AUSTRALIAN CENTRE FOR ROBOTIC VISION

ANNUAL REPORT **2016**

creating robots that see...

AN ARC CENTRE OF EXCELLENCE

Australian Government
Australian Research Council

Australian Centre for Robotic Vision at our third annual symposium RoboVis2016, held in Lorne, Victoria.

Our Vision, Mission and Values

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*
Empower	*We energise, motivate and support our people to be knowledge leaders*
Collaborate	*We work together and partner to solve grand challenges*
Impact	*We make a difference, by applying our transformational research and turning our ideas into reality*

Cover image caption: Our front cover shows a quadrotor aerial vehicle that is used in the Centre's Vision and Action research program to study vision based navigation in complex unstructured 3D environments. The vehicle uses an omnidirectional catadioptric (mirror-based) camera mounted above the central avionics stack (image courtesy of the Australian National University).

Table of Contents

About the Australian Centre for Robotic Vision

Who we are

We are an Australian research centre that leads the world in the new discipline of robotic vision, applying computer vision to robotics. While robotics is about machines that interact with the physical world, computer vision is about analysing and understanding the world through images. Robotic vision expands the capabilities of robots, allowing them to see and understand the world in which they are working. We believe it is the key technology that will allow robotics to change the way we live and work. Already we are looking to apply our technologies to solve real challenges in the monitoring and protection of the natural and built environments, the provision of healthcare in hospitals and in the home, sustainable food production, and efficiently harnessing our natural resources.

Formed in 2014, our Centre is an unincorporated collaborative venture between interdisciplinary researchers from leading Australian and international research organisations; Queensland University of Technology (QUT), The University of Adelaide (Adelaide), The Australian National University (ANU), and Monash University, CSIRO (Data61), the French Institute for Research in Computer Science and Automation (INRIA) Rennes, Georgia Institute of Technology, Imperial College London, the Swiss Federal Institute of Technology (ETH) Zürich, and the University of Oxford. We have been given carriage of $25.6m funding over seven years to form the largest group of its kind in the world, generating: internationally impactful science; new research capacity; highly skilled people; an increased reputation for Australian research; intellectual property; technologies that will transform important Australian industries, and solutions to some of the hard challenges facing Australia and the rest of the world.

Centres of Excellence

We receive investment of $19m in public funding from the Australian Research Council's Centres of Excellence program. The Australian Research Council (ARC) is a statutory agency responsible for funding excellent research and research training, and manages the National Competitive Grants Program (NCGP), which funds basic and applied research across all disciplines. ARC Centres of Excellence are prestigious foci of expertise through which high-quality researchers collaboratively maintain and develop Australia's international standing in research areas of national priority (see National Benefit p. 42). ARC Centres of Excellence involve significant collaboration allowing outstanding research to be supported by the complementary research resources of universities, publicly-funded research organisations, other research bodies, governments and industry.

World map highlighting partner locations

About this Report

Report Description

Our report covers the activities of the ARC Centre of Excellence for Robotic Vision for the 2016 calendar year. Activities encompass research, training, outreach, industry engagement, operations and finance. Our reporting period aligns with the requirements of the Australian Research Council, our primary source of funding, and the report forms part of our official reporting (and accounting) requirements.

Aims of the Report

Our Centre has been given carriage of $19m of public funds from the Australian Research Council, matched with a further $6.6m in funding from partner organisations, the Australian universities: Queensland University of Technology (QUT), The Australian National University (ANU), Monash University (Monash), and The University of Adelaide (Adelaide). In return for this funding our Centre has committed to an ambitious research program and to achieving a range of key performance indicators that will be described in this report, covering: research findings, research training and professional education, international, national and regional links and networks, end-user links, organisational support, governance, and national benefit. Our report outlines our vision as well as highlighting our achievements and providing an overview of our operations for 2016. We also identify and map out the various stakeholders that form part of, or have an interest in, the activities of the Centre, and outline our strategy for engaging with our stakeholders, researchers, and the wider community.

Anticipated Readership

The primary audiences for this report are our funders and stakeholders, and we also hope it will be of interest to the broader community in both Australia and overseas. Subject matter has been selected in line with our vision and strategic plan and in accordance with the expectations of the Australian Research Council. Unless otherwise stated, the use of the words 'we, 'us, 'our' and 'the Centre' refers to the ARC Centre of Excellence for Robotic Vision, known as the Australian Centre for Robotic Vision. You will also find this report, and various other Centre publications on our website at www.roboticvision.org. To provide feedback on this report please visit www.roboticvision.org/feedback

Vision-enabled wheeled robot used for beach experiments at our Robotic Vision Summer School. Image courtesy of Michael Milford.

History

The technologies of robotics and computer vision are each over 50 years old. From the 1970s, Artificial Intelligence (AI) labs around the world investigated both robotics and computer vision. Robots collected images of their surroundings and computer vision provided the understanding required for robots to physically interact with the world around them. Early robots were primitive and slow and computer vision was limited by the speed of computers of the day — particularly for metric problems like understanding the geometry of the world. In the 1990s affordable laser rangefinders provided the metric information required to enable robots to navigate in the world. Since that time, laser-based perception has displaced vision as the sensor of choice for roboticists while computer vision became dominated by interpretation of medical images and social/media photographs, not images captured by robots. Our Centre is rekindling the early partnership between robotics and vision, using advanced modern computers and algorithms with the aim to establish vision as a powerful and cost-effective sensor for robots.

Why is vision so important for robots? For the same reason that the sense of vision is so important for living beings. Vision has been a principal driver of evolution, providing animals with a map of their external world and an awareness of how they fit in the world. This concurrently invokes self-awareness - the recognition that the "self" viewing the world is also separate from it. Robots need this capability to understand their environment, to make decisions and to perform useful tasks. It is not until robots can harness visual information that they can work safely and effectively alongside humans in the complex, unstructured and dynamically changing environments in which we live and work.

