests and career objectives with a panel of established researchers in robotics. They worked closely with forum participants in exploring trends in robotics and provided insights into their own career trajectories.
- Research Fellow Dr Feras Dayoub and Chief Investigator Dr Niko Sünderhauf co-organised the workshop "Long-term Autonomy and Deployment of Intelligent Robots in the Real-world". The workshop was a collaboration between QUT, the Oxford Robotics Institute, Czech Technical University, and Carnegie Mellon University.
- Research Fellow Dr Juxi Leitner co-organised the "Advances in Robotic Warehouse Automation: Solutions, Lessons Learned and the Future after the Amazon Robotics Challenge" workshop with Joey Durham (Amazon Robotics), Albert Causo (Nanyang Technological University) and Alberto Rodriguez (MIT).
- Centre Research Affiliates Dr Chris McCool and Dr Chris Lehnert organised the workshop "Robotic Vision and Action in Agriculture: the future of agri-food systems and its deployment to the real-world".
- Research Fellow Juxi Leitner and Chief Investigator Dr Niko Sünderhauf ran the "Tidy Up My Room" Challenge in which teams deployed their robots into a room set-up (think IKEA showroom) to perform specific tasks.
- Research Engineer Steve Martin organised the public event "Robowars" involving robots big and small, built from metal, power tools and remote-controlled toy parts to battle in the ultimate test of human and machine.
- Chief Investigator Professor Matt Dunbabin gave a keynote talk on "Robot vision to action in the wild – challenges and opportunities for vision in natural environments".
- Partner Investigator Professor Paul Newman's keynote talk "The Road to Anywhere Autonomy" addressed the remaining challenges facing extensive deployment of driverless vehicles and progress aided by machine learning, optimisation and data fusion.

Centre researchers participated in other key conferences in 2018:

- AAAI Conference on Artificial Intelligence in New Orleans, United States;
- Computer Vision and Pattern Recognition (CVPR) conference in Salt Lake City, United States;
- Robotics: Science & Systems (RSS) in Pittsburgh, United States;
- British Machine Vision Conference (BMCV) in Northumbria, United Kingdom;
- European Conference on Computer Vision (ECCV) in Munich, Germany;
- International Conference on Intelligent Robots and Systems (IROS) in Madrid Spain;
- Neural Information Processing Systems (NIPS) in Montreal, Canada; and
- Australasian Conference on Robotics and Automation (ACRA) in New Zealand.

Research Fellow Dr Valerio Ortenzi organised the workshop *"Human-Robot Cooperation and Collaboration in Manipulation: Advancements and Challenges"* during the International Conference on Intelligent Robots and Systems (IROS 2018) in Madrid, Spain. The workshop was a joint collaboration between QUT (Valerio Ortenzi and Peter Corke) and the Biorobotics Institute at Scuola Superiore Sant'Anna in Italy (Marco Controzzi and Francesca Cini).

At IROS 2018, Centre Director Peter Corke presented a tutorial on "Screw Theory for Robotics" and its application to commercial manipulators along with speakers from Northwestern University, London South Bank University, University of Naples Federico II, and the University Carlos III Madrid.

The Australasian Conference on Computer Vision (ACCV) was held in Perth in December 2018. Deputy Director Ian Reid was a General Chair of the committee, and Chief Investigator Hongdong Li was a program chair. Chief Investigator Richard Hartley gave a keynote address on the usefulness of deep learning vs math in computer vision. Associate Investigator Qi Wu and Chief Investigator Anton van den Hengel were organisers of the "Combining Vision and Language" workshop.

Standing ovation: ICRA 2018 takes out 'best conference' award

Australia's first staging of the world's premier robotics research forum, **International Conference on Robotics and Automation (ICRA)**, in Brisbane, was the result of four years' planning by the Centre's Advisory Committee Chair, Alex Zelinsky, and Centre Director Peter Corke.

It marked the first time the premier international forum touched down in the Southern Hemisphere in its 35-year history.

The Centre was a proud supporter of the conference and our researchers were involved in forums, workshops, tutorials and competitions, and as volunteers.

Centre Director Peter Corke was Program Chair of the conference that attracted more than 3,000 delegates from 60 countries; showcased the presentation of 1,052 papers; and created $10.5 million in economic benefit to Queensland.

"I attended my first ICRA in 1988 and most of the subsequent ones," said Distinguished Professor Corke.

"It was an honour to host so many international researchers, many of them long-term friends and colleagues, and I'm really proud of the work done by a wonderful team of volunteers.

"We showed Australia in a very positive light, the local technical work, our ability to organise a massive conference and a fabulous venue and host city."

DID YOU KNOW?

ICRA 2018 was awarded the prestigious Australian Events Award for Best Congress or Conference and Best Event (Queensland Category) in the Meetings & Events Australia Industry Awards. The accolades acknowledged the work of ICMS Australasia in its management of the conference staged at Brisbane Convention and Exhibition Centre featuring four robotic competitions and seven robotic forums covering industry, education, government, robotic art installations, robot ethics, social robotics and sustainable humanitarian technologies. In addition, a public day attracted over 3,000 to the Brisbane Southbank precinct for ICRA X.

COMMERCIALISATION

A core objective of the Centre is to develop intellectual property that we can license or sell in commercial markets. We have established guidelines and policies on how to manage our intellectual property, along with approaches for open-sourcing some of our technologies. The Centre maintains a register of all intellectual property that we develop during the course of our activities, along with any dependencies on background intellectual property. Individual researchers assign the intellectual property they have developed to their host university, as collaborating partners in the Centre Agreement. The Agreement sets out the terms of ownership among project parties, who consult with the Centre Executive on protecting and commercialising intellectual property as it is developed.

In addition to commercialising our intellectual property, we can achieve great impact with our research through our open source policy. In cases where intellectual property we create has low potential for commercialisation, the Centre recommends the 'BSD-3-Clause License' created by the Open Source Initiative (opensource.org).

The Centre believes open source software is increasingly important to broadening the global impact of research results, particularly in robotics and computer vision. As appropriate, we disseminate our research knowledge as open-source code and web-based demonstrations.

One of the principles we adopt in the Centre is that of entrepreneurship. We encourage and support our researchers to attend entrepreneurship master classes, business workshops and start-up forums. We want them to keep their minds alert and open to alternative commercialisation pathways that are constantly evolving in our 21st Century marketplaces.

In 2018, a Centre trio – Research Fellows Jürgen 'Juxi' Leitner and Nicole Robinson, and Master of Philosophy Researcher Norton Kelly-Boxall – ventured into the world of start-ups. They formed LYRO Robotics (www.lyro.io) as a separate entity to commercialise the next generation of smart robotic systems. Of note, they were inspired by the experience of a Centre team winning the 2017 Amazon Robotics Challenge in Japan with the only custom-built Cartesian robot in the global competition (see page 34 for full story). Two separate start-ups were launched by Centre researchers in 2017 (AlphaOne.ai and Nuarda), as detailed in our last annual report.

Open Innovation

Centre Director Peter Corke and Chief Investigator Jonathan Roberts form part of a team that established Design Robotics in 2018. This collaborative innovation project aims to reduce the integration time between real-world robotics design and custom manufacturing.

Design Robotics is a collaboration between UAP (Urban Art Projects), a world leader in public art and architectural design solutions; QUT (School of Design); Innovative Manufacturing Cooperative Research Centre; and RMIT University.

"This is an exciting venture because it marries research and commercialisation," said Professor Jonathan Roberts.

The Design Robotics Collaboration will certainly enhance UAP's ability to manufacture high-value products in less time and at reduced cost.

"The technology, however, will also greatly benefit other Australian manufacturers across industries where mass-customised products are in high demand, such as medical devices, in construction and aerospace."

The project includes the Design Robotics Open Innovation Network. This serves as an interface to wider industry, both to disseminate project findings and to gain feedback on industry needs.

"All too often meaningful and innovative collaboration is hindered by a closed-model of knowledge production, including intellectual property and other restrictions around knowledge sharing," said Distinguished Professor Peter Corke.

"The Design Robotics project embraces openness in research and tangible impact and value creation for end users – both industry stakeholders and, ultimately, the wider public."

For more visit **www.designrobotics.net**

International and National Links & Networks

Building research networks at home, and overseas, is an important part of sharing and expanding our knowledge in the computer vision and robotics sectors, together with growing our Centre's profile as a leader in the field of robotic vision.

We actively encourage our Chief Investigators, Research Fellows and PhD Researchers to visit each other's laboratories, including those of our international partners. The Centre hosts visits in return. This allows us to work more effectively together and build a truly international community around robotic vision.

PhD Researcher Fahimeh Rezazadegan spent six months at the University of Maryland Institute for Advanced Computer Studies (UMIACS) hosted by Professor Larry Davis.

"I met researchers in both the computer vision and robotics fields, developed my research topic (resulting in a paper in an A+ AI conference), showcased our Centre's research, and forged connections for future collaborations."

Chief Investigator Steve Gould spent a year at Amazon, Seattle

"Two of the key benefits of being an academic researcher are that you get to work on important and interesting problems and that you get to travel to new places and meet interesting people. Make sure you take advantage of both."

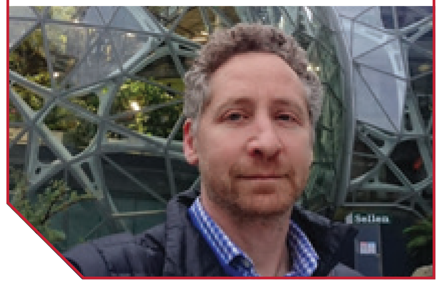

Centre Director Peter Corke spent four months at MathWorks in Boston (see page 102)

"MathWorks' MATLAB language is the 12th most popular programming language. I've used it since the company was founded and have written a number of popular software tools based on it. It was a real treat to spend time inside 'the mothership' where the software is made."

Chief Investigator Hongdong Li spent seven months on sabbatical at Carnegie Mellon University

"I worked with some of the world's best minds in computer vision and robotics on new research problems that I had never encountered before."

PhD Researcher Cedric Scheerlinck is spending 12 months at University of Zürich

"I love the feeling when I tell someone I'm from Australia and they automatically think I'm so cool! I think it's important to build connections with people around the world because it gives both parties insight into the culture of the other, and facilitates sharing of ideas."

PhD Researcher Doug Morrison is spending four months at Amazon Robotics in Berlin, Germany

"I have really enjoyed working alongside a tight-knit team of researchers and engineers on a wide variety of topics, which has made this a fantastic learning experience, as well as applying my own skills to real-world robotics problems."

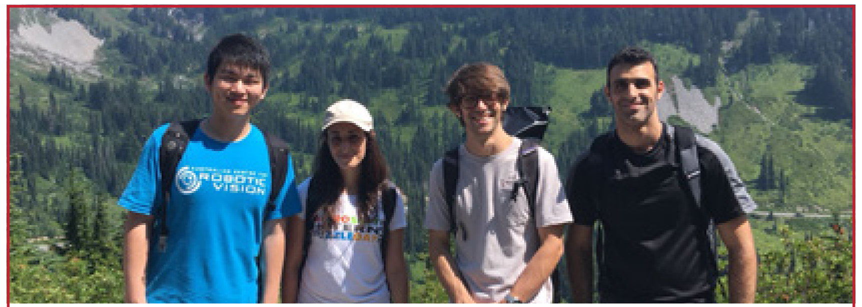

PhD Researchers Gil Avraham (pictured far right) and Huangying Zhang (pictured far left) both completed a three-month research internship at Microsoft HoloLens in Washington State

"During my internship, I worked with the spatial mapping team at Microsoft HoloLens on a project related to deep learning and 3D reconstruction, which is closely aligned to my PhD topic. It helped me understand the working styles in academic and industry environments, and was a great opportunity to meet and learn from researchers and experts outside the Centre." Huangying Zhang

"My internship was a very valuable experience, working with a diverse team of researchers and engineers with an incredible set of skills, in a cutting-edge industrial setting." Gil Avraham

PhD Researcher Timo Stoffregen spent seven months at the Robotics and Perception Group of Davide Scaramuzza, resulting in two Computer Vision and Pattern Recognition (CVPR) conference submissions and a planned collaborative journal submission

"The experience was invaluable for me and resulted in several collaborations and research outputs. I would strongly recommend doing something like this to everyone doing their PhD at the Centre."

Worldly-wise!

Creating a vibrant international robotic vision community is a key Centre goal. Leading by example, ***Centre Director Peter Corke*** *spent four months on sabbatical leave in Boston (June-September 2018). There he joined MathWorks, the world's leading developer of mathematical computing software used by engineers and scientists. He reflects on the importance of building research networks at home and overseas.*

Boston is a technology hothouse. It is home to 35 colleges and universities and boasts a large and diverse population of technology companies including bio, software, AI and robotics.

I spent three days per week at MathWorks, a software company that makes mathematical and software tools for industry, research and education. Their MATLAB language is the 12th most popular programming language.

They see automation as a key megatrend and are actively developing tools for computer vision, robotics, sensor fusion, deep learning and system integration tools.

MathWorks is a 35-year-old, privately-owned company with an annual revenue of nearly US$1 billion.

"It was a real treat to spend time inside 'the mothership' where the software is made."

I've used MATLAB since the company was founded and have written a number of popular software tools based on it. It was a real treat to spend time inside 'the mothership' where the software is made.

Being there gave me insights into the complexities of creating and maintaining an ever-growing piece of software.

The company takes feedback from users very seriously and once word got out there was a user inside the mothership, I received many invitations to preview and critique upcoming functionality through formal usability studies and informal white-boarding sessions. I also learnt a lot in this process and that will be folded into current and future courses at QUT.

MathWorks has a great internal culture and is still growing – while I was there a new campus was opened to accommodate the growing headcount.

While on sabbatical I also worked on a new book project and spent time on some teaching-related software projects.

I took the opportunity to visit Boston Dynamics, MassRobotics (which is incubating robotic companies), Northeastern University and MIT and did a side trip to our Centre partner Georgia Tech.

Section 6 **Our People**

2018 Honours and Accolades

- Associate Investigator Professor Elizabeth Croft was awarded the RA McLachlan Award by the Engineers and Geoscientists of British Columbia. This is a peak career award for Professional Engineering in the Province of British Columbia, Canada.
- The International Conference on Robotics and Automation (ICRA) won the Australian Events Award for Best Congress or Conference of 2018, and Best Event (Queensland Category) in the Meetings & Events Australia Industry Awards.
- Chief Operating Officer Dr Sue Keay was recognised in the 'Power 100' list (published in *The Courier-Mail*) for putting robotics on the map in Australia.
- Research Affiliate Dr Anjali Jaiprakash won the MIT Technology Review innovator under 35 for the Asia Pacific region.
- Chief Investigator Gustavo Carneiro was promoted to Professor at the University of Adelaide.
- Chief Investigator Dr Niko Sünderhauf was promoted to senior lecturer at QUT and received a Google Faculty Research Award.
- Centre Deputy Director Professor Ian Reid was appointed as Head of the School of Computer Science at the University of Adelaide.
- PhD Researcher Cedric Scheerlinck, from ANU, was awarded a Swiss Government Excellence Scholarship. This 12-month scholarship will allow Cedric to pursue his PhD research at the University of Zürich (UZH), working with Professor Davide Scaramuzza. Cedric will be based in Zürich from September 2018 to August 2019. His PhD topic is the development of novel motion estimation algorithms using event cameras, for use in robotics.
- Chief Investigator Professor Anton van den Hengel was a finalist in the South Australian Scientist of the Year awards.
- Associated Investigator Qi Wu and his colleagues Yan Huang and Liang Wang from the National Laboratory of Pattern Recognition, China, were presented with an NVIDIA Pioneering Research Award during the Computer Vision and Pattern Recognition (CVPR) conference for their work "Learning Semantic Concepts and Order for Image and Sentence Matching". The awards went to those who've used NVIDIA's AI platform to support great work featured in papers accepted by CVPR and other leading academic conferences.
- Associate Investigator Associate Professor Nick Barnes won the CSIRO Digital and National Facilities Science Excellence Award.
- Centre PhD graduates (and now Research Fellows) Dr Hui Li and Dr Bohan Zhuang both received Deans Commendations for Doctoral Thesis Excellence at the University of Adelaide.
- Chief Investigator Professor Hongdong Li has been appointed as Area Chair for both CVPR and International Conference on Computer Vision (ICCV) in 2019.
- Research Fellow Dr Feras Dayoub became Associated Editor for the IEEE Robotics and Automation Letters, and Guest Associate Editor (Special Issue on Precision Agricultural Robotics and Autonomous Farming Technologies) in IEEE Robotics and Automation Letters.
- Centre PhD Researchers Liu Liu, Liyuan Pan and Chief Investigator Hongdong Li came third in the Google Landmark Retrieval challenge with their team, SevenSpace.