In July 2012 the ARC invited applications to form Centres of Excellence. Professors Peter Corke (QUT) and Robert Mahony (ANU) identified the opportunity to finally bring the robotics and computer vision disciplines together, working on the common challenge of creating robots that see. A team of researchers (see People p. 57) was assembled and the Australian Centre for Robotic Vision was approved at the end of 2013. We commenced operations on 1st January 2014 and became fully operational on 21st July 2014 following the signing of agreements between all collaborating organisations. Our Centre was officially launched in early 2015 and now has research strength of more than 100 people, including 13 Chief Investigators, 16 Research Fellows, 44 PhD researchers, 9 Honours Students, 18 Associate Investigators, 6 Partner Investigators, 9 Professional Staff, 2 Research Engineers and 9 Research Affiliates.

Centre Performance

Both robotics and computer vision are seen as powerful disruptive agents of change. We believe that vision is the key to unleashing the full potential of robots.

The robotic tipping point

Advances in sensors and actuators, computation and machine-to-machine communication contribute to the development of robots that are able to perceive the world around them. This is the step required to make robots truly useful and safe to humans. The hype around robotics and computer vision has been building for more than 50 years. To date, technologies have fallen short of our imagination — advanced machines in movies far surpass the abilities of modern robots. By applying software algorithms for computer vision, robots are able to analyse and extract a range of useful information from images they capture using cameras. We are then able to create and deploy robots that can adapt and understand their surroundings, making it possible for robots to finally develop some of the abilities we have previously only imagined. This makes it possible to deploy robots safely alongside people. Advanced robotics is predicted to be one of the top nine disruptive technologies (Gartner) and yet to date, technologies have fallen short of our imagination. Robotics and computer vision have teetered on the cusp of achieving their potential for more than 50 years. We aim to undertake the science, and create the technologies, that will allow this next generation of robots to see: using cameras and advanced computer vision techniques to understand their environment, adapt to change, and learn to cooperate. When we are able to develop truly robotic vision solutions to problems, then robotics will reach the tipping point and will start to become more commonplace in our everyday lives.

What is a robot?

Robots are autonomous machines that can move within their physical environment and manipulate objects. The range of tasks that a robot can perform is limited by the robot's knowledge of its surroundings. How do robots learn about their environment?

As humans we use our eyes and brains to visually sense and understand the world around us. Visual perception is widely believed to hold an evolutionary advantage that led to the Cambrian explosion of life on planet Earth more than 500 million years ago. Computer vision aims to give machines and computers similar, if not better, visual perception than humans.

"Robots are autonomous machines that can move within their physical environment and manipulate objects. The range of tasks that a robot can perform is limited by the robot's knowledge of its surroundings."

In our Centre, we can envisage a world where robots:

- understand tasks and can learn how to adapt to changes
- watch a person perform a task once and then repeat it
- recognise people, understand what they are doing and help with what they need next
- fly underneath large infrastructure assets such as bridges, autonomously inspect them and monitor changes
- have superhuman hand/eye coordination for handling goods

What is a robot?

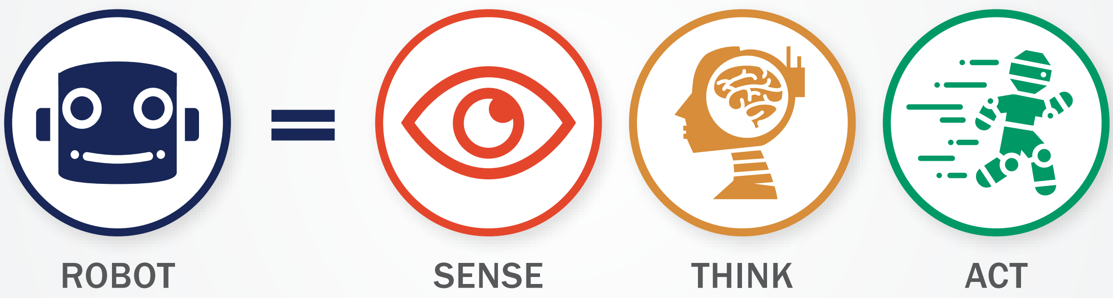

Centre Strategy

Culture

A vibrant, high energy, future-focused, collaborative robotic vision community developing knowledge leaders for both industry and academia

Science

Leading the world in transformational research in the new field of robotic vision

Integrate

Bringing the disciplines of robotics and computer vision together to create new robotic vision technologies

Engage

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

Transform

Applying robotic vision technologies to solve innovation challenges and transform the world

- reduce the cost of healthcare with low cost robots
- have superhuman vision that operate in all weather and lighting conditions recognising where they are and what is going on around them
- navigate using signs created for humans
- recognise complex objects after having seen them just once and remember how to interact with them
- engage in lifelong continuous learning about their immediate environment and the world
- share not just physical work but what they see and know about the world
- are deployed on very large scales, automatically discovering and using sensing and computational resources around them

While some of these things might seem like science fiction, already within the Centre we have:

- created an underwater robot that culls crown of thorns starfish – addressing an impending ecological crisis, without the cost and risk involved in using human divers;
- created an agricultural robot that can see the difference between weeds and crops and apply the correct dose of the correct herbicide precisely on the weed – potentially reducing herbicide usage by 70%;
- created a fruit picking robot that picks the right fruit the right way every time 24x7;
- earned 1st place in the international Pascal visual object classes (VOC) challenge;
- earned 6th place in the international Amazon Picking challenge using vision to compete against 15 other research groups to complete a complex robotics task;
- launched free global university level courses in robotic vision – nearly 100,000 students from over 100 countries have taken the courses so far;
- held a series of robotic vision workshops at both robotics and computer vision conferences to raise the profile of this new field; and
- been active in the debate about robots, jobs and society.

By 2030 computers will be almost 30x faster than today, cameras will be cheap, and the need for robots will be pressing. Smart machines harnessing the power of deep learning is one of the top 10 strategic technology trends predicted by market research company Gartner and the potential rise of new business models involving M2M (machine-to-machine) sales. The pace of change is relentless and the world needs to be able to adapt and make the most of the opportunities that will eventuate using robots that understand their environment, adapt to change and are able to cooperate effortlessly with human co-workers.

Highlights for 2016

10 AWARDS

17 NEW PHD STUDENTS

84 MEDIA PIECES

250k SOCIAL MEDIA HITS

64 BRIEFINGS

32 JOURNAL ARTICLES

10 ROBOTIC VISION WORKSHOPS *(ACADEMIC & INDUSTRY)*

Results in global competitions

1st Place	Google Impact Challenge ($750k)
1st Place	Pascal visual object classes (VOC) Challenge
1st Place	European Conference on Computer Vision (ECCV) Visual Object Tracking-TIR Challenge
2nd Place	Maritime RobotX Challenge
6th Place	Amazon Picking Challenge

Centre Director Peter Corke met with Professor Oussama Khatib, Director of Stanford Robotics Lab, during his visit to Stanford University.

Director's Report

2016 was the year where we positioned ourselves for future success

Centre Growth and Development

I'm proud to reflect on a year where we have made real progress in bringing the disciplines of robotics and computer vision together and building a robotic vision community. More than 80% of the Centre's resources are in the form of people and it's no secret that people with skills in the emergent field of robotic vision are highly sought after due to the transformative nature of their skills set. This has presented our Centre with significant recruitment challenges but also highlights the enormous opportunity (and responsibility) we have to develop people with robotic vision knowledge. Two of our PhD graduates are now working in industry research positions, including Zongyuan Ge (Postdoctoral Researcher at IBM) and Zetao Chen (Postdoctoral Researcher at the APRIL Laboratory – linked to the Toyota Research Institute). Two of our Chief Investigators (CIs) have taken temporary positions with industry, Steven Gould with Amazon, and Ben Upcroft with Oxbotica, an Oxford University spinout. We have enacted our succession plan and will welcome three new Chief Investigators to the Centre in 2017, Jonathan Roberts, Tat-jun Chin and Matt Dunbabin. All three of our new CIs have a strong track record of contribution to the Centre through their former roles as Associate Investigators.

We now have our full quota of research fellows and professional staff (welcoming a Centre Administrator at ANU in 2016) and we have more than 40 PhD students working alongside our 13 Chief Investigators, 18 Associate Investigators and 9 Research Affiliates (see PEOPLE p. 57). We expect our early career researchers will go on to work in industry, government and academia and become the knowledge leaders of the future in the new field of robotic vision. They will help translate research results in robotics and vision to existing industries and convert their research into new products, services and enterprises so as to forge whole new industries. We will support our early career researchers in this journey to ensure their success (see Research Training p. 58).

Developing Centre Culture

In January we started a monthly Centre News providing snapshots of our Centre: our people, our projects and our science, as well as what's happening in the wider world of robotic vision. The newsletter goes to 160 people. We have held three "RoboVis" symposia since our formation in July 2014. Our symposia are an opportunity for all members of our Centre; CIs, AIs, research fellows, PhD researchers, and friends to gather together in one place. We named our annual symposium "RoboVis" as the term symbolises what the Centre is all about, bringing researchers from the fields of robotics and computer vision together. RoboVis2016 was held in Victoria in September 2016 and was attended by 80 people, with a number of Centre Advisory Committee and End-User Advisory Board members attending. New awards were presented to members of the Centre for exceptional performance in research, collaboration, outreach and profile raising. RoboVis incorporated a full training day for research fellows and PhD researchers around leadership and career development. In 2017 we will deliver a concentrated program of training in knowledge leadership to our early career researchers. We also held our second Robotic Vision Summer School (RVSS) featuring international speakers from our partner and other overseas organisations (see CASE STUDY p. 55).

Building a Robotic Vision Community

The seven-year funding of our Centre allows us to undertake the long-term research and cultural change required to successfully meld the disciplines of robotics and computer vision together. This gives us an advantage over other research organisations working in this space. We believe that scientific frontiers will be most advanced by combining both disciplines, with breakthrough technological progress to follow. Large-scale projects associated with the Centre in the areas of environment, healthcare, food production and resources provide significant opportunities for Australian researchers. The most effective large-scale project we have tackled so far is the Amazon Picking Challenge (APC). This competition, sponsored by Amazon, seeks to develop robotic and perception technology that allows a robot to pick a specified item from a shelf compartment that contains a mixture of items. Our team comprised researchers from QUT and Adelaide, and represented Australia in the international competition held in Leipzig, Germany (see CASE STUDY p. 30).

In an effort to bring the global robotics and computer vision communities together the Centre has run a number of successful workshops at international conferences, promoting vision in robotics and vice versa. These have attracted high-profile invited speakers (including several PIs and CIs)

Centre Director Peter Corke presents the "Centre Citizen Award" to Trung Pham at RoboVis, the Centre's annual symposium in September 2016, in recognition of Trung's outstanding contributions to the Centre.

Centre Director Peter Corke hosts one of the many tours showcasing our robotic vision capabilities

and large audiences. Nearly 200 people attended our half-day tutorial on Robotic Vision at the International Conference on Robotics and Automation (2016) organized by Dr François Chaumette (Centre PI) with speakers including me as Centre Director. More than 150 people attended our full day workshop, "Are the Sceptics Right? Limits and Potentials of Deep Learning in Robotics" held at the Robotics Science and Systems Conference. The workshop was organised by Centre research fellows Dr Juxi Leitner and Dr Niko Sünderhauf, with speakers from Google DeepMind, Berkeley, Stanford, MIT and Cornell University.