Research Training and **Professional Development**

CREATING KNOWLEDGE LEADERS

As a Centre of Excellence, we have an important remit from the Australian Research Council to provide high-quality postgraduate and postdoctoral training environments for the next generation of researchers and research leaders.

Since our inception, we have strategically chosen to devote as much of the Centre's personnel budget as possible to research and research training. This extends to development of our people as future talent for industry and academia, including Knowledge Leadership courses at all Centre nodes.

Our objectives are clear:

- We want to develop and support the next generation of robotic vision experts.
- We want to prepare our leaders for diverse career pathways.
- We want to see our graduates and alumni form new companies in Australia's fledgling robotic vision industry and give them the skills and support they need to take robotic vision to the world.

The Centre has a target of training 24 postdoctoral fellows and 80 PhD researchers over its life to 2020. We have already achieved this. The Centre has delivered a total of 16 knowledge leadership workshops to our PhD researchers and Research Fellows. Our cohort of PhD researchers has risen from six (in 2014) to 72 (in 2018). A further 19 PhD researchers have graduated from the Centre. Our postdoctoral fellows currently number 22.

Our Centre's alumni (numbering 50 at the end of the reporting period) are powerful global ambassadors for the advancement of robotic vision. Our alumni can be found working in diverse locations across Australia, the United States, South America, China, Singapore, Sweden, Switzerland, the Netherlands, Germany, the United Kingdom, Pakistan and Canada. This network spans 13 countries and 31 organisations.

The Centre strives to provide a rich training experience for our early career researchers. We offer a range of professional development opportunities for researchers and staff, including:

- travel to other Centre nodes or our overseas partner universities, to experience different research environments (see page 100);
- the opportunity to attend top international conferences and regular seminars at Centre nodes; and
- meeting and interacting with eminent visiting researchers from our national and international partners.
- In 2016, the Centre partnered with Evexia's workplace psychologists, Jo Karabitsios and David Whittingham, to launch a custom-developed Knowledge Leadership program. This started with a one-day workshop at our annual symposium, RoboVis, and has since expanded to an additional eight annual workshops (two at each node). The program focuses on personal and professional development, including modules on: Your Brand, Leading with Influence, Your Development and The Entrepreneur. As a result, researchers are equipped with better knowledge of intra-personal skills required for a successful research leadership career including social influencing and experiential problem solving. The program has the added advantage of strengthening connections across the Centre.

OUR ALUMNI AROUND THE WORLD

"Having just commenced my PhD studies with an exciting journey ahead, I found the workshop an invaluable opportunity to gear myself with communication skills and boost my self-confidence."
Mahsa Ehsanpour, University of Adelaide

"A big aspect of the workshops for many were discussions about our 'personal brand' and how we might go about explaining our values and why we wanted to do research in our particular field. We were given the opportunity to think about how we would communicate to the public the value of our research and the impact we are hoping to make. We were each able to video and review ourselves trying to explain our work to a particular target audience."
Luke Ditria, Monash University

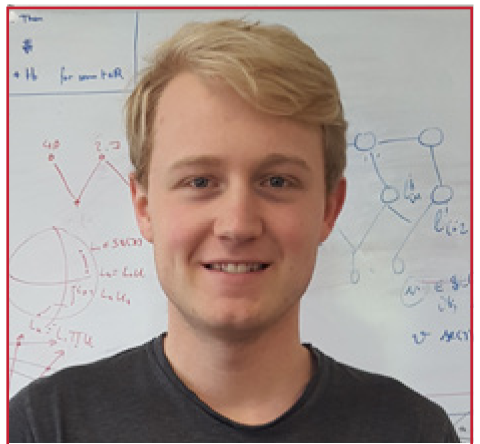

"The Knowledge Leadership workshops have been as much about developing our skills as they have been about forming connections throughout the Centre."
Pieter van Goor, Australian National University

"As a PhD student we often come across times of stress. This stress may be caused by a looming deadline, a conflict, or might be coming from the challenge of balancing your studies with your personal life. As a result there may be times when we wish to reject reality and substitute our own; times where we want to simply escape the stresses. While, escaping from our reality by taking time out is a great way to refresh; ultimately the stress-inducing problem is still there and needs to be dealt with. To help navigate through the challenges of a career in research we have consumed the tips and tricks provided by the Evexia's Knowledge Leadership workshops."
James Mount, QUT

OPERATIONAL EXCELLENCE

The Centre Operations Team hosted the 2018 Centres of Excellence (CoE) Professional Staff Development day on 29 August in Brisbane. This workshop followed the success of the inaugural event in 2017 hosted by our colleagues from the ARC Centre of Excellence for Engineered Quantum Systems.

Led by Centre Coordinator, Kate Aldridge, it involved 45 attendees, representing operational staff from five ARC CoEs and two ARC research training centres. It succeeded in its aim of further developing peer networks across Centres of Excellence and sharing best practice in operations and research entity management.

Chief Operating Officers gave a brief introduction of their respective Centres, followed by presentations on key initiatives. Guest speakers included: QUT Vice-Chancellor Professor Margaret Sheil, and, from the ARC, Liz Visher (Director, Program Partnerships) and Luisa Powell (Assistance Director, NCGP Major Investments). We took our colleagues on a tour of the Australian Centre for Robotic Vision's headquarters at QUT, where they met our robots and even took 'a selfie' with SoftBank's social robot, Pepper. We closed the day with a fun personal development exercise led by Jo Karabitsios and David Whittingham from Evexia.

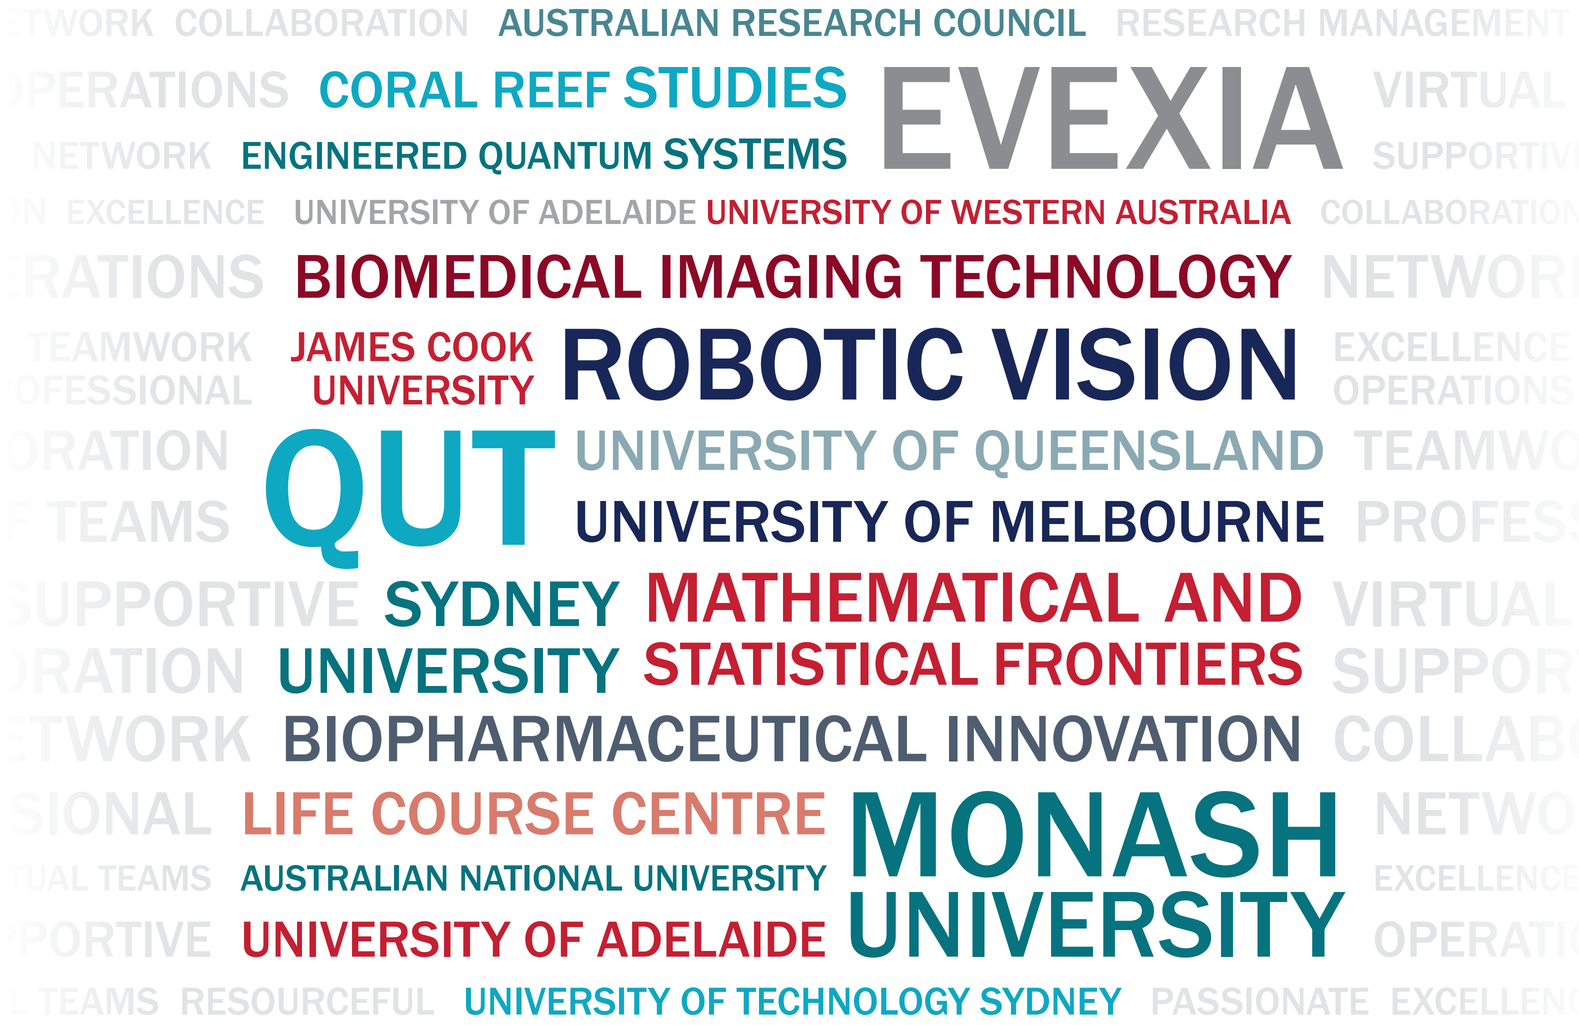

GENDER DIVERSITY

Gender imbalances are acute in computer science and only slightly better in engineering. It is still a huge issue facing our Centre. We are actively pursuing ways to redress this, by employing a number of tactics, such as implementing Gender-based KPI targets (below), addressing unconscious bias in advertising when recruiting for research fellows and PhD researchers. We also engaged external consultants, Gender Matters, to help us investigate other ways to build gender diversity in our field. In October 2017, at our annual RoboVis symposium, the Centre launched a Robotic Vision Gender Equity Plan. This plan incorporates feedback from across the Centre, in addition to the recommendations from the Gender Matters' report. A copy of our Robotic Vision Gender Equity Plan can be downloaded from the Centre's website.

In 2018 we welcomed Centre Advisory Board member Kylie Ahern; Associate Investigators Professor Elizabeth Croft and Dr Miaomiao Liu; Research Fellows Dr Nicole Robinson and Dr Hui Li; Research Engineer Garima Samvedi; and 12 new female PhD researchers.

Performance Measure	Reporting Frequency	Target 2018	Outcome 2018
Whole of Centre event with >80% attendance (80-100 people) Gender Equality and Diversity Event at the Centre's annual symposium, RoboVis. Includes a high-profile speaker describing organisational change that can be achieved in support of increasing diversity and impacts of individual actions	Annually	1	**1**
Centre travel support for staff with caring responsibilities accessed by eli-gible staff as per Centre Gender Equity Policy	Annually	1	**1**
% Female delegates attending RoboVis, the Centre's annual symposium	Annually	20%	**23%**
% Female speakers at RoboVis, the Centre's annual symposium	Annually	20%	**33%**
% Female session chairs at RoboVis, the Centre's annual symposium	Annually	20%	**20%**

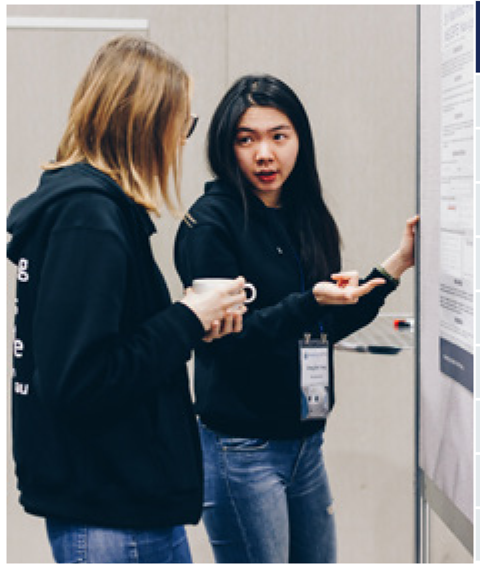

Postion	#	FTE	Gender Ratio (female:male)
Chief Investigator	16	3.6	0:16
Partner Investigator	6	0.3	0:6
Associate Investigator	19	-	2:17
Operations Team - Professional Staff*	11	7.6	10:1
Research Engineer/Software Engineer	11	10.1	1:10
Project Manager/Officer	2	2	2:0
Research Fellow	31	31	6:25
PhD Researcher	74	74	16:58

**Professional Staff include full-time and part-time operations staff*

Diversity in STEM

One of the Centre's Gender KPIs is a gender equality and diversity event at our annual symposium, RoboVis.

Each year the event will include a high-profile speaker describing organisational change that can be achieved in support of increasing diversity and the impacts of individual actions. The 'Diversity in STEM' talk was a key highlight of this year's program, and we were fortunate to welcome two guest speakers: Professor Brian P. Schmidt (AC FAA FRS) and Professor Elanor Huntington.

Their talks centred around four key themes. Firstly, how to make our Centre's culture and environment welcoming, nurturing and empowering. Secondly, how to successfully advertise and recruit and create an inclusive work culture, followed by strategies for systematically removing barriers for advancement across all levels of our respective institutions. Finally, the discussions focused on how the Centre can create a pipeline for the future and ways to make significant changes from primary school level right through to our universities and beyond.

Professor Schmidt is the Vice-Chancellor and President of The Australian National University. He is committed to growing the profile of Science and Technology in the community and attracting more young women to take up careers in STEM related fields. He is on a mission to encourage young women to get involved in sciences and maths.

Professor Huntington became the first female Dean of Engineering and Computer Science at The Australian National University when she was appointed in June 2014. She is committed to growing the profile of Engineering in the community and is passionate about attracting more young women to take up careers in STEM related fields.

EDUCATION

Our Robotic Vision Summer School (RVSS) is fast gaining an international reputation as a 'bucket list' experience for graduate students and industry researchers to learn about fundamental and advanced topics in computer vision and robotics. Staged at ANU's Kioloa Coastal Campus on the southern coast of New South Wales, it also promises a unique immersion into the beauty of Australia, where beach meets bush and Aussie wildlife abound. Since its launch in 2015, the event has drawn a total of 315 attendees including both international and domestic students and expert speakers. In 2018, the summer school hosted 83 attendees, made up of 34 domestic and 16 international students (including six from overseas partner organisations). Four international speakers also attended. RVSS comprises an intensive week of lectures, technical deep dives and a hands-on robotics workshop. For many students, the workshop provides their first opportunity to experiment with computer vision algorithms on actual robotic hardware.