Our Centre encourages, and has funding to support, visits between the laboratories for our investigators, research fellows and PhD researchers, creating a transnational research community around robotic vision. We have hosted more than 20 international researchers and visited more than 30 international labs and facilities. (see Linkages p. 53). We are also looking to create a network of companies in Australia that are in the robotic vision space. A major project for 2017 will be the creation of a robotic vision roadmap for Australia, raising the industry's profile and allowing measurement of the significant benefit of robotic vision to the Australian economy.

Taking Robotic Vision to the World

In 2016 the Centre ran a series of workshops for industry in both capital cities and regional areas. We partnered with AusIndustry, to promote and educate people about robotic vision. The Centre garnered media coverage, provided globally accessible robotic vision education programs, formed networks to help support people learning code, and have hosted a range of visits to Centre nodes from a wide range of stakeholders (see Engagement p. 45). At the end of 2016 we established a quarterly external newsletter to promote the Centre's work.

We have identified three impact areas where the Centre can make a difference, broadly in the areas of environment, healthcare and food production/resources. The Centre brings together the top Australian research groups in robotic vision and is creating the largest group of its kind in the world. The Centre is generating: internationally impactful science; intellectual property and technologies that will transform important Australian industries; new research capacity; highly skilled people; and an increased reputation for Australian research.

Centre Impact

To ensure the Centre is tracking towards success we have instituted node level performance targets to increase visibility and responsibility around our key performance metrics. Regular review of these figures by the Centre Executive will ensure that we make changes to our direction as required. Following a review of our research, in May 2016 we established a research committee, comprising the executive and the program leaders, which meets every month. This has provided better visibility of progress at project level and the group is empowered to resolve project-level challenges.

To test the effectiveness of the communication of our impact, we held a Mock Mid-term Review on the 31st August chaired by Professor Ian Mackinnon, who has participated in 39 reviews for the ARC. We asked the panel to look at the Centre's

performance to date. The review panel challenged us to demonstrate the value of our activities and we found that by the end of our third year the total value of the Centre's activities exceeded $20m with over 30% being from industry sources.

The review led us to improve the messaging around the purpose of our Centre and our path to impact and in response we developed a set of Impact areas - societal needs that our technology can address (see Impact Areas p. 36).

We encourage our younger researchers to think about impact beyond academic publications. This was the main topic of my talk at our annual conference and was reinforced with a short talk by Advisory Committee member Hugh Durrant-Whyte. We continue to push this message through the Centre newsletter and executive team conversations with researchers.

The high quality of our Centre's research is demonstrated by the recognition we receive. Amongst many awards and achievements this year, two of our senior researchers, and members of the Centre Executive, Prof Rob Mahony and Prof Ian Reid were recognised as Fellows by two of Australia's top scientific societies in 2016 (see Research Performance p. 19)

Challenges

With both robotics and computer vision booming, retention of key research staff will be increasingly challenging. A number of personnel have left the Centre during 2016 and recruitment will be a priority in 2017. We aim to encourage retention by providing a vibrant, high energy, future-focused and collaborative environment so that we can continue to attract and retain high quality researchers.

Outlook for the Centre

2017 is the year where our research begins to get traction. With an increased focus on impact, the next twelve months will see us delivering breakthrough technologies as a result of the partnership between robotics and computer vision, and promoting the importance of robotic vision amongst industry, government and the general community. We will benchmark ourselves against the best in the world. We are commissioning an international review of the Centre's science in July. Our funding body, the Australian Research Council, will conduct a mid-term review of the Centre's activities in September. We will continue to develop the knowledge leaders of the future and will expand our influence and leadership role in the robotic and vision communities. By the end of the year we hope to be well on the way to establishing the framework of a robotic vision roadmap for Australia.

I would like to thank all our colleagues - Centre researchers, operations staff and executive - for their contributions to a Centre that is already achieving impact. We are now starting to scale into an enterprise that will have influence beyond the bounds of academia and positively contribute to the sustainable well-being of people and the environments they live in.

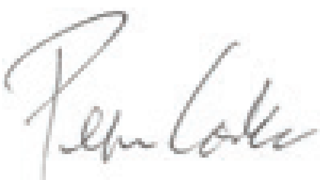

Peter Corke
Centre Director

Centre Operations Team (from l to r, Thuy Mai, Rachel Sinnott, Kate Aldridge, Sue Keay, Sandra Pedersen, Sarah Allen) with @Google Impact Challenge winners Matt Dunbabin and Feras Dayoub #rvlogo wear it and share it

CASE STUDY: Google Impact Challenge Winners

Centre Researchers Dr Matthew Dunbabin and Dr Feras Dayoub teamed up with the Great Barrier Reef Foundation to win the People's Choice Award of the Google Impact Challenge Australia.

Their project to create a low-cost 'robo reef protector' won the people's choice award in the challenge. The award is worth $750,000.

The foundation says the team will build on the researchers' successful COTSbot platform, which was designed to tackle one of the greatest threats to the reef, the crown-of-thorns starfish, or COTS.

"To be recognised in this way is pretty awesome," said Dr Dunbanin. "We learnt a lot from COTSbot, what works, what doesn't work. What we learnt will be brought into the new design."

The team is now working to create the RangerBot, a low-cost, more versatile version of the COTSbot. It will do that by shrinking COTSbot, adding a suite of vision-based sensors and developing a range of attachments to tackle various monitoring and management activities along the Great Barrier Reef

"This is a fantastic opportunity that opens the door to building more robots that will help protect the Reef," said Dr Dayoub, a Research Fellow with the Centre. "This motivates us to answer the support of the people who voted for us by working very hard towards building a great robot, the RangerBot."

The Google Impact Challenge Australia was created to help not-for-profit organisations develop technologies that can help tackle the world's biggest social challenges. Certainly, the Great Barrier Reef and its health is an important challenge facing Australia. In fact, the foundation says the reef contributes $6-billion annually to the Australian economy.

"We wouldn't be here without the support of the Great Barrier Reef Foundation, an organisation truly dedicated to reef conservation," said Dr Dunbabin.

The funding will also allow the team to drive down the cost of building the robot, making it affordable for communities.

Market Trends

Moore's law predicts that the overall processing power for computers will double approximately every two years, simultaneously computing power available per dollar has been increasing by a factor of ten approximately every four years.

Vision sensors (cameras) are becoming cheaper, more highly integrated with computation, and can be found almost everywhere – moving from the insides of pockets and bags to being always on and outward looking. Networking bandwidth will continue to increase and the cost of memory will continue to fall and the world population of robots will continue to grow. The market for computer vision and robotic technologies is booming. Sales of machine vision gear that help computers understand their surroundings and identify images, much like a human can, are rapidly growing, with the machine vision camera market set to be worth more than USD$2b by 2021 (Source: Industry ARC). Large internet companies like Google and Facebook are key investors in vision technology and the software to make it work. The trend for these leading companies, to acquire companies that specialise in robotics, vision and machine learning, increased in 2016 with more than USD$18.9b worth of acquisitions made in robotics alone (see Figure 1).

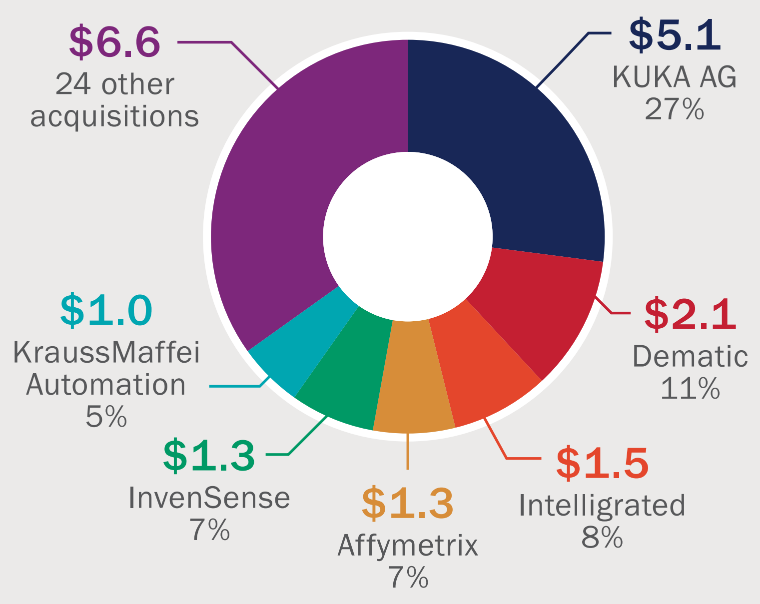

Figure 1: USD$18.9 billion spent on acquisition of 30 robotics companies in 2016 (20 more robotics companies acquired that did not disclose transaction amounts). Source: The Robot Report [https://www.therobotreport.com/news/18.5-billion-paid-to-acquire-49-robotics-companies-in-2016]

The new wave of robotic and computer vision technologies are in Phase 2 of Gartner's hype cycle, the peak of inflated expectations. Gartner's hype cycle charts emerging technologies as they mature and are adopted and applied by business. Each year scores of robotics start-ups are acquired. The story is similar for computer vision start-ups.Some are acquired with little fanfare, such as the speculated acquisition of ZurichEye by Facebook. Computer vision is a key application of artificial intelligence (AI) and machine learning, and all technologies seem to be hitting the peak of inflated expectations.

Robotic vision is the gateway to a whole new set of technologies, developed by bringing robotics and computer vision together. While robotics is about machines that perceive and interact with the physical world, computer vision involves methods for acquiring, processing, analysing and understanding images using a computer. Combining the two produces the key technologies that will allow robotics to transform the way we live and work by giving robots visual perception. So where is robotic vision on the hype cycle?

Robotic Vision is in Phase 1, as evidenced by our proof-of-concept projects like the crown-of-thorns starfish robot, COTSbot and the agricultural robots AgBotII and Harvey. These technologies will soon be heading into Phase 2, the peak of inflated expectations — but how long will the hype of Phase 2 last before we enter into Phase 3, the trough of disillusionment?

The risk for robotic vision is that by the time start-ups have formed around newly created technologies, the hype cycle for related technologies will change. Robotic vision technologies may skip the peak of inflated expectations altogether, instead enjoying an accelerated path towards the slope of enlightenment on the heels of related technologies. Good news for the application of robotic vision to solve real world challenges.

In the current climate of demand, Australia is in a position to create a new export industry in robotic and vision technologies leveraging the unique capabilities we develop to suit the Australian environment. Although the makers of such technologies in Australia are currently fragmented, we believe that Australia will follow world trends in the increased demand for robots and robotic vision technologies. This will trigger the need for improved coordination amongst industry players, more cohesive local supply chains, and the need for more people with the skills and expertise to lead a new era of companies that combine robotics and vision. We plan to continue our leading role in the creation of Australian robotic vision technologies and the development of a robotic vision industry, improving collaboration between researchers and industry.