Our knowledge sharing extends to Centre Director Peter Corke's collaboration with QUT to develop an open online robotics education resource, the QUT Robot Academy (**www.robotacademy.net.au**).

The academy celebrated its first birthday in 2018, drawing more than 70,000 international users and over 500,000 lesson views. It is an education resource that provides more than 200 short lessons on specific robotics topics, including robot arms and computer vision, with courses at the equivalent of a second or third year university level. Distinguished Professor Corke said: "The fact that we are able to help people to master this subject material, any time of the day or night, and no matter where they are located in the world, is great news for the Robot Academy and it is also great news for the future of robotics."

Meet our people

Chief Investigators (CIs)

Gustavo Carneiro
University of Adelaide

Tat-Jun Chin
University of Adelaide

Peter Corke
QUT

Tom Drummond
Monash

Matt Dunbabin
QUT

Stephen Gould
ANU

Richard Hartley
ANU

Hongdong Li
ANU

Robert Mahony
ANU

Michael Milford
QUT

Ian Reid
University of Adelaide

Jonathan Roberts
QUT

Chunhua Shen
University of Adelaide

Niko Sünderhauf
QUT

Anton van den Hengel
University of Adelaide

Gordon Wyeth
QUT

Associate Investigators (AIs)

Luis Mejias Alvarez
QUT

Nick Barnes
ANU/Data61

Ross Crawford
QUT

Anthony Dick
University of Adelaide

Anders Eriksson
QUT

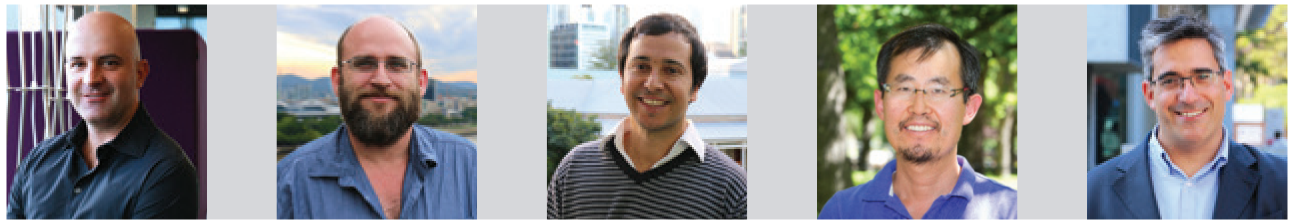

Clinton Fookes
QUT

Jason Ford
QUT

Felipe Gonzalez
QUT

Jonghyuk Kim
ANU

Tristan Perez
QUT

Associate Investigators (AIs) (continued)

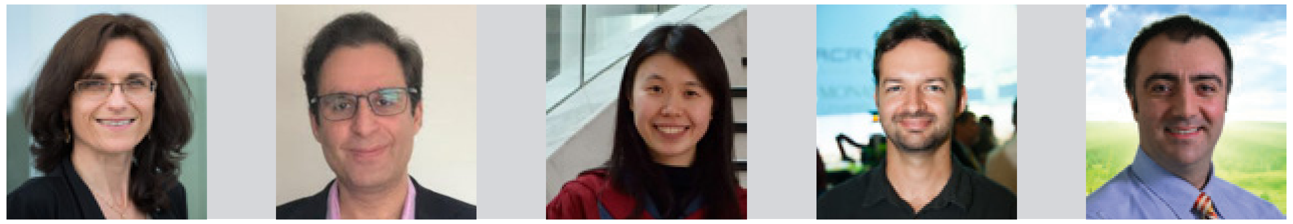

Elizabeth Croft
Monash

Mehrtash Harandi
Monash

Miaomiao Liu
ANU

Thierry Peynot
QUT

Fatih Porikli
ANU/DATA61

Qinfeng 'Javen' Shi
University of Adelaide

David Suter
University of Adelaide

Jochen Trumph
ANU

Qi Wu
University of Adelaide

Partner Investigators (PIs)

François Chaumette
INRIA, France

Andrew Davison
Imperial College
London, UK

Frank Dellaert
Georgia Tech, USA

Paul Newman
University of Oxford, UK

Marc Pollefeys
ETH Zurich, Switzerland

Philip Torr
University of Oxford, UK

Masters Researchers

Norton Kelly-Boxall
Master of Philosophy
Researcher
QUT

Attila Lengyel
Visiting Masters
Researcher
QUT

Project Officers

Belinda Ward
Project Leader
Social Robotics
QUT

Sandy Holmes
Project Officer
Robotics Roadmap
QUT

Research Fellows

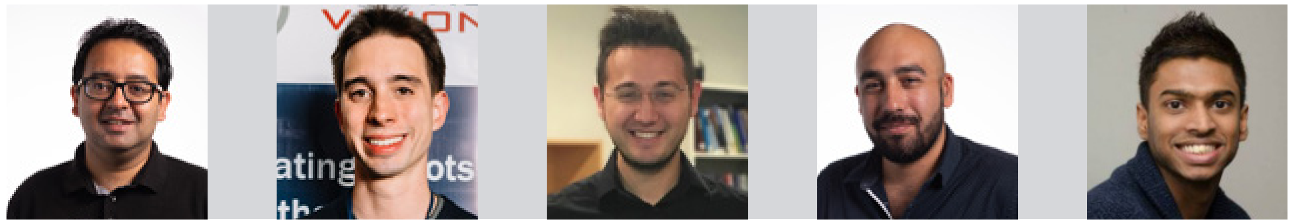

Suman Bista
QUT

Dylan Campbell
ANU

Akansel Cosgun
Monash

Feras Dayoub
QUT

Thanuja Dharmasiri
Monash

Thanh-Toan Do
University of Adelaide

Masoud Faraki
Monash

Basura Fernando
ANU

David Hall
QUT

Ben Harwood
Monash

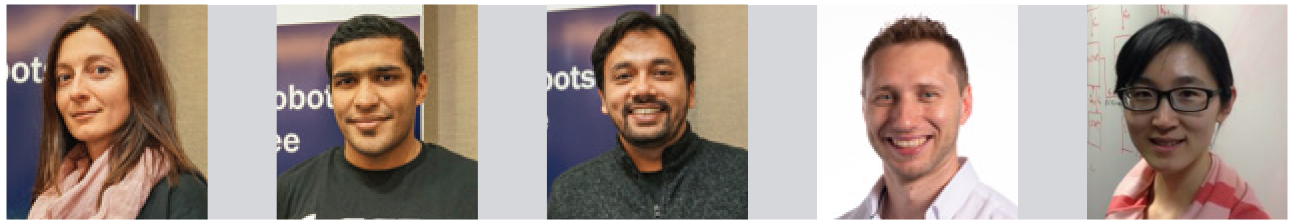

Viorela Ila
ANU

Vijay Kumar
University of Adelaide

Yasir Latif
University of Adelaide

Jürgen 'Juxi' Leitner
QUT

Hui Li
University of Adelaide

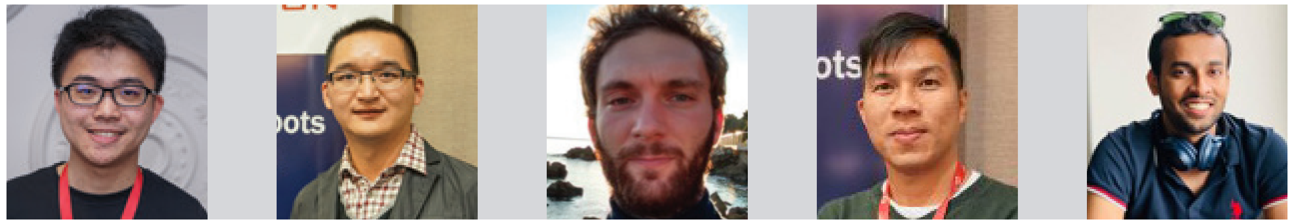

Vincent Lui
Monash

Chao Ma
University of Adelaide

Valerio Ortenzi
QUT

Trung Than Pham
University of Adelaide

Pulak Purkait
University of Adelaide

Yuankai Qi
University of Adelaide

Nicole Robinson
QUT

Sareh Rowlands
QUT

Michele 'Mike' Sasdelli
University of Adelaide

Andrew Spek
Monash

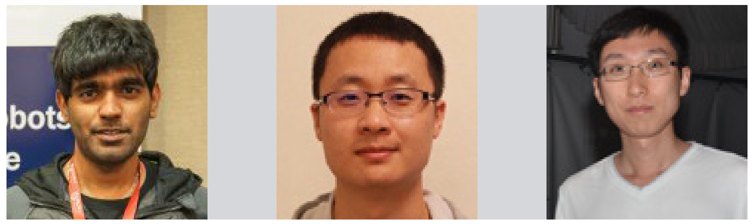

Saroj Weerasekera
University of Adelaide

Haoyang Zhang
QUT

Bohan Zhuang
University of Adelaide

Associated Research Fellows

Ravi Garg
University of Adelaide

Adam Jacobson
QUT

Hamid Rezatofighi
University of Adelaide

Fan Zeng
QUT

Research Affiliates

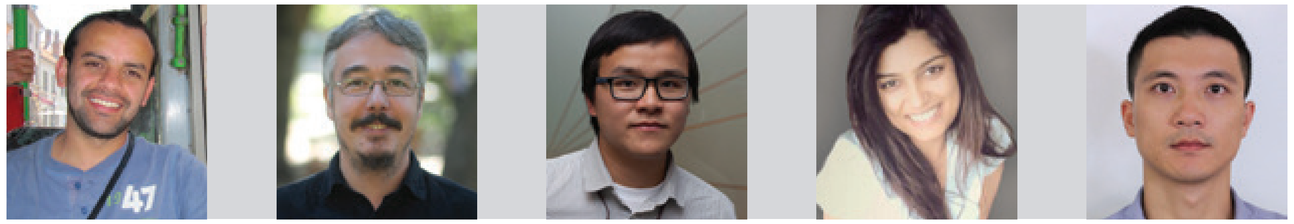

Cesar Cadena
ETH Zurich, Switzerland

Donald Dansereau
University of Sydney

Zongyuan Ge
Monash/NVidia

Anjali Jaiprakash
QUT

Pan Ji
NEC Laboratories

Laurent Kneip
Shanghai Tech, China

Chris Lehnert
QUT

Chris McCool
University of Bonn
Germany

Mark McDonnell
UniSA

Anton Milan
Amazon, Germany

Chuong Nguyen
CSIRO

Ajay Pandey
QUT

Ahmet Sekercioglu
Monash / Sorbonne
University, France

Ben Upcroft
Oxbotica

Liao 'Leo' Wu
University of
New South Wales

PhD Researchers

Shahnewaz Ali, QUT | Rafael Felix Alves, University of Adelaide | Gil Avraham, Monash | Artur Banach, QUT | Ming Cai, University of Adelaide | William Chamberlain, QUT

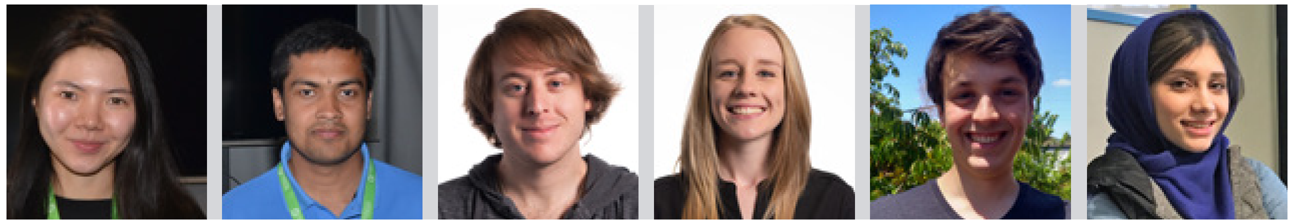

Shin Fang Ch'ng, University of Adelaide | Arif Chowdhury, ANU | Tom Coppin, QUT | Emily Corser, QUT | Luke Ditria, Monash | Mahsa Ehsanpour, University of Adelaide

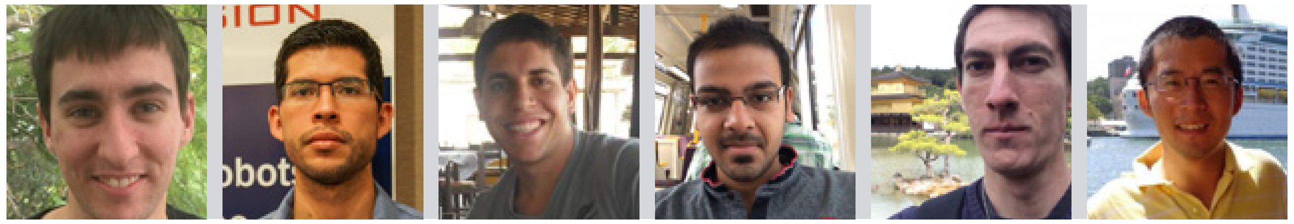

Jordan Erskine, QUT | Andres Felipe Marmol Velez, QUT | Rodrigo Fonseca Santa Cruz Oliveira, ANU | Sourav Garg, QUT | Luis Guerra Fernandez, Monash | Jian 'Edison' Guo, ANU

Augustus 'Gus' Hebblewhite, Monash | Mina Henein, ANU | William Hooper, QUT | Mehdi Hosseinzadeh, University of Adelaide | Natalie Jablonsky, QUT | Jasmin James, QUT

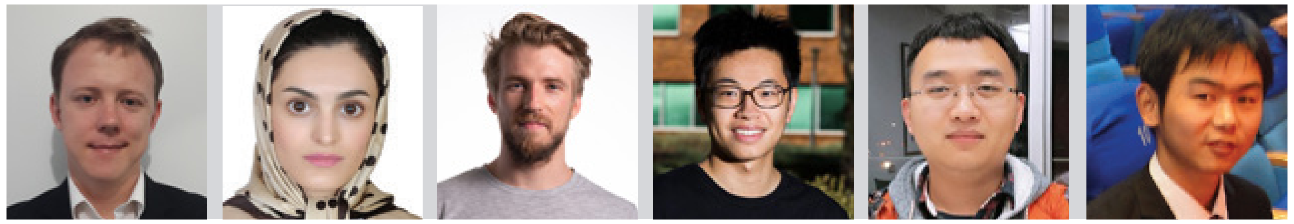

Chris Jeffrey, QUT | Samira Kaviani, ANU | Robert Lee, QUT | Kejie 'Nic' Li, University of Adelaide | Liu Liu, ANU | Yu Liu, University of Adelaide

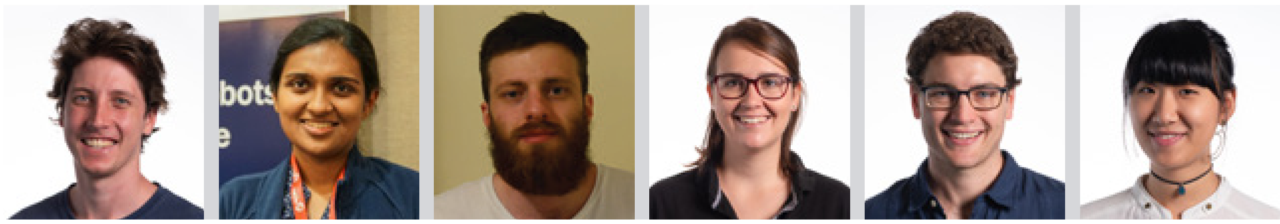

Sean McMahon, QUT | Mehdani Menikdiwela, ANU | Ben Meyer, Monash | Dimity Miller, QUT | Douglas Morrison, QUT | Ya 'Serena' Mou, QUT

PhD Researchers (continued)

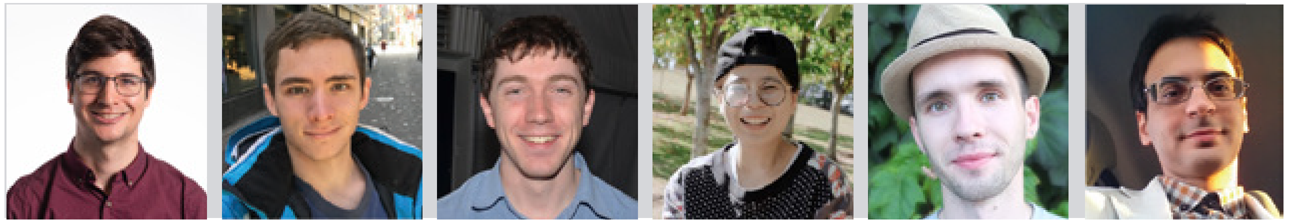

James Mount, QUT | Lachlan Nicholson, QUT | Sean O'Brien, ANU | Liyuan Pan, ANU | Vladimir Nekrasov, University of Adelaide | Amir Rahimi, ANU