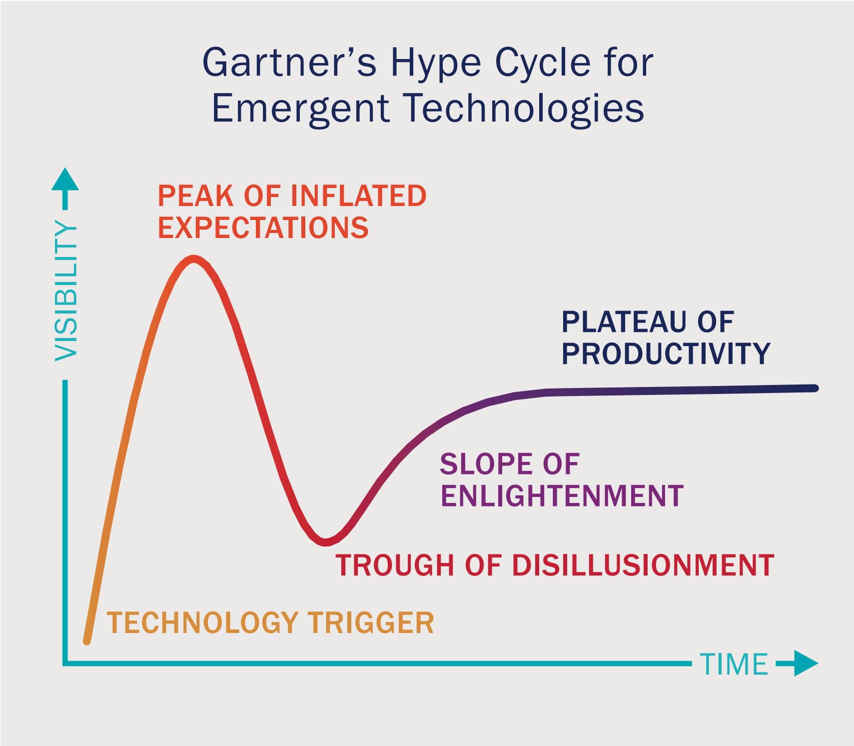

Social Impact of Robots

"Will robots take our jobs?" is one of the first questions asked by visitors to the Australian Centre for Robotic Vision. More recently the question has slightly changed to, "When will robots take our jobs?"

Researchers tend to see the many limitations of robots and dismiss these concerns, for example, our drinks-carrying robot is trained to avoid obstacles - including obstacles in the form of people who want a drink. But while robots are a long way from having the balance and dexterity of humans, people do have cause for concern. The pace of technological change is overwhelming. Only 10 years ago the iPhone did not exist and the first autonomous vehicles bristled with so many sensors, and so much onboard computing hardware, they would have struggled to carry a passenger. Ten years later and today your smart phone is orders of magnitude more powerful than the mainframe computers that put man on the moon, and autonomous vehicles are sleek and ready for passengers.

While it takes human's 3-4 years to gain enough mastery of a subject to gain a university degree, IBM Watson can process 500Gb of data, the equivalent of reading a million books, per second. And while we humans can gain competence through our years of work experience, deep learning enables intelligent machines to also learn from their experiences, or indeed from ours (see robots learning to cook by watching YouTube videos). The only difference is that once one machine learns, that knowledge can be transferred to all networked machines. Imagine being able to share all the information you have gained from your life experiences with every other human on the planet.

The loss of millions of manufacturing jobs in the U.S. due to offshoring was driven by lower labour costs in other countries, and not, initially, because of robots. Indeed robots have been developed as a specific onshoring strategy to bring manufacturing back to the U.S.. Australian Rodney Brooks, the inventor of the Roomba vacuum cleaning robot, founded Rethink Robotics with the aim of making American manufacturing more competitive. Rethink produces the Baxter robot, a cheap, flexible, human-safe collaborative robot that we use in one of our labs.

Ironically, Bill Gates has recently called for robots who replace human workers to be taxed at the same rate that human workers are taxed. This seems fair, but how exactly will we know when robots have taken our jobs? While the U.S. installed about 135,000 new industrial robots between 2010 and 2015, the number of employees in the automotive sector increased by 230,000 during the same period (IFR International Federation of Robotics) and there is a complex relationship between automation, employment growth and productivity (A3 White Paper).

Indeed it is automation rather than robots per se that are the biggest threat to human jobs, it's just that robots are the obvious, physical manifestation of artificial intelligence.

Gartner recognises many other manifestations of AI and machine learning in their analysis of top strategic technology trends for 2017 - virtual assistants (chatbots) and operational applications (robotic process automation). Threats to job security may be invisible as virtual and inconspicuous intelligent apps take over white collar work.

"Taking the robot out of the human," is how Leslie Wilcocks, London School of Economics, describes robotic process automation. Repetitive tasks are not a strong point for humans so why not let robots do that work? Unlike the job losses of the past, concentrated amongst blue collar workers, labour disruption caused by AI and machine learning will impact white collar workers.

White collar or "professional" workers tend to be more resilient to job loss. Arguably it is easier for white collar workers to move and retrain for different jobs if displaced by "robots". Does this mean we shouldn't be worried about robots taking our jobs?

Although the effect of automation on the workplace may be overstated the discontent in other countries over job security and wealth distribution means that we should be worried.

We are living in an era where earnings inequality continues to grow. Over the past 40 years in Australia, wages have risen by 59% for the top 10th of the population and by only 15% for the bottom 10th. According to labour economist John Mangan, "the pay gap is now so vast that while people may be sharing the same geographical area, they may as well be living in a different society". We must strive to ensure that the benefits of automation are shared amongst all of humanity.