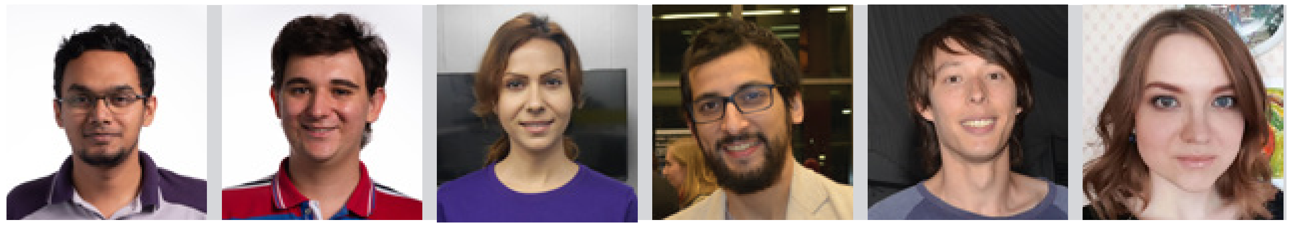

Quazi Marufur 'Maruf' Rahman, QUT | Andrew Razjigaev, QUT | Fahimeh Rezazadegan, QUT | Cristian Rodriguez Opazo, ANU | Cedric Scheerlinck, ANU | Violetta Shevchenko, University of Adelaide

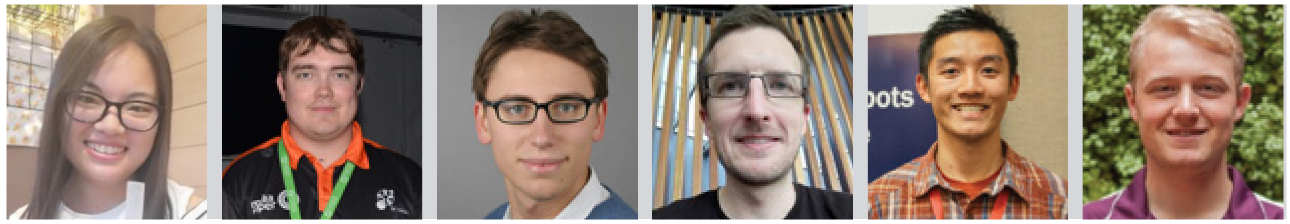

Yujiao Shi, ANU | Jean-Luc Stevens, ANU | Timo Stoffregen, Monash | Brendan Tidd, QUT | Dorian Tsai, QUT | Pieter van Goor, ANU

Yunyan Xing, Monash | Huangying Zhan, University of Adelaide | Jun Zhang, ANU | Zhen 'Frederic' Zhang, ANU | Zheyu Zhang, ANU | Tianyu 'Alan' Zhu, Monash

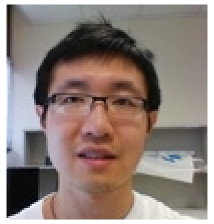

Yan Zuo, Monash

Associated PhD Researchers

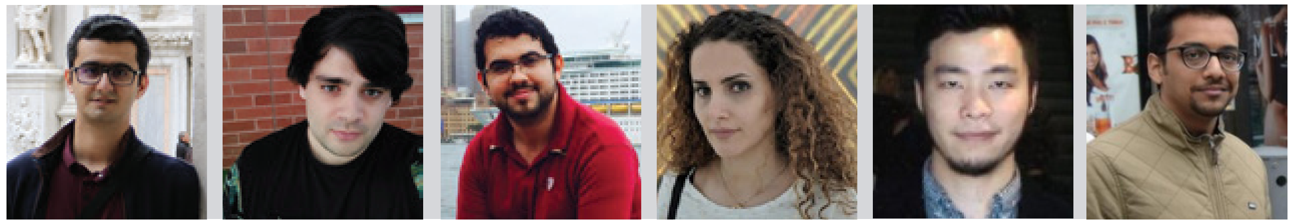

- Mohammad Sadegh Aliakbarian – ANU
- Adrian Johnston – University of Adelaide
- Ehab Salahat – ANU
- Fatemeh Shiri – ANU
- Yao Lu – ANU
- Kartik Gupta – ANU

- Jing Zhang – ANU
- Shihao Jiang – ANU
- Tong Zhang – ANU
- Yiran Zhong – ANU
- Xin Yu – ANU
- Zhiwei Xu – ANU

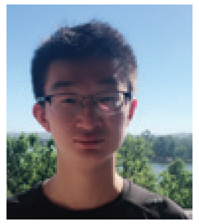

- Ziang Cheng – ANU

PHD Graduates // Congratulations to all our 2018 graduates!

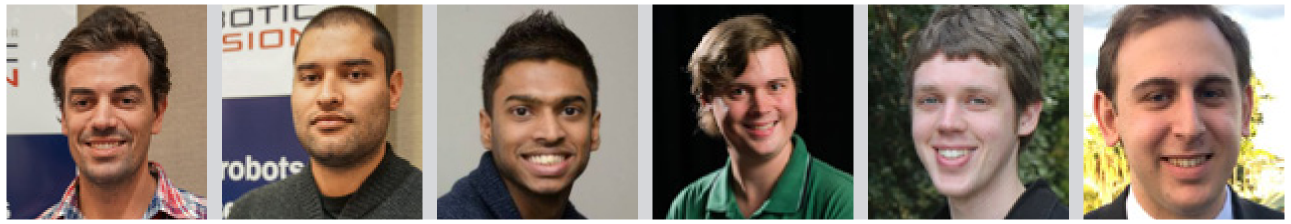

- Peter Anderson – ANU
- Juan David Adarve Bermudez – ANU
- Thanuja Dharmasiri – Monash
- David Hall – QUT
- Ben Harwood – Monash
- Adam Jacobson – QUT

- Hui Li – University of Adelaide
- Andrew Spek – Monash
- Saroj Weerasekera – University of Adelaide
- Josh Weberruss – Monash
- Fangyi Zhang – QUT
- Yi Zhou – ANU

Engineers

- Sam Bahrami – Research Programmer – University of Adelaide
- Tao Hu – Software Engineer – ANU
- Gerard Kennedy – Research Assistant – ANU
- Alex Martin – Research Engineer – ANU
- Steve Martin – Research Engineer – QUT
- Tom Rowntree – Research Engineer – University of Adelaide

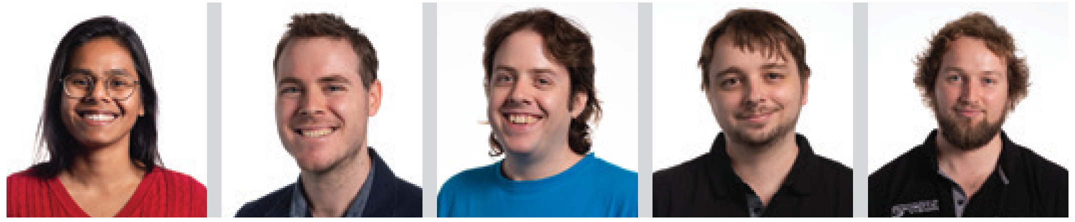

- Garima Samvedi – Research Engineer – QUT
- John Skinner – Research Associate – QUT
- Rohan Smith – Research Engineer – QUT
- Gavin Suddrey – Software Engineer – QUT
- Ben Talbot – Research Engineer – QUT

Centre Operations team

- Kate Aldridge – Centre Coordinator – QUT
- Sarah Allen – QUT Node Administration Officer (Part-Time) – QUT
- Shani Fernando – Finance & Administration Officer (Part-Time) – QUT
- Sue Keay – Chief Operating Officer – QUT
- Ireen Khan – PA to Centre Director – QUT
- Tracy Kelly – Finance & Administration Officer (Part-Time) – QUT

- Tim Macuga – Communications Specialist (Part-Time) – QUT
- Thuy Mai – University of Adelaide Node Administration Officer (Part-Time) – University of Adelaide
- Sandra Pedersen – Monash University Node Administration Officer (Part-Time) – Monash
- Carol Taylor – ANU Node Administration Officer (Part-Time) – ANU
- Shelley Thomas – Communications Specialist – QUT

Legacy of a 'positive disruptor'

Australian Centre for Robotic Vision Chief Operating Officer Dr Sue Keay, after four-and-a-half years at the world-first research centre, bade farewell at the end of 2018 to forge a new path with Australia's national science agency.

From January 2019, the bona fide Superstar of STEM – one of the first named by Science & Technology Australia in 2017 – joins CSIRO's Data61 as Physical Systems Research Director.

"We're certainly going to miss Sue, but Australia will continue to benefit from her passion and tenacity as a positive disruptor," said Centre Director Peter Corke.

"Sue helped build our ARC Centre of Excellence from scratch, understanding that the breakthrough science and technologies needed to create a new generation of 'truly useful' robots – able to see and understand the environments they work in – could only be achieved through concerted, large-scale and collaborative effort.

"As part of this, Sue has not only successfully driven the Centre's strategic direction, in the process drawing national recognition as a powerful advocate for transformational change and champion for diversity in all its forms, but helped shape a national agenda developing Australia's first Robotics Roadmap."

Recognised as one of Queensland's most influential people, Dr Keay has been the Chief Operating Officer of the Australian Centre for Robotic Vision since it was created in 2014. Over the past four years, her achievements include development of a successful $1.5-million R&D project on humanoid robotics supported by the Queensland Government to explore the vision capabilities of SoftBank's social robot, Pepper; and delivery of *A Robotics Roadmap for Australia 2018*.

Further, in a personal 'win' as a champion for diversity, Dr Keay has been the driver behind the realisation of Hopper Down Under coming to Brisbane (29-31 July 2019). The event is best described as a little sister of sorts to the Grace Hopper Celebration – the world's largest gathering of women technologists produced by AnitaB.org in the United States. Staged annually, it draws 18,000+ attendees from 81 countries.

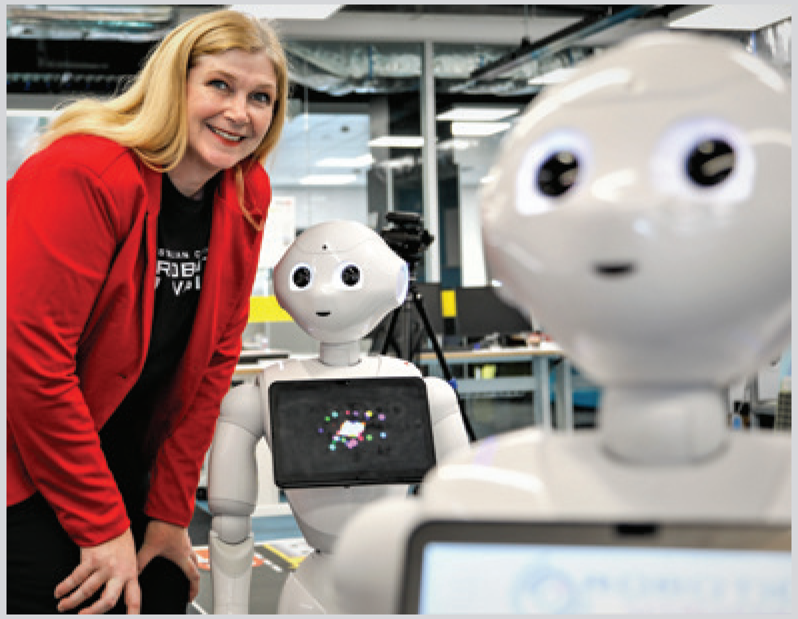

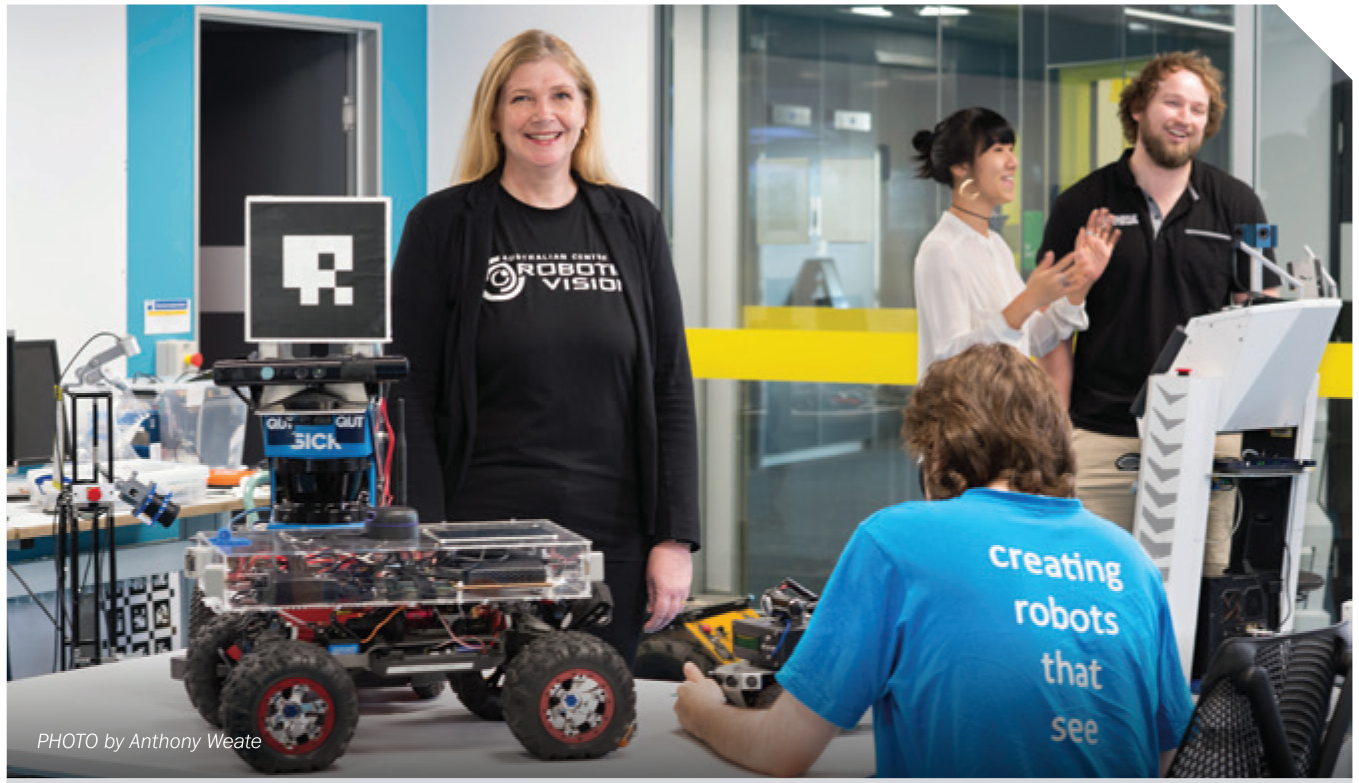

PHOTO by Anthony Weate

Dr Keay travelled to Houston to meet with AnitaB.org in September 2018, inviting them to bring the Grace Hopper Celebration to Australia. Hopper Down Under is set to become an annual celebration and, in its inaugural year, has expanded to cover the Asia-Pacific region. It will include technical and career development sessions, including Poster Sessions, a Career Fair, a Tech Expo, Mentoring Circles, and more.

In her new role at CSIRO's Data61 – itself a partner organisation of the Australian Centre for Robotic Vision – Dr Keay will lead research encompassing robotics and autonomous systems, distributed sensing networks, 3D mapping, AI-enabled computer vision and cybernetics.

A strong advocate of a collaborative, multi-sector approach, she will also continue to advocate recommendations contained in the Robotics Roadmap, focusing on ways industry, government, educational institutions, investors and the wider public can better harness new and emerging technologies.

"I'm proud of my work at the Australian Centre for Robotic Vision and indebted to lessons learned there, not least being that a leader is only as good as his or her team," Dr Keay said.

"As Chief Operating Officer at the Centre, I have been surrounded by the world's best researchers. I look forward to watching them continue to achieve the vision of creating truly useful robots able to see and understand for the good of all people and our planet.

"I also look forward to continuing to play a part in building Australia's future. A prosperous Australia that embraces new and transformative technologies to bolster national health, wellbeing and sustainability, in turn overcoming challenges of our vast and remote geography."

With interests in entrepreneurship and disruptive technologies, Dr Keay recently completed her MBA with UQ Business School. A graduate of the Australian Institute of Company Directors, she serves on the Board of the CRC for Optimising Resource Extraction and the expert advisory panel of Queensland.AI and mentors female-led start-up companies.

A university medallist and Jaeger scholar, Dr Keay has more than 20 years' experience in the research sector. She has a PhD in Earth Sciences from the Australian National University and was an ARC post-doctoral fellow at The University of Queensland, before escaping the lab and moving first to science communication and then into research management, research commercialisation and IP management.

Section 7 **Governance**

Good governance of the Centre is the responsibility of our Centre Executive Committee, with oversight provided by the Centre Advisory Board.

In 2018, the Centre Advisory Committee and End User Advisory Board were restructured to form a new Centre Advisory Board.

CENTRE ADVISORY BOARD

Our Centre Advisory Board oversees the Centre's overall strategic direction; provides independent scientific expertise and advice regarding the Centre's Research Program and engagement with industry, government and the general public. Its membership comprises an independent Chair and up to six members representing academia and industry with expertise in the relevant science and with a track record in technology commercialisation.

PROFESSOR ALEX ZELINSKY AO (CHAIR)

Professor Zelinsky is the Chair of the Centre's Advisory Board, and Vice-Chancellor and President of The University of Newcastle.

Prior to joining the University, he was Australia's Chief Defence Scientist and head of the Defence Science and Technology (DST) Group. Before joining Defence, Professor Zelinsky was Group Executive for Information Sciences at the Commonwealth Scientific and Industrial Research Organisation (CSIRO) and Director of CSIRO's Information and Communication Technologies (ICT) Centre. He was Chief Executive Officer and co-founder of Seeing Machines, a high-technology company developing computer vision systems. The company is listed on the London Stock Exchange and was a start-up from the Australian National University in Canberra, Australia, where Professor Zelinsky was Professor of Systems Engineering. Previously he researched robotics and computer vision at the AIST Electro-technical Laboratory in Japan and has taught and conducted research in computer science at the University of Wollongong. Professor Zelinsky has extensive experience in advising Federal and State governments in Australia, including as a member of the Australian Government's Defence Industry Innovation Board. He is a Fellow of the Institute of Electrical and Electronics Engineers, the Australian Academy of Technological Sciences, Engineers Australia, and the Australian Institute of Company Directors. He received the Engineers Australia M A Sargent Medal 2015 – the most prestigious award made by the College of Electrical Engineers and in 2013 he was awarded the prestigious Pearcey Medal, the ICT industry's premier prize for lifetime achievement.

Professor Zelinsky was appointed an Officer in the Order of Australia (AO) in the 2017 Queen's Birthday Honours "for distinguished service to defence science and technology, to systems engineering, and to education as an academic and researcher".

KYLIE AHERN

Ms Ahern is an award-winning science publisher and entrepreneur with experience that spans media, telecommunications, science and education. In 2004 she co-founded Cosmos Media and launched Australia's top-selling science magazine, *Cosmos*. The company was recognised through 54 awards and commendations for the high quality of its publications, websites, science outreach programs, publishing and journalism. Ms Ahern created educational products – teachers' notes, study guides, posters, career guides – to which 70 per cent of Australia's high schools and hundreds of overseas subscribed to reach more than 130,000 students. She transitioned the business from print to multiple media platforms, with Apple naming it among its 'Best of 2012' app launches.

Since selling Cosmos Media in 2013, she has been consulting for various organisations. She helped set up the Nature Publishing Group in Australia, conceived of the Queensland Brain Institute's Brain series and, in 2016, founded STEM Matters, a strategy, communications and content agency. Clients include Westpac, the Federal Government and universities and research institutes across Australia.

DR HENRIK I. CHRISTENSEN

Dr Christensen is a Professor of Computer Science at Dept. of Computer Science and Engineering UC San Diego. He is also the director of the Institute for Contextual Robotics. Dr Christensen conducts research on systems integration, human-robot interaction, mapping and robot vision. The research is performed within the Cognitive Robotics Laboratory. He has published more than 350 contributions across AI, robotics and vision. His research has a strong emphasis on 'real problems with real solutions'.

He is actively engaged in the setup and coordination of robotics research in the US (and worldwide). Dr Christensen received the Engelberger Award 2011, the highest honour awarded by the robotics industry. He was also awarded the "Boeing Supplier of the Year 2011". Dr Christensen is a fellow of American Association for Advancement of Science (AAAS) and Institute of Electrical and Electronic Engineers (IEEE). His research has been featured in major media including CNN, *The New York Times* and BBC.

TRENT LUND

Mr Lund is the Lead Partner for Innovation & Digital Ventures at PwC Australia. He helps organisations leverage emerging technologies to transform ideas into customer-centred, commercial outcomes. With two decades of industry knowledge, he has worked across the globe – including in the Asia-Pacific, the United Kingdom and the Middle East. He has worked in business consultancy and new ventures where innovation is leveraged to identify new sources of value

ANDREW HARRIS

Andrew Harris is a Professor of Chemical and Biomolecular Engineering at The University of Sydney, and the Australian director of Laing O'Rourke's future engineering and innovation consultancy, the Engineering Excellence Group. Laing O'Rourke is Australia's largest private engineering and construction business.

Professor Harris received his PhD from the University of Cambridge in 2002 and is a Chartered Engineer and Fellow of the Institution of Chemical Engineers (IChemE) and Engineers Australia (IEAust).

Throughout his career he has worked at the interface between industry and academia.

He is a non-executive Director of Hazer Group (ASX:HZR), a listed clean tech, and serves on the industry advisory board of the Australian Research Council Centre for Robotic Vision.

Professor Harris was recognised as one of Australia's 50 most innovative engineers by peak body, Engineers Australia, in 2016.

RUSSEL RANKIN

Mr Rankin has more than 30 years' experience in the food and beverage industry in various senior commercial and research positions. He is Director and Founder of Food Innovation Partners Pty Ltd, a company that provides business, innovation and commercialisation services to the food industry, along with business development services for companies and research organisations.

Prior to starting Food Innovation Partners, Mr Rankin was General Manager – Innovation with the National Food Industry Strategy: a federal government initiative established to provide leadership to Australia's food industry. He has also worked for CSIRO for more than 25 years in the area of food research before venturing into the commercial arena.

He is Chair of Queensland Department of Agriculture and QUT's Agricultural Robotics program Advisory Board; member of Government of South Australia's Advisory and Assessment Board for Advanced Food Manufacturing program; member of the Advisory Board to KFSU Pty Ltd, a company making dietary fibre from sugar cane; Director of The Food Market Company; Director of Freshly & Co; and Director of Beauty Drink Pty Ltd.

CENTRE EXECUTIVE COMMITTEE

The Executive Committee drives the Centres' responsible governance, with oversight from the Centre Advisory Board. The Executive Committee is accountable to the Australian Research Council, a statutory agency responsible for Australia's National Competitive Grants program and major contributor to the Centre to the value of $19 million in public funding.

The committee includes representatives from the Centre's four domestic partners - QUT, the University of Adelaide, the Australian National University and Monash University. It provides the leadership and direction that is critical to the successful operation of our Centre and achievement of our ambitious research program.

The Executive Committee meets monthly – more regularly when needed – via video conference and a quarterly face-to-face meetings located at different Centre nodes in rotation. Meetings cover a range of operational, research, collaboration and commercialisation matters. All meeting documentation, such as agendas, actions, resolutions and notes, is recorded on the Centre's intranet.

The members of our Executive Committee are:

Centre Director,
Distinguished Professor Peter Corke
(QUT)

Centre Deputy Director, ARC Laureate
Professor Ian Reid
(University of Adelaide)

Chief Operating Officer,
Dr Sue Keay

Chief Investigator,
Professor Stephen Gould
(Australian National University)

Chief Investigator,
Professor Tom Drummond
(Monash University)

CENTRE RESEARCH COMMITTEE

The Centre's Research Committee includes all Research Project Leaders and Demonstrator Leaders. The Committee meets monthly via videoconference, and face-to-face twice a year, to review research progress and make decisions on the Centre's research direction. In 2018, a series of face-to-face project reviews were conducted by members of the Centre Executive Committee at each of the Centre's nodes.

2018 Australian Centre for Robotic Vision Key Performance Indicators

Performance Measure	Reporting Frequency	Target	Outcome
Research Performance			
Journal Articles	Annually	25	**28**
Conference publications	Annually	50	**61**
Disclosures/Patents	Annually	3	**4**
Paper prizes or awards	Annually	5	**5**
Percentage of publications relevant to Strategic Research Priorities (SRPs) [%]	Annually	90%	**90%**
Centre h-index	Annually	41	**59**
Research Training & Development			
Careers Training	Annually	1	**1**
Entrepreneurial (Start Up) Training	Annually	1	**1**
Knowledge Leadership	Annually	1	**6**
Project Management	Annually	1	**1**
Media Training	Annually	1	**1**
Mentoring Programs			
Knowledge Leadership Program	Annually	1	**1**
Welcome/Induction Program	Annually	1	**1**
Engagement			
Annual symposium - RoboVis	Annually	1	**1**
Robotic Vision Summer School	Annually	1	**1**
Workshops held at International Conferences	Annually	2	**8**
Industry Workshops	Annually	1	**1**
Presentations and briefings			
To the public	Annually	5	**61**
To government (parliamentarians and departments/agencies at both State and Federal level)	Annually	5	**28**
To industry/business/end-users	Annually	5	**84**
To non-govenment organisations	Annually	5	**26**
To professional organisations & bodies	Annually	5	**6**

Performance Measure	Reporting Frequency	Target	Outcome
New organisation collaborations			
New participating organisations or organisations we have established a collaborative relationship with	Annually	2	**11**
Additional funding secured			
External income (includes industry, administering organisation, government & other ARC funding) [$mil]	Annually	$2 million	**$9,497,707**
People			
Number of additional researchers working on Centre research			
Postdoctoral researchers	Annually	7	**11**
Honours students	Annually	8	**26**
PhD students	Annually	10	**25**
Postgraduate completions			
PhD & Masters completions	Annually	10	**12**
Honours completions	Annually	8	**19**
Gender Equity			
Whole of Centre event with >80% attendance (80-100 people). Gender Equality and Diversity Event at the Centre's annual symposium, RoboVis. Includes a high profile speaker describing organisational change that can be achieved in support of increasing diversity and impacts of individual actions.	Annually	1	**1**
Centre travel support for staff with caring responsibilities accessed by eligible staff as per Centre Gender Equity Policy	Annually	1	**1**
% Female delegates attending RoboVis, the Centre's annual symposium	Annually	20%	**23%**
% Female speakers at RoboVis, the Centre's annual symposium	Annually	20%	**33%**
% Female session chairs at RoboVis, the Centre's annual symposium	Annually	20%	**20%**

Section 8 **Finance and Operations**

Financial **Performance**

Our Financial Statement provides a summary of our financial performance for the 2018 calendar year.

The Centre receives funding from two main sources, the Australian Research Council (ARC) and the Centre's collaborating organisations. In addition to cash, our collaborating and partner organisations provide significant resources through 'in-kind' contributions, mainly in the form of our Chief Investigators and Associate Investigators' time, space, central computing and networking. We also source additional funding via industry engagement and relevant industry projects, raising over $32 million over the life of the Centre so far.

FINANCIAL PERFORMANCE

The Centre's financial performance for the calendar year is summarised in our 2018 Income and Expenditure Report (see page 132). Our expenditure is predominantly allocated to salaries and staff expenses, with the remainder allocated to equipment purchases, travel and professional development and operating expenses. The Centre accumulated a significant carry forward of funds in 2014 due to delays in establishing the legal agreements governing the operation of the Centre which prevented us from recruiting. Carry-forward continued into 2015 and 2016 but we were actively spending both income and carry forward in 2017 and 2018. The carry-forward is fully allocated in our budget and we project it to be fully expended with annual expenditure forecast to exceed 'new' annual income for the remainder of the Centre's life to ensure funds will be fully expended.

New income is approximately $3.7million, in cash, that we are contracted to receive each year until 2021. The Centre Executive determines budget allocations based on the Centre's original bid submission, with all nodes contributing to operating expenses. We have centralised some expenditure at QUT, with budgets across the life of the Centre developed by the mutual consent of all members of the Executive. We have created a pool of untied cash to support strategic initiatives by using the (not guaranteed) indexation funds sent to the Administering Organisation by the ARC. Some of these funds were used in 2018 to support the Centre's new Gender Equity Plan which funds 50 per cent of the salary of female postdoctoral research fellows, and travel support for researchers with parenting commitments. These initiatives are one response designed to address some of the structural barriers preventing female participation in robotics and computer vision.

Another key activity for our strategic funds in 2018 was the creation and launch of the Robotics Roadmap for Australia.

SUMMARY OF CONTRIBUTIONS FROM ALL CENTRE PARTNERS

Each year, our administrative and collaborative partner organisations contribute $940,000 as cash, a total of more than $6.86 million over the life of the Centre. Over $1 million is contributed as in-kind, or $6.98 million over the life of the Centre. Together, our international partner organisations contribute $139,000 per annum of in-kind, totalling $973,000 over the Centre's seven-year lifespan. The Centre's collaborating Partners, where most of our researchers are based, also provide access to a broad range of robotic vision equipment, which is conservatively valued at over $1 million per annum ($7 million in total, over the life of the Centre). A summary of the total cash and in-kind contributions from our Partners over the Centre's seven-year term is shown on page 133. All figures quoted are in Australian dollars.