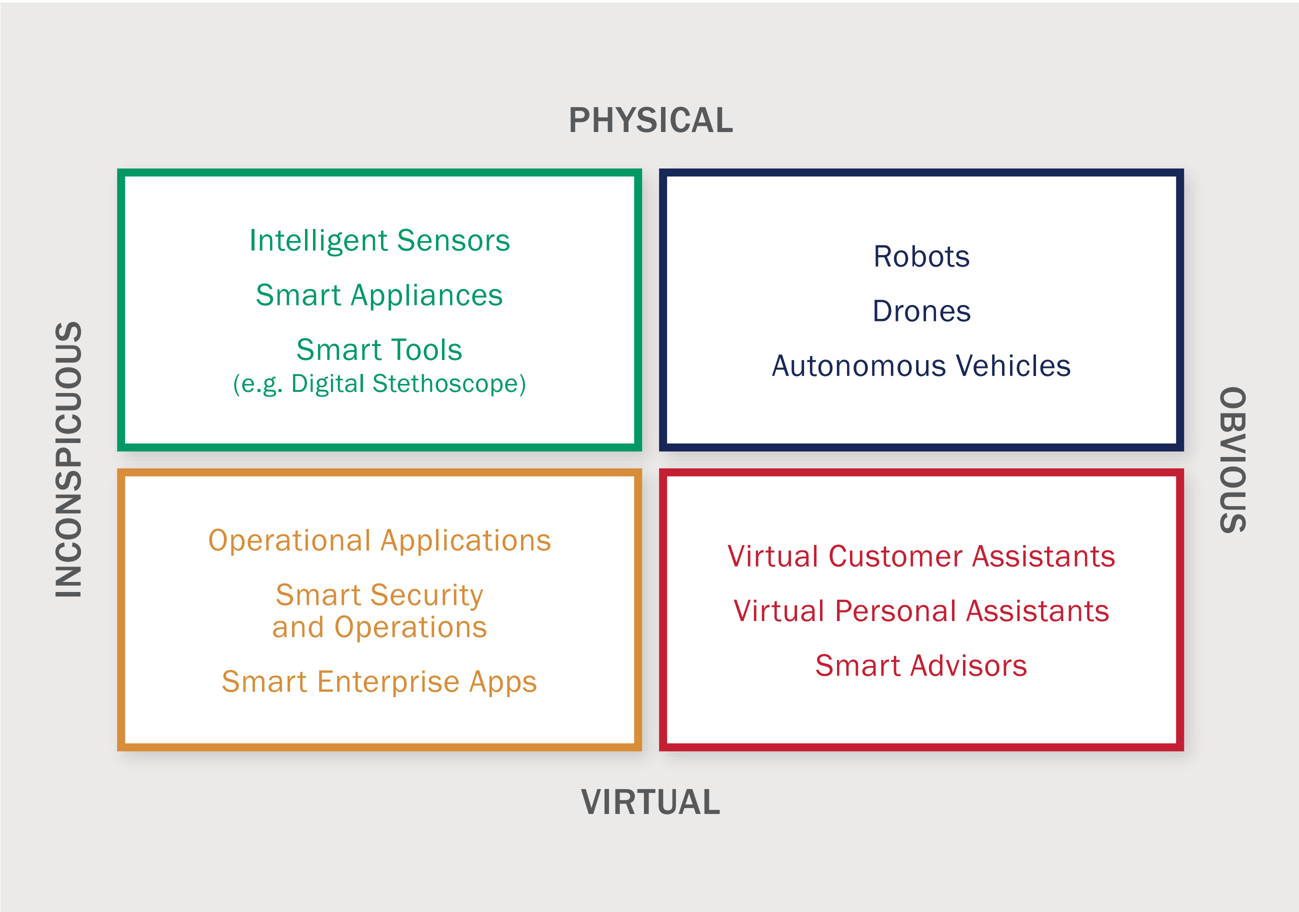

Source: Uses of AI and Machine Learning, Gartner (October 2016)

Research Performance

Research Structure

Our Centre will achieve breakthrough science and technology in robotic vision by addressing five key research programs: Sensing, Understanding, Acting, Learning, and Technology (Figure 2). Programs are a portfolio of projects led by a Centre CI who helps project leaders to craft project plans with challenging but achievable and measurable goals, and a realistic timeline. Program leaders meet frequently with their project leaders; represent the projects in research committee meetings; resolve equipment, staffing and other project resourcing issues; as well as highlighting project achievements. Programs meet assigned Centre KPIs.

Centre projects are our quantum of research (PhD students, RFs, CIs, AIs) within each of our five research programs. Each project has: a leader; a clear objective; a duration; and clear milestones. Ideally a project includes researchers from more than one node. Projects have variable duration and are dynamic. New projects are created by our Research Committee, which determines research staffing and project duration. All Centre PhD researchers must belong to, and contribute to, a Centre project. The objectives of PhD research must significantly overlap with the objectives of the Centre project. We strongly encourage undergraduate Honours projects supervised by a CI, RF or AI, that fall within the research areas of the Centre. All projects must develop technology that demonstrates the capability of robotic vision, that is, can be demonstrated on a robot. This is the litmus test that can be applied to all our projects to ensure they are relevant to the Centre's vision.

A new research committee, comprising the executive and the program leaders, was established in May 2016 and has met every month to review research progress, to make decisions on the direction of the Centre's research, and to annually review and modify the strategic direction of the Centre's research to ensure the Centre can meet its many objectives. Our new research committee structure has provided better visibility of progress at project level and the group is empowered to resolve project-level challenges.

Chad Odsey demonstrates some of the Centre's research at the Real Futures Conference, Brisbane, October 2016.

Centre Research Programs and Projects

Our ambitious research programs will develop technologies that will harness the rich information from visual data to allow robots to perceive the world and be truly useful to humans. We will deliver our vision through five cross-connected research programs:

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
Robust Vision (RV)	**Semantic Representation (SR)**	**Vision & Action (VA)**	**Visual Learning (VL)**	**Algorithms & Architecture (AA)**
RV1 *Robust robotic visual recognition*	**SR1** *Video Representations*	**VA1** *Learning Robust Hand-eye Coordination for Grasping in Novel Environments*	**VL1** *Fundamental Deep Learning*	**AA1** *VOS-distributed robotic vision*
	SR2 *Representing Human-Object Interactions*			**AA2** *ACRV SLAM Framework*
RV2 *Novel visual sensing & hybrid hardware for robotic operation in adverse conditions & for difficult objects*	**SR3** *Scene Understanding for Robotic Vision*	**VA2** *High Performance Visual Control and Navigation in Complex Cluttered Environments*	**VL2** *Learning for Robotic Vision*	**AA3** *Computer graphics simulation for robotic vision*
	SR4 *Joint Vision & Language Reps for Robots*			

Figure 2: Centre Research Programs

Activity Plan for 2016

CULTURE

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

Strategic Objectives	Key Tasks
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**	• Maintain full complement of Research Fellows • Aim for 90% 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 • Centre Executive meetings at least fortnightly with quarterly F2F meetings