Each year, our administrative and collaborative partner organisations contribute

$940,000

as cash,

...a total of more than

$6.86 million

over the life of the Centre

STATEMENT OF OPERATING INCOME AND EXPENDITURE FOR YEAR END 31 DECEMBER 2018

Income	2014	2015	2016	2017	2018
ARC Centre Grant	2,714,290	2,714,284	2,714,285	2,714,285	2,714,285
Monash	116,929	116,929	116,929	116,929	116,929
University of Adelaide	244,400	244,400	244,400	244,400	243,163
ANU	230,000	230,000	230,000	230,000	230,000
QUT	360,000	350,000	350,000	350,000	350,000
ARC Indexation	82,027	132,110	180,499	223,919	267,993
Interest from QUT (commenced 2016)			6,359	64,175	86,889
Other Income to Administrating Organisation (QUT)				220,471	44,990
Total Income	**3,747,646**	**3,787,723**	**3,842,472**	**4,164,179**	**4,054,249**

Expenditure	2014	2015	2016	2017	2018
Purchased Equipment	42,256	51,829	57,325	1,638	52,471
Shared Equipment/Facilities	-	815	-	-	-
Travel and Professional Development	24,276	253,180	321,077	609,953	645,568
Maintenance (IT and Lab)	8,350	19,639	71,769	109,516	87,884
Salaries/Personnel	373,061	2,337,787	2,813,795	3,138,116	3,880,286
Other	38,002	180,876	192,309	248,253	235,899
Total Expenditure	**485,946**	**2,844,126**	**3,456,274**	**4,107,476**	**4,902,107**
Surplus/Deficit	3,261,700	943,597	386,198	56,703	- 847,858
Previous year carry forward	-	3,261,700	4,205,297	4,591,495	4,648,198
Total carry forward surplus*	**3,261,700**	**4,205,297**	**4,591,495**	**4,648,198**	**3,800,339**

FINANCE KPIS

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017	Target 2018	Outcome 2018
Annual cash contributions from administering and collaborating organisations	Annually										
ANU		230,000	230,000	230,000	230,000	230,000	230,000	230,000	230,000	230,000	230,000
Monash		116,400	116,929	116,400	116,929	116,400	116,929	116,400	116,929	116,400	116,929
QUT		350,000	360,000	350,000	350,000	350,000	350,000	350,000	350,000	350,000	350,000
University of Adelaide		244,400	244,400	244,400	244,400	244,400	244,400	244,400	244,400	244,400	243,163
Annual in-kind contributions from administering and collaborating organisations	Annually										
ANU		332,000	114,000	332,000	150,000	316,000	411,926	322,000	544,445	329,000	569,667
Monash		151,000	151,000	151,000	229,000	159,000	180,275	163,000	193,090	168,000	221,493
QUT		309,000	309,000	309,000	250,000	309,000	490,759	309,000	645,186	309,000	713,558
University of Adelaide		153,000	200,000	153,000	205,000	184,000	211,172	184,000	267,411	184,000	330,623
Annual cash contributions from Partner organisations	Annually	Nil	Nil	Nil	Nil	Nil	Nil	Nil	Nil	Nil	Nil
Annual in-kind contributions from Partner organisations	Annually										
Georgia Institute of Technology		18,000	18,000	18,000	18,000	18,000	-	18,000	-	18,000	18,000
Imperial College London		18,100	3,600	18,100	18,100	18,100	18,100	18,100	18,100	18,100	18,100
INRIA		14,500	-	14,500	14,500	14,500	14,500	14,500	14,500	14,500	14,500
Data61/CSIRO		30,000	30,000	30,000	30,000	30,000	30,000	30,000	30,000	30,000	30,000
University of Oxford		32,000	32,000	32,000	32,000	32,000	32,000	32,000	16,000	32,000	32,000
ETH Zürich		26,400	-	26,400	26,400	26,400	26,400	26,400	-	26,400	26,400
Other research income sourced by the Centre - End User (industry, public sector, ARC Linkage and Discovery in non-core areas, CRC)	Annually	$0	$1m	$1m	$4.1m	$1.5m	$8m	$1.5m	$9.5m	$2m	$9.5m
Number of new organisations collaborating with, or involved in, the Centre	Annually	2	2	2	2	2	4	2	13	2	11
Level and quality of infrastructure provided to the Centre	All review										

Outputs

2018 Publications

*denotes Core Centre Research Output

JOURNAL ARTICLES (40)

*Bilen, H., Fernando, B., Gavves, E., & Vedaldi, A. (2018). Action Recognition with Dynamic Image Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, 40(12), 2799–2813. http://doi.org/10.1109/TPAMI.2017.2769085

*Campbell, D. J., Petersson, L., Kneip, L., & Li, H. (2018). Globally-Optimal Inlier Set Maximisation for Camera Pose and Correspondence Estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1–1. http://doi.org/10.1109/TPAMI.2018.2848650 *In Press

*Chen, Z., Liu, L., Sa, I., Ge, Z., & Chli, M. (2018). Learning Context Flexible Attention Model for Long-Term Visual Place Recognition. IEEE Robotics and Automation Letters, 3(4), 4015–4022. http://doi.org/10.1109/LRA.2018.2859916

*de Marco, S., Marconi, L., Mahony, R., & Hamel, T. (2018). Output regulation for systems on matrix Lie-groups. Automatica, 87, 8–16. https://doi.org/10.1016/J.AUTOMATICA.2017.08.006

Guo, G., Wang, H., Shen, C., Yan, Y., & Liao, H.-Y. M. (2018). Automatic Image Cropping for Visual Aesthetic Enhancement Using Deep Neural Networks and Cascaded Regression. IEEE Transactions on Multimedia, 20(8), 2073–2085. http://doi.org/10.1109/TMM.2018.2794262

*Hall, D., Dayoub, F., Perez, T., & McCool, C. (2018). A rapidly deployable classification system using visual data for the application of precision weed management. Computers and Electronics in Agriculture, 148, 107–120. http://doi.org/10.1016/J.COMPAG.2018.02.023

Han, X., Lu, J., Zhao, C., You, S., & Li, H. (2018). Semi supervised and Weakly Supervised Road Detection Based on Generative Adversarial Networks. IEEE Signal Processing Letters, 25(4), 551–555. http://doi.org/10.1109/LSP.2018.2809685

Hodgson, J. C., Mott, R., Baylis, S. M., Pham, T. T., Wotherspoon, S., Kilpatrick, A. D., Ramesh, R.S., Reid, I., Terauds, A., & Koh, L. P. (2018). Drones count wildlife more accurately and precisely than humans. Methods in Ecology and Evolution, 9(5), 1160–1167. http://doi.org/10.1111/2041-210X.12974

*Jacobson, A., Chen, Z., & Milford, M. (2018). Leveraging variable sensor spatial acuity with a homogeneous, multi-scale place recognition framework. Biological Cybernetics, 1–17. http://doi.org/10.1007/s00422-017-0745-7

*James, J., Ford, J. J., & Molloy, T. L. (2018). Learning to Detect Aircraft for Long-Range Vision-Based Sense-and-Avoid Systems. IEEE Robotics and Automation Letters, 3(4), 4383–4390. http://doi.org/10.1109/LRA.2018.2867237

*James, J., Ford, J. J., & Molloy, T. L. (2018). Quickest Detection of Intermittent Signals With Application to Vision-Based Aircraft Detection. IEEE Transactions on Control Systems Technology, 1–8. http://doi.org/10.1109/TCST.2018.2872468

*Li, B., He, M., Dai, Y., Cheng, X., & Chen, Y. (2018). 3D skeleton based action recognition by video-domain translation-scale invariant mapping and multi-scale dilated CNN. Multimedia Tools and Applications. https://doi.org/10.1007/s11042-018-5642-0

Li, H., Wang, P., You, M., & Shen, C. (2018). Reading car license plates using deep neural networks. Image and Vision Computing, 72, 14–23. http://doi.org/10.1016/J.IMAVIS.2018.02.002

Li, X., Zhao, L., Ji, W., Wu, Y., Wu, F., Yang, M.-H., Dacheng, T., Reid, I. (2018). Multi-Task Structure-aware Context Modeling for Robust Keypoint-based Object Tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1–1. https://doi.org/10.1109/TPAMI.2018.2818132 *In Press

*Liao, Z., Drummond, T., Reid, I., & Carneiro, G. (2018). Approximate Fisher Information Matrix to Characterise the Training of Deep Neural Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1–1. http://doi.org/10.1109/TPAMI.2018.2876413

*Lin, G., Shen, C., van den Hengel, A., & Reid, I. (2018). Exploring Context with Deep Structured Models for Semantic Segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 40(6), 1352–1366. http://doi.org/10.1109/TPAMI.2017.2708714

*Liu, L., Li, H., Dai, Y., & Pan, Q. (2018). Robust and Efficient Relative Pose With a Multi-Camera System for Autonomous Driving in Highly Dynamic Environments. IEEE Transactions on Intelligent Transportation Systems, 19(8), 2432–2444. http://doi.org/10.1109/TITS.2017.2749409

Liu, W., Zhang, P., Chen, X., Shen, C., Huang, X., & Yang, J. (2018). Embedding Bilateral Filter in Least Squares for Efficient Edge-preserving Image Smoothing. IEEE Transactions on Circuits and Systems for Video Technology, 1–1. http://doi.org/10.1109/TCSVT.2018.2890202

*Lu, H., Shen, C., Cao, Z., Xiao, Y., & van den Hengel, A. (2018). An Embarrassingly Simple Approach to Visual Domain Adaptation. IEEE Transactions on Image Processing, 27(7), 3403–3417. https://doi.org/10.1109/TIP.2018.2819503

Mao, J., Hu, X., & Milford, M. (2018). An adaptive localization system for image storage and localization latency requirements. Robotics and Autonomous Systems, 107, 246–261. http://doi.org/10.1016/J.ROBOT.2018.06.007

*McCool, C., Beattie, J., Milford, M., Bakker, J. D., Moore, J. L., & Firn, J. (2018). Automating analysis of vegetation with computer vision: Cover estimates and classification. Ecology and Evolution, 8(12), 6005–6015. http://doi.org/10.1002/ece3.4135

*McMahon, S., Sunderhauf, N., Upcroft, B., & Milford, M. (2018). Multimodal Trip Hazard Affordance Detection on Construction Sites. IEEE Robotics and Automation Letters, 3(1), 1–8. http://doi.org/10.1109/LRA.2017.2719763

*Nash, W., Drummond, T., & Birbilis, N. (2018). A review of deep learning in the study of materials degradation. Npj Materials Degradation, 2(1), 37. http://doi.org/10.1038/s41529-018-0058-x

*Nicholson, L., Milford, M., & Sunderhauf, N. (2018). QuadricSLAM: Dual Quadrics From Object Detections as Landmarks in Object-Oriented SLAM. IEEE Robotics and Automation Letters, 4(1), 1–8. http://doi.org/10.1109/LRA.2018.2866205

Palmer, D. W., Coppin, T., Rana, K., Dansereau, D. G., Suheimat, M., Maynard, M. Atchison, D. A., Roberts, J., Crawford, R., & Jaiprakash, A. (2018). Glare-free retinal imaging using a portable light field fundus camera. Biomedical Optics Express, 9(7), 3178. http://doi.org/10.1364/BOE.9.003178

Robinson, N. L., Connolly, J., Johnson, G. M., Kim, Y., Hides, L., & Kavanagh, D. J. (2018). Measures of incentives and confidence in using a social robot. Science Robotics, 3(21), eaat6963. http://doi.org/10.1126/scirobotics.aat6963

*Saha, S. K., Fernando, B., Cuadros, J., Xiao, D., & Kanagasingam, Y. (2018). Automated Quality Assessment of Colour Fundus Images for Diabetic Retinopathy Screening in Telemedicine. Journal of Digital Imaging, 31(6), 869–878. http://doi.org/10.1007/s10278-018-0084-9

*Santa Cruz, R., Fernando, B., Cherian, A., & Gould, S. (2018). Visual Permutation Learning. IEEE Transactions on Pattern Analysis and Machine Intelligence (Vol. PP). IEEE. https://doi.org/10.1109/TPAMI.2018.2873701 *Early Access

*Sünderhauf, N., Brock, O., Scheirer, W., Hadsell, R., Fox, D., Leitner, J., Upcroft, B., Abbeel, P., Burgard, W., Milford, M., & Corke, P. (2018). The limits and potentials of deep learning for robotics. The International Journal of Robotics Research, 37(4–5), 405–420. http://doi.org/10.1177/0278364918770733

Varamin, A. A., Abbasnejad, E., Shi, Q., Ranasinghe, D. C., & Rezatofighi, H. (2018). Deep Auto-Set: A Deep Auto-Encoder-Set Network for Activity Recognition Using Wearables (Vol. 18). Retrieved from https://doi.org/10.475/123_4

*Wang, H., Guobao, X., Yan, Y., & Suter, D. (2018). Searching for Representative Modes on Hypergraphs for Robust Geometric Model Fitting. IEEE Transactions on Pattern Analysis and Machine Intelligence. https://doi.org/10.1109/TPAMI.2018.2803173 *In Press

*Wang, X., Şekercioğlu, Y., Drummond, T., Frémont, V., Natalizio, E., & Fantoni, I. (2018). Relative Pose Based Redundancy Removal: Collaborative RGB-D Data Transmission in Mobile Visual Sensor Networks. Sensors, 18(8), 2430. http://doi.org/10.3390/s18082430

Yan, Y., Tan, M., Tsang, I., Yang, Y., Shi, Q., & Zhang, C. (2018). Fast and Low Memory Cost Matrix Factorization: Algorithm, Analysis and Case Study. IEEE Transactions on Knowledge and Data Engineering, 1–1. https://doi.org/10.1109/TKDE.2018.2882197

*Yao, R., Lin, G., Shen, C., Zhang, Y., & Shi, Q. (2018). Semantics-Aware Visual Object Tracking. IEEE Transactions on Circuits and Systems for Video Technology, 1–1. https://doi.org/10.1109/TCSVT.2018.2848358 *In Press

*You, M., Zhang, Y., Shen, C., & Zhang, X. (2018). An Extended Filtered Channel Framework for Pedestrian Detection. IEEE Transactions on Intelligent Transportation Systems, 19(5), 1640–1651. https://doi.org/10.1109/TITS.2018.2807199

Yu, L., Jacobson, A., & Milford, M. (2018). Rhythmic Representations: Learning Periodic Patterns for Scalable Place Recognition at a Sublinear Storage Cost. IEEE Robotics and Automation Letters, 3(2), 811–818. http://doi.org/10.1109/LRA.2018.2792144

*Zhang, L., Wang, P., Wei, W., Lu, H., Shen, C., van den Hengel, A., & Zhang, Y. (2018). Unsupervised Domain Adaptation Using Robust Class-Wise Matching. IEEE Transactions on Circuits and Systems for Video Technology. https://doi.org/10.1109/TCSVT.2018.2842206 *In Press

*Zhang, L., Wei, W., Zhang, Y., Shen, C., van den Hengel, A., & Shi, Q. (2018). Cluster Sparsity Field: An Internal Hyperspectral Imagery Prior for Reconstruction. International Journal of Computer Vision, 126(8), 797–821. https://doi.org/10.1007/s11263-018-1080-8

*Zhu, G., Porikli, F., & Li, H. (2018). Not All Negatives Are Equal: Learning to Track With Multiple Background Clusters. IEEE Transactions on Circuits and Systems for Video Technology, 28(2), 314–326. http://doi.org/10.1109/TCSVT.2016.2615518

*Zhuang, N., Yan, Y., Chen, S., Wang, H., & Shen, C. (2018). Multi-label learning based deep transfer neural network for facial attribute classification. Pattern Recognition, 80, 225–240. https://doi.org/10.1016/J.PATCOG.2018.03.018

CONFERENCE PAPERS (70)

*Adjigble, M., Marturi, N., Ortenzi, V., Rajasekaran, V., Corke, P., & Stolkin, R. (2018). Model-free and learning-free grasping by Local Contact Moment matching. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 2933–2940). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8594226

*Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., & Zhang, L. (2018). Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 6077–6086). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00636

*Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I., Gould, S., & van den Hengel, A. (2018). Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 3674–3683). IEEE. http://doi.org/10.1109/CVPR.2018.00387

*Bansal, S., Cosgun, A., Nakhaei, A., & Fujimura, K. (2018). Collaborative Planning for Mixed-Autonomy Lane Merging. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 4449–4455). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8594197

*Bateux, Q., Marchand, E., Leitner, J., Chaumette, F., & Corke, P. (2018). Training Deep Neural Networks for Visual Servoing. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1–8). Brisbane, Australia: IEEE. http://doi.org/10.1109/ICRA.2018.8461068

Cai Z., Chin TJ., Le H., Suter D. (2018) Deterministic Consensus Maximization with Biconvex Programming. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11216. Springer.

Camps, S. M., Houben, T., Fontanarosa, D., Edwards, C., Antico, M., Dunnhofer, M., Martens, E.G.H.J, Baeza, J.A., Vanneste, B.G.L., van Limbergen, E.J., de W., Peter, H.N., Verhaegen, F., & Carneiro, G. (2018). One-class Gaussian process regressor for quality assessment of transperineal ultrasound images. In International Conference on Medical Imaging with Deep Learning (MIDL). Amsterdam. Retrieved from https://eprints.qut.edu.au/120113/

*Cherian, A., Sra, S., Gould, S., & Hartley, R. (2018). Non-linear Temporal Subspace Representations for Activity Recognition. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 2197–2206). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00234

*Cherubini, A., Leitner, J., Ortenzi, V., & Corke, P. (2018). Towards vision-based manipulation of plastic materials. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 485–490). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8594108

Chin TJ., Cai Z., Neumann F. (2018) Robust Fitting in Computer Vision: Easy or Hard?. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11216. Springer.

*Cruz, R. S., Fernando, B., Cherian, A., & Gould, S. (2018). Neural Algebra of Classifiers. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 729–737). IEEE. https://doi.org/10.1109/WACV.2018.00085

*De Marco, S., Hua, M.-D., Mahony, R., & Hamel, T. (2018). Homography estimation of a moving planar scene from direct point correspondence. In 2018 IEEE Conference on Decision and Control (CDC) (pp. 565–570). Florida, United States: IEEE. http://doi.org/10.1109/CDC.2018.8619386

*Deng, C., Wu, Q., Wu, Q., Hu, F., Lyu, F., & Tan, M. (2018). Visual Grounding via Accumulated Attention. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 7746–7755). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00808

Deng R., Shen C., Liu S., Wang H., Liu X. (2018) Learning to Predict Crisp Boundaries. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11210. Springer.

*Dharmasiri, T., Spek, A., & Drummond, T. (2018). ENG: End-to-end Neural Geometry for Robust Depth and Pose Estimation using CNNs. Asian Conference on Computer Vision (ACCV). Retrieved from http://arxiv.org/abs/1807.05705

*Do, T.-T., Nguyen, A., & Reid, I. (2018). AffordanceNet: An End-to-End Deep Learning Approach for Object Affordance Detection. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1–5). Brisbane, Australia: IEEE. http://doi.org/10.1109/ICRA.2018.8460902

*Eriksson, A., Olsson, C., Kahl, F., & Chin, T.-J. (2018). Rotation Averaging and Strong Duality. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 127–135). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00021

*Felix R., Vijay Kumar B.G., Reid I., Carneiro G. (2018) Multi-modal Cycle-Consistent Generalized Zero-Shot Learning. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11210. Springer.

*Garg, S., Suenderhauf, N., & Milford, M. (2018). Don't Look Back: Robustifying Place Categorization for Viewpoint- and Condition-Invariant Place Recognition. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 3645–3652). Brisbane, Australia: IEEE. http://doi.org/10.1109/ICRA.2018.8461051

*He, T., Tian, Z., Huang, W., Shen, C., Qiao, Y., & Sun, C. (2018). An End-to-End TextSpotter with Explicit Alignment and Attention. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 5020–5029). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00527

*Jacobson, A., Zeng, F., Smith, D., Boswell, N., Peynot, T., & Milford, M. (2018). Semi-Supervised SLAM: Leveraging Low-Cost Sensors on Underground Autonomous Vehicles for Position Tracking. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 3970–3977). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8593750

*Jiang, S., Hartley, R., & Fernando, B. (2018). Kernel Support Vector Machines and Convolutional Neural Networks. In 2018 Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–7). Canberra, Australia: IEEE. http://doi.org/10.1109/DICTA.2018.8615840

Khan, S. H., Hayat, M., & Barnes, N. (2018). Adversarial Training of Variational Auto-Encoders for High Fidelity Image Generation. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 1312–1320). Lake Tahoe, United States: IEEE. https://doi.org/10.1109/WACV.2018.00148

*Kumar, S., Cherian, A., Dai, Y., & Li, H. (2018). Scalable Dense Non-rigid Structure-from-Motion: A Grassmannian Perspective. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 254–263). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00034

*Latif, Y., Garg, R., Milford, M., & Reid, I. (2018). Addressing Challenging Place Recognition Tasks Using Generative Adversarial Networks. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 2349–2355). Brisbane, Australia: IEEE. http://doi.org/10.1109/ICRA.2018.8461081

*Li K., Pham T., Zhan H., Reid I. (2018) Efficient Dense Point Cloud Object Reconstruction Using Deformation Vector Fields. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11216. Springer.

*Li, X., Li, H., Joo, H., Liu, Y., & Sheikh, Y. (2018). Structure from Recurrent Motion: From Rigidity to Recurrency. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 3032–3040). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00320

Liu Y., Dong W., Gong D., Zhang L., Shi Q. (2018) Deblurring Natural Image Using Super-Gaussian Fields. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11205. Springer.

*Lu X., Ma C., Ni B., Yang X., Reid I., Yang MH. (2018) Deep Regression Tracking with Shrinkage Loss. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11218. Springer.

*Lui, V., Geeves, J., Yii, W., & Drummond, T. (2018). Efficient Subpixel Refinement with Symbolic Linear Predictors. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 8165–8173). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00852

*Ma, C., Shen, C., Dick, A., Wu, Q., Wang, P., Hengel, A. van den, & Reid, I. (2018). Visual Question Answering with Memory-Augmented Networks. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 6975–6984). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00729

Maicas G., Bradley A.P., Nascimento J.C., Reid I., Carneiro G. (2018) Training Medical Image Analysis Systems like Radiologists. In: Frangi A., Schnabel J., Davatzikos C., Alberola-López C., Fichtinger G. (eds) Medical Image Computing and Computer Assisted Intervention – MICCAI 2018. MICCAI 2018. Lecture Notes in Computer Science, vol 11070. Springer.

*Marmol, A., Corke, P., & Peynot, T. (2018). ArthroSLAM: Multi-Sensor Robust Visual Localization for Minimally Invasive Orthopedic Surgery. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 3882–3889). Madrid, Spain: IEEE. https://doi.org/10.1109/IROS.2018.8593501

*McFadyen, A., Dayoub, F., Martin, S., Ford, J., & Corke, P. (2018). Assisted Control for Semi-Autonomous Power Infrastructure Inspection Using Aerial Vehicles. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 5719–5726). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8593529

*Meyer, B. J., Harwood, B., & Drummond, T. (2018). Deep Metric Learning and Image Classification with Nearest Neighbour Gaussian Kernels. In IEEE International Conference on Image Processing (ICIP) (pp. 151–155). Athens, Greece: IEEE. http://doi.org/10.1109/ICIP.2018.8451297

*Milan, A., Pham, T., Vijay, K., Morrison, D., Tow, A. W., Liu, L., Erskine, J., Grinover, R., Gurman, A., Hunn, T., Kelly-Boxall, K., Lee, D., McTaggart, M., Rallos, G., Razjigaev, A., Rowntree, T., Shen, T., Smith, R., Wade-McCue, S., Zhuang, Z., Lehnert, C., Lin, G., Reid, I., Corke, P., and Leitner, J. (2018). Semantic Segmentation from Limited Training Data. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1908–1915). Brisbane: IEEE. http://doi.org/10.1109/ICRA.2018.8461082

*Miller, D., Nicholson, L., Dayoub, F., & Sunderhauf, N. (2018). Dropout Sampling for Robust Object Detection in Open-Set Conditions. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1–7). Brisbane: IEEE. http://doi.org/10.1109/ICRA.2018.8460700

*Morrison, D., Tow, A. W., McTaggart, M., Smith, R., Kelly-Boxall, N., Wade-McCue, S., Erskine, J., Grinover, R., Gurman, A., Hunn, T., Lee, D., Milan, A., Pham, T., Rallos, G., Razjigaev, A., Rowntree, T., Kumar, V., Zhuang, Z., Lehnert, C., Reid, I., Corke, P., and Leitner, J. (2018). Cartman: The low-cost Cartesian Manipulator that won the Amazon Robotics Challenge. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 7757–7764). Brisbane: IEEE. http://doi.org/10.1109/ICRA.2018.8463191

*O'Brien, S. G. P., Trumpf, J., Ila, V., & Mahony, R. (2018). A Geometric Observer for Scene Reconstruction Using Plenoptic Cameras. In 2018 IEEE Conference on Decision and Control (CDC) (pp. 557–564). Florida, United States: IEEE. http://doi.org/10.1109/CDC.2018.8618954

*O'brien, S., Trumpf, J., Ila, V., & Mahony, R. (2018). Calibrating Light-Field Cameras Using Plenoptic Disc Features. In 2018 International Conference on 3D Vision (3DV) (pp. 286–294). Verona, Italy: IEEE. http://doi.org/10.1109/3DV.2018.00041

*Pham T., Vijay Kumar B.G., Do TT., Carneiro G., Reid I. (2018) Bayesian Semantic Instance Segmentation in Open Set World. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11214. Springer.

*Pham, T. T., Do, T.-T., Sunderhauf, N., & Reid, I. (2018). SceneCut: Joint Geometric and Object Segmentation for Indoor Scenes. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1–9). Brisbane: IEEE. http://doi.org/10.1109/ICRA.2018.8461108

*Rodriguez C., Fernando B., Li H. (2018) Action Anticipation by Predicting Future Dynamic Images. In: Leal-Taixé L., Roth S. (eds) Computer Vision – ECCV 2018 Workshops. ECCV 2018. Lecture Notes in Computer Science, vol 11131. Springer.

*Rubino, C., Del Bue, A., & Chin, T.-J. (2018). Practical Motion Segmentation for Urban Street View Scenes. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1879–1886). Brisbane, Australia: IEEE. http://doi.org/10.1109/ICRA.2018.8460993

*Shen, T., Lin, G., Shen, C., & Reid, I. (2018). Bootstrapping the Performance of Webly Supervised Semantic Segmentation. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 1363–1371). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00148

*Shiri, F., Porikli, F., Hartley, R., & Koniusz, P. (2018). Identity-Preserving Face Recovery from Portraits. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 102–111). IEEE. https://doi.org/10.1109/WACV.2018.00018

*Song, Y., Ma, C., Wu, X., Gong, L., Bao, L., Zuo, W., Shen, C., Lau, Rynson W.H., & Yang, M.-H. (2018). VITAL: Visual Tracking via Adversarial Learning. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 8990–8999). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00937

*Spek, A., Dharmasiri, T., & Drummond, T. (2018). CReaM: Condensed Real-time Models for Depth Prediction using Convolutional Neural Networks. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 540–547). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8594243

*Stevens, J.-L., & Mahony, R. (2018). Vision Based Forward Sensitive Reactive Control for a Quadrotor VTOL. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 5232–5238). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8593606

*Talbot, B., Garg, S., & Milford, M. (2018). OpenSeqSLAM2.0: An Open Source Toolbox for Visual Place Recognition Under Changing Conditions. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 7758–7765). Madrid, Spain: IEEE. http://doi.org/10.1109/IROS.2018.8593761

Teney, D., Anderson, P., He, X., & Hengel, A. van den. (2018). Tips and Tricks for Visual Question Answering: Learnings from the 2017 Challenge. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 4223–4232). IEEE. http://doi.org/10.1109/CVPR.2018.00444

*Wang J., Cherian A. (2018) Learning Discriminative Video Representations Using Adversarial Perturbations. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11208. Springer.

*Wang, J., Cherian, A., Porikli, F., & Gould, S. (2018). Video Representation Learning Using Discriminative Pooling. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 1149–1158). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00126

*Weerasekera, C. S., Dharmasiri, T., Garg, R., Drummond, T., & Reid, I. (2018). Just-in-Time Reconstruction: Inpainting Sparse Maps Using Single View Depth Predictors as Priors. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1–9). Brisbane, Australia: IEEE. http://doi.org/10.1109/ICRA.2018.8460549

*Wu, Q., Wang, P., Shen, C., Reid, I., & Hengel, A. van den. (2018). Are You Talking to Me? Reasoned Visual Dialog Generation Through Adversarial Learning. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 6106–6115). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00639

Yang, J., Gong, D., Liu, L., & Shi, Q. (2018). Seeing Deeply and Bidirectionally: A Deep Learning Approach for Single Image Reflection Removal. Retrieved from http://openaccess.thecvf.com/content_ECCV_2018/papers/Jie_Yang_Seeing_Deeply_and_ECCV_2018_paper.pdf

*Yu X., Fernando B., Ghanem B., Porikli F., Hartley R. (2018) Face Super-Resolution Guided by Facial Component Heatmaps. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11213. Springer.

*Yu, X., Fernando, B., Hartley, R., & Porikli, F. (2018). Super-Resolving Very Low-Resolution Face Images with Supplementary Attributes. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 908–917). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00101

*Zhan, H., Garg, R., Weerasekera, C. S., Li, K., Agarwal, H., & Reid, I. M. (2018). Unsupervised Learning of Monocular Depth Estimation and Visual Odometry with Deep Feature Reconstruction. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 340–349). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00043

*Zhang J., Wu Q., Shen C., Zhang J., Lu J., van den Hengel A. (2018) Goal-Oriented Visual Question Generation via Intermediate Rewards. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11209. Springer.

*Zhang, J., Ila, V., & Kneip, L. (2018). Robust Visual Odometry in Underwater Environment. In 2018 OCEANS - MTS/IEEE Kobe Techno-Oceans (OTO) (pp. 1–9). Kobe, Japan: IEEE. http://doi.org/10.1109/OCEANSKOBE.2018.8559452

*Zhang, J., Zhang, T., Daf, Y., Harandi, M., & Hartley, R. (2018). Deep Unsupervised Saliency Detection: A Multiple Noisy Labeling Perspective. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 9029–9038). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00941

*Zhang, Q., Chin, T.-J., & Le, H. M. (2018). A Fast Resection-Intersection Method for the Known Rotation Problem. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 3012–3021). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00318

*Zhong Y., Dai Y., Li H. (2018) Stereo Computation for a Single Mixture Image. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11213. Springer.

*Zhong Y., Li H., Dai Y. (2018) Open-World Stereo Video Matching with Deep RNN. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11206. Springer.

*Zhong, Y., Dai, Y., & Li, H. (2018). 3D Geometry-Aware Semantic Labeling of Outdoor Street Scenes. In 2018 24th International Conference on Pattern Recognition (ICPR) (pp. 2343–2349). IEEE. http://doi.org/10.1109/ICPR.2018.8545378

*Zhou Y., Gallego G., Rebecq H., Kneip L., Li H., Scaramuzza D. (2018) Semi-dense 3D Reconstruction with a Stereo Event Camera. In: Ferrari V., Hebert M., Sminchisescu C., Weiss Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science, vol 11205. Springer.

*Zhuang, B., Shen, C., Tan, M., Liu, L., & Reid, I. (2018). Towards Effective Low-Bitwidth Convolutional Neural Networks. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 7920–7928). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00826

*Zhuang, B., Wu, Q., Shen, C., Reid, I., & Hengel, A. van den. (2018). Parallel Attention: A Unified Framework for Visual Object Discovery Through Dialogs and Queries. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 4252–4261). Salt Lake City, United States: IEEE. http://doi.org/10.1109/CVPR.2018.00447

*Zuo, Y., Avraham, G., & Drummond, T. (2018). Traversing Latent Space using Decision Ferns. Asian Conference on Computer Vision (ACCV). Retrieved from http://arxiv.org/abs/1812.02636

SUBMITTED PAPERS (71)

*Abbasnejad, E., Wu, Q., Abbasnejad, I., Shi, J., & Van Den Hengel, A. (2018). An Active Information Seeking Model for Goal-oriented Vision-and-Language Tasks. Retrieved from https://arxiv.org/pdf/1812.06398

*Abbasnejad, E., Wu, Q., Shi, J., & Van Den Hengel, A. (2018). What's to know? Uncertainty as a Guide to Asking Goal-oriented Questions. Retrieved from https://arxiv.org/pdf/1812.06401

*Ahn, H. S., Dayoub, F., Popovic, M., MacDonald, B., Siegwart, R., & Sa, I. (2018). An Overview of Perception Methods for Horticultural Robots: From Pollination to Harvest. Retrieved from http://arxiv.org/abs/1807.03124

*Aliakbarian, M. S., Saleh, F. S., Salzmann, M., Fernando, B., Petersson, L., & Andersson, L. (2018). VIENA 2 : A Driving Anticipation Dataset. Retrieved from https://arxiv.org/pdf/1810.09044

*Baktashmotlagh, M., Faraki, M., Drummond, T., & Salzmann, M. (2018). Learning Factorized Representations for Open-set Domain Adaptation. Retrieved from http://arxiv.org/abs/1805.12277

*Bruce, J., Sünderhauf, N., Mirowski, P., Hadsell, R., & Milford, M. (2018). Learning Deployable Navigation Policies at Kilometer Scale from a Single Traversal. Retrieved from http://arxiv.org/abs/1807.05211

Cao, J., Guo, Y., Wu, Q., Shen, C., Huang, J., & Tan, M. (2018). Adversarial Learning with Local Coordinate Coding. Retrieved from http://arxiv.org/abs/1806.04895

Cao, Y., Zhao, T., Xian, K., Shen, C., & Cao, Z. (2018). Monocular Depth Estimation with Augmented Ordinal Depth Relationships. Retrieved from http://arxiv.org/abs/1806.00585

*Collins, J., Howard, D., & Leitner, J. (2018). Quantifying the Reality Gap in Robotic Manipulation Tasks. Retrieved from https://arxiv.org/abs/1811.01484

*Do, T.-T., Hoang, T., Tan, D.-K. Le, & Cheung, N.-M. (2018). From Selective Deep Convolutional Features to Compact Binary Representations for Image Retrieval. Retrieved from http://arxiv.org/abs/1802.02899

*Doan, A.-D., Do, T.-T., Latif, Y., Chin, T.-J., & Reid, I. (2018). Practical Visual Localization for Autonomous Driving: Why Not Filter? Retrieved from https://arxiv.org/pdf/1811.08063

*Doan, A.-D., Jawaid, A. M., Do, T.-T., & Chin, T.-J. (2018). G2D: from GTA to Data. Retrieved from http://arxiv.org/abs/1806.07381

Gale, W., Oakden-Rayner, L., Carneiro, G., Bradley, A. P., & Palmer, L. J. (2018). Producing radiologist-quality reports for interpretable artificial intelligence. Retrieved from http://arxiv.org/abs/1806.00340

*Garg, S., Suenderhauf, N., & Milford, M. (2018). LoST? Appearance-Invariant Place Recognition for Opposite Viewpoints using Visual Semantics. Retrieved from http://arxiv.org/abs/1804.05526

Gong, D., Zhang, Z., Shi, Q., Hengel, A. van den, Shen, C., & Zhang, Y. (2018). Learning an Optimizer for Image Deconvolution. Retrieved from https://arxiv.org/abs/1804.03368

Jack, D., Pontes, J. K., Sridharan, S., Fookes, C., Shirazi, S., Maire, F., & Eriksson, A. (2018). Learning Free-Form Deformations for 3D Object Reconstruction. Retrieved from http://arxiv.org/abs/1803.10932

Jiang, S., Lu, X., Lei, Y., & Liu, L. (2018). Mask-aware networks for crowd counting. Retrieved from http://arxiv.org/abs/1901.00039

*Hall, D., Dayoub, F., Skinner, J., Corke, P., Carneiro, G., & Sünderhauf, N. (2018). Probability-based Detection Quality (PDQ): A Probabilistic Approach to Detection Evaluation. Retrieved from http://arxiv.org/abs/1811.10800

*Hausler, S., Jacobson, A., & Milford, M. (2018). Feature Map Filtering: Improving Visual Place Recognition with Convolutional Calibration. Retrieved from http://arxiv.org/abs/1810.12465

*Henein, M., Kennedy, G., Ila, V., & Mahony, R. (2018). Simultaneous Localization and Mapping with Dynamic Rigid Objects. Retrieved from http://arxiv.org/abs/1805.03800

*Hoang, T., Do, T.-T., Le-Tan, D.-K., & Cheung, N.-M. (2018). Simultaneous Compression and Quantization: A Joint Approach for Efficient Unsupervised Hashing. Retrieved from http://arxiv.org/abs/1802.06645

*Hosseinzadeh, M., Latif, Y., Pham, T., Suenderhauf, N., & Reid, I. (2018). Towards Semantic SLAM: Points, Planes and Objects. Retrieved from http://arxiv.org/abs/1804.09111

*Hosseinzadeh, M., Li, K., Latif, Y., & Reid, I. (2018). Real-Time Monocular Object-Model Aware Sparse SLAM. Retrieved from https://arxiv.org/pdf/1809.09149

*Jablonsky, N., Milford, M., & Sünderhauf, N. (2018). An Orientation Factor for Object-Oriented SLAM. Retrieved from http://arxiv.org/abs/1809.06977

Khaliq, A., Ehsan, S., Milford, M., & McDonald-Maier, K. (2018). A Holistic Visual Place Recognition Approach using Lightweight CNNs for Severe ViewPoint and Appearance Changes. Retrieved from http://arxiv.org/abs/1811.03032

*Korthals, T., Leitner, J., & Rückert, U. (2018). Coordinated Heterogeneous Distributed Perception based on Latent Space Representation. Retrieved from https://arxiv.org/abs/1809.04558

Le, H., Eriksson, A., Do, T.-T., & Milford, M. (2018). A Binary Optimization Approach for Constrained K-Means Clustering. Retrieved from http://arxiv.org/abs/1810.10134

*Le, H., & Milford, M. (2018). Large scale visual place recognition with sub-linear storage growth. Retrieved from http://arxiv.org/abs/1810.09660

*Lee, R., Mou, S., Dasagi, V., Bruce, J., Leitner, J., & Sünderhauf, N. (2018). Zero-shot Sim-to-Real Transfer with Modular Priors. Retrieved from http://arxiv.org/abs/1809.07480

*Lehnert, C., Tsai, D., Eriksson, A., & McCool, C. (2018). 3D Move to See: Multi-perspective visual servoing for improving object views with semantic segmentation. Retrieved from http://arxiv.org/abs/1809.07896

*Li, H., Wang, P., Shen, C., & Van Den Hengel, A. (2018). Visual Question Answering as Reading Comprehension. Retrieved from https://arxiv.org/pdf/1811.11903

Li, R., Xian, K., Shen, C., Cao, Z., Lu, H., & Hang, L. (2018). Deep attention-based classification network for robust depth prediction. Retrieved from https://arxiv.org/abs/1807.03959

*Li, H., Wang, P., Shen, C., & Zhang, G. (2018). Show, Attend and Read: A Simple and Strong Baseline for Irregular Text Recognition. Retrieved from https://arxiv.org/pdf/1811.00751

*Li, K., Garg, R., Cai, M., & Reid, I. (2018). Optimizable Object Reconstruction from a Single View. Retrieved from https://arxiv.org/pdf/1811.11921

*Liu, L., Li, H., & Dai, Y. (2018). Deep Stochastic Attraction and Repulsion Embedding for Image Based Localization. Retrieved from https://arxiv.org/pdf/1808.08779

Liu, W., Xu, W., Chen, X., Huang, X., Shen, C., & Yang, J. (2018). Edge-Preserving Piecewise Linear Image Smoothing Using Piecewise Constant Filters. Retrieved from http://arxiv.org/abs/1801.06928

*Liu, Z., Lin, G., Ling Goh, W., Liu, F., Shen, C., & Yang, X. (2018). Correlation Propagation Networks for Scene Text Detection. Retrieved from https://arxiv.org/pdf/1810.00304

*Lu, Y., Harandi, M., Hartley, R., & Pascanu, R. (2018). Block Mean Approximation for Efficient Second Order Optimization. Retrieved from http://arxiv.org/abs/1804.05484

Maicas, G., Bradley, A. P., Nascimento, J. C., Reid, I., & Carneiro, G. (2018). Pre and Post-hoc Diagnosis and Interpretation of Malignancy from Breast DCE-MRI. Retrieved from https://arxiv.org/pdf/1809.09404

*Miller, D., Dayoub, F., Milford, M., & Sünderhauf, N. (2018). Evaluating Merging Strategies for Sampling-based Uncertainty Techniques in Object Detection. Retrieved from http://arxiv.org/abs/1809.06006

*Morrison, D., Corke, P., & Leitner, J. (2018). Closing the Loop for Robotic Grasping: A Real-time, Generative Grasp Synthesis Approach. Retrieved from http://arxiv.org/abs/1804.05172

*Morrison, D., Corke, P., & Leitner, J. (2018). Multi-View Picking: Next-best-view Reaching for Improved Grasping in Clutter. Retrieved from http://arxiv.org/abs/1809.08564

*Nekrasov, V., Chen, H., Shen, C., & Reid, I. (2018). Fast Neural Architecture Search of Compact Semantic Segmentation Models via Auxiliary Cells *. Retrieved from https://arxiv.org/pdf/1810.10804

*Nekrasov, V., Dharmasiri, T., Spek, A., Drummond, T., Shen, C., & Reid, I. (2018). Real-Time Joint Semantic Segmentation and Depth Estimation Using Asymmetric Annotations. Retrieved from http://arxiv.org/abs/1809.04766

*Nekrasov, V., Shen, C., & Reid, I. (2018). Diagnostics in Semantic Segmentation. Retrieved from https://arxiv.org/pdf/1809.10328

*Nekrasov, V., Shen, C., & Reid, I. (2018). Light-Weight RefineNet for Real-Time Semantic Segmentation. Retrieved from http://arxiv.org/abs/1810.03272

*Nguyen, A., Do, T.-T., Reid, I., Caldwell, D. G., & Tsagarakis, N. G. (2018). Object Captioning and Retrieval with Natural Language. Retrieved from https://arxiv.org/abs/1803.06152

Nguyen, H. Van, Rezatofighi, S. H., Vo, B.-N., & Ranasinghe, D. C. (2018). Online UAV Path Planning for Joint Detection and Tracking of Multiple Radio-tagged Objects. Retrieved from https://arxiv.org/pdf/1808.04445

*Pan, L., Hartley, R., Liu, M., & Dai, Y. (2018). Phase-only Image Based Kernel Estimation for Single-image Blind Deblurring. Retrieved from https://arxiv.org/pdf/1811.10185

*Rezatofighi, S. H., Kaskman, R., Motlagh, F. T., Shi, Q., Cremers, D., Leal-Taixé, L., & Reid, I. (2018). Deep Perm-Set Net: Learn to predict sets with unknown permutation and cardinality using deep neural networks. Retrieved from https://arxiv.org/abs/1805.00613

*Rezazadegan, F., Shirazi, S., Baktashmotlagh, M., & Davis, L. S. (2018). On Encoding Temporal Evolution for Real-time Action Prediction. Retrieved from http://arxiv.org/abs/1709.07894

*Scheerlinck, C., Barnes, N., & Mahony, R. (2018). Continuous-time Intensity Estimation Using Event Cameras. Retrieved from https://arxiv.org/pdf/1811.00386

Snaauw, G., Gong, D., Maicas, G., Hengel, A. van den, Niessen, W. J., Verjans, J., & Carneiro, G. (2018). End-to-End Diagnosis and Segmentation Learning from Cardiac Magnetic Resonance Imaging. Retrieved from https://arxiv.org/abs/1810.10117

*Song, X., Wang, P., Zhou, D., Zhu, R., Guan, C., Dai, Y., Su, H., Li, H., & Yang, R. (2018). ApolloCar3D: A Large 3D Car Instance Understanding Benchmark for Autonomous Driving. Retrieved from http://arxiv.org/abs/1811.12222

*Stoffregen, T., & Kleeman, L. (2018). Simultaneous Optical Flow and Segmentation (SOFAS) using Dynamic Vision Sensor. Retrieved from http://arxiv.org/abs/1805.12326

*Suddrey, G., Jacobson, A., & Ward, B. (2018). Enabling a Pepper Robot to provide Automated and Interactive Tours of a Robotics Laboratory. Retrieved from http://arxiv.org/abs/1804.03288

Tursun, O., Denman, S., Sivipalan, S., Sridharan, S., Fookes, C., & Mau, S. (2018). Component-based Attention for Large-scale Trademark Retrieval. Retrieved from http://arxiv.org/abs/1811.02746

*Wang, P., Wu, Q., Cao, J., Shen, C., Gao, L., & Van Den Hengel, A. (2018). Neighbourhood Watch: Referring Expression Comprehension via Language-guided Graph Attention Networks. Retrieved from https://arxiv.org/pdf/1812.04794

*Zaffar, M., Ehsan, S., Milford, M., & Maier, K. M. (2018). Memorable Maps: A Framework for Re-defining Places in Visual Place Recognition. Retrieved from https://arxiv.org/pdf/1811.03529

*Zhang, H., Li, Y., Wang, P., Liu, Y., & Shen, C. (2018). RGB-D Based Action Recognition with Light-weight 3D Convolutional Networks. Retrieved from https://arxiv.org/pdf/1811.09908

*Zhang, L., Wang, P., Liu, L., Shen, C., Wei, W., Zhang, Y., & Van Den Hengel, A. (2018). Towards Effective Deep Embedding for Zero-Shot Learning. Retrieved from https://arxiv.org/pdf/1808.10075

*Zhang, T., Ji, P., Harandi, M., Hartley, R., & Reid, I. (2018). Scalable Deep k-Subspace Clustering. Retrieved from https://arxiv.org/abs/1811.01045

*Zhang, T., Lin, G., Cai, J., Shen, T., Shen, C., & Kot, A. C. (2018). Decoupled Spatial Neural Attention for Weakly Supervised Semantic Segmentation. Retrieved from http://arxiv.org/abs/1803.02563

*Zhuang, B., Shen, C., & Reid, I. (2018). Training Compact Neural Networks with Binary Weights and Low Precision Activations. Retrieved from https://arxiv.org/pdf/1808.02631

*Zhuang, B., Shen, C., Tan, M., Liu, L., & Reid, I. (2018). Structured Binary Neural Networks for Accurate Image Classification and Semantic Segmentation. Retrieved from https://arxiv.org/abs/1811.10413

*Zhuang, Z., Tan, M., Zhuang, B., Liu, J., Guo, Y., Wu, Q., Huang, J., & Zhu, J. (2018). Discrimination-aware Channel Pruning for Deep Neural Networks. Retrieved from https://arxiv.org/pdf/1810.11809

*Zuo, Y., Avraham, G., & Drummond, T. (2018). Generative Adversarial Forests for Better Conditioned Adversarial Learning. Retrieved from http://arxiv.org/abs/1805.05185

Zhang, L., Wang, P., Shen, C., Liu, L., Wei, W., Zhang, Y., & Hengel, A. van den. (2018). Adaptive Importance Learning for Improving Lightweight Image Super-resolution Network. Retrieved from http://arxiv.org/abs/1806.01576

Zhang, P., Liu, W., Lu, H., & Shen, C. (2018). Salient Object Detection by Lossless Feature Reflection. Retrieved from https://arxiv.org/pdf/1802.06527

Zhang, P., Lu, H., & Shen, C. (2018). HyperFusion-Net: Densely Reflective Fusion for Salient Object Detection. Retrieved from https://arxiv.org/abs/1804.05142

Zhang, P., Wang, L., Wang, D., Lu, H., & Shen, C. (2018). Agile Amulet: Real-Time Salient Object Detection with Contextual Attention. Retrieved from http://arxiv.org/abs/1802.06960

The Story of our Logo

Our logo represents the reunification of robotics and computer vision. It symbolises how robots might see in the future and recognises the importance of vision in the evolution of life on Earth.

AUSTRALIAN CENTRE FOR ROBOTIC VISION

540 million years ago, during the critical time period known as the Cambrian, the sense of vision, with its advanced and complex neurological network, was at the center of the Darwinian struggle for survival. Vision was a principal driver of evolution, providing animals with a map of their external world and concurrently invoking self-awareness - the recognition that the 'self' viewing the world was also separate from it.

Vision also allowed animals to recognise similar forms and to associate with them, producing the inherent survival advantages involved in being part of a group.

Eventually, after 540 million years, humans and the human eye evolved. Humans then developed the technology to capture images using cameras, which mimic the human eye.

As the purpose of the Australian Centre for Robotic Vision is to give robots the gift of sight, our logo incorporates the most important elements of the eye.

Our Centre sits at the aperture (or opening) that allows light into the eye.

The silver outer circle represents the sclera, the protective, outer layer of the eye.

The blue circles represent the iris and the pupil, which control the amount of light entering the eye's natural crystalline lens. This clear, flexible structure works like the lens in a camera, shortening and lengthening its width in order to focus light rays.

The red shape represents a cross-section through an eye and symbolises the retina, where light rays come to a focusing point.

Embedded in the retina are millions of light sensitive cells, responsible for capturing light rays and processing them into impulses that are sent to the optic nerve. In a robot's eye these are digital sensors.

Just as vision played a major role in the evolution of life on Earth it can also spark the intelligence required for robots to be able to understand their environment, to make decisions and to perform useful tasks in the complex, unstructured and dynamically changing environments in which we live and work. Just as the minds of animals developed around the need to support a sense of vision, so to will the capabilities of robots.

Robotic Vision Partner Organisations

QUT

The Australian National University

The University of Adelaide

Monash University

Imperial College London

Data 61 CSIRO

University of Oxford

Inria — Inventeurs du monde numérique

Georgia Tech

ETH zürich

Australian Government — Australian Research Council

Robotic Vision Industry Partners

Intel RealSense Technology

MathWorks

Australian Centre for Robotic Vision Head Office Contact Details

Queensland University of Technology
2 George Street
Brisbane QLD 4000 AUSTRALIA

p +61 7 3138 7549
e info@roboticvision.org
w roboticvision.org

Design by Studio 55