SCIENCE

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

Strategic Objectives	Key Tasks
Deliver internationally recognised research in robotic vision **Create and implement projects based on collaboration and innovation that enhance research outcomes**	**Research Program: Sensing** • High accuracy (few centimetres) camera-based localisation module demo • Polarized vision water detection for water hazard traversal and avoidance on a mobile robot **Research Program: Understanding** • Video representations and human action recognition for human-robot interaction • Recognise and discover objects/instances and their affordances from real-world videos of human-object interactions; predict human action and provide active support during a simple assembly task • Real-time high-level (semantic) scene mapping, via development of new representations of objects and scenes • Vision-based human-robot dialogue for describing and following a navigation route, or describing and executing an assembly task **Research Program: Acting** • Pick-and-place demonstrator to operate in a complex unstructured environment such as picking pens from a desk • Development of infinite dimensional lightfield depth and flow algorithms • Reactive control strategies for high speed aerial obstacle avoidance based on flow. Implementation on an aerial robotic demonstrator **Research Program: Learning** • Vision tasks for the APC - open set recognition + uncertainty estimation • Video-rate approximate semantic segmentation on a mobile, power-constrained **Research Program: Technology** • Vision Web Service Framework that demonstrates running vision algorithms as web services • Semantic SLAM with object detection demonstrator • Underwater robot visual odometry and SLAM demonstrator

INTEGRATE

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

Strategic Objectives	Key Tasks
Connect research organisations, governments, industry and the private sector to build critical mass in robotic vision Lead RV in Australia and overseas	• Create a 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

ENGAGE

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

Strategic Objectives	Key Tasks
Identify and engage with key stakeholders on the potential applications of RV Establish vibrant national & international RV communities Increase RV educational opportunities	• Engage Centre's End-User Advisory Board at least twice yearly • Host a visit from the Australian Research Council • Create a Robotic Vision resources hub • Support PD for teachers in robotics and coding • Host visits and tours of robotic vision facilities by government, industry and the community (including school groups) • 4 x media releases by partners related to robotic vision • 4 x public lectures on robotic vision

TRANSFORM

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

Strategic Objectives	Key Tasks
Demonstrate how research can advance products and services Generate downstream investment to take RV technology into industry Foster innovation, entrepreneurship and new enterprises to advance RV	• Provide knowledge leadership training to all early career researchers to help in their career development • Engage Centre's End-User Advisory Board at least twice yearly • Continue to build network of industry contacts and maintain CRM • Present on robotic vision at key industry events • Targeted publications on Centre work • Develop framework for a robotic vision roadmap of Australia

Research Performance (KPIs)

ARC Centre of Excellence Key Performance Indicators (KPIs)	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Conference publications	25	42	50	53	50	66
Journal publications	12	15	25	22	25	32
Paper prizes or awards	1	6	3	1	5	4
Disclosures/Patents	0	0	0	0	0	1
Quality of research outputs % of papers published in peer-reviewed outlets	100	100	100	100	100	100
Percentage of publications relevant to SRPs	80	80	80	100	80	100
Number of invited talks/papers/keynote lectures given at major international meetings (incl. int'l conferences held in Australia)	5	6	10	41	10	38
Number and nature of commentaries about the Centre's achievements (in professional magazines and the popular press - newspapers, TV etc)	5	49	10	94	20	84

Centre researchers explore each other's work at demo day, an important component of RoboVis, the Centre's annual symposium

Recognising our People

Our researchers **Matt Dunbabin, Feras Dayoub** and the Great Barrier Reef Foundation won the public vote in the Google Impact Challenge and a $750,000 prize. This will go a long way toward helping them achieve their dream of a low-cost totally visually guided robot for monitoring the health of the reef.

The Centre had two finalists in the 2016 Australian Museum Eureka Prizes. Centre **Chief Investigator Michael Milford** was a finalist for the Eureka Prize for Outstanding Early Career Researcher. Our COTSbot team of **Matt Dunbabin, Feras Dayoub, and Peter Corke** were finalists for the Eureka Prize for Environmental Science. The Eureka Prizes are Australia's most prestigious science awards, and celebrate excellence in several areas, including research and innovation

The Centre had three great successes at the European Conference on Computer Vision (ECCV) held in October 2016. ECCV is a biennial computer vision conference considered to be one of the top three in the field, along with Computer Vision and Pattern Recognition (CVPR) and International Conference on Computer Vision (ICCV). Centre researchers **Gustavo Carneiro (CI), Hongdong Li (CI), Laurent Kneip (Associate Investigator) Peter Anderson (PhD Researcher) and Ravi Garg (Research Fellow)** travelled to Amsterdam and were able to celebrate these successes in person. The prestigious Koenderink Paper Prize was awarded to **Ed Rosten and CI Tom Drummond** for their 2006 ECCV paper "Machine learning for high-speed corner detection". This prize is given to fundamental contributions in computer vision that have stood the test of time. Google Scholar shows 2400+ citations for this work. **Partner Investigator Professor Andy Davison** and his ICL team **Hanme Kim and Stefan Leutenegge**r won the ECCV 2016 Best Paper Award with the paper "Real-Time 3D Reconstruction and 6-DoF Tracking with an Event Camera". Outstanding Reviewer awards were presented to **RF Anoop Cherian and AI Anders Eriksson**, while **Gao Zhu, Associate Investigator Fatih Porkili and CI Hongdong Li** won 1st place and "Best Performing Tracker Award" at the Visual Object Tracking-TIR Challenge.

Best paper award at the International Conference on Digital Image Computing (DICTA 2016) went to Centre researchers **Zongyuan Ge, Chris McCool, Conrad Sanderson, Peng Wang, Lingqiao Liu, Ian Reid and Peter Corke** for their paper ''Exploiting Temporal Information for DCNN-based Fine-Grained Object Classification'', and an honourable mention to a student supervised by **AI Fatih Porikli**, Tong Zhang, at the Asian Conference on Computer Vision (ACCV) 2016.

Google Impact Challenge Vote winners and Eureka Prize Finalists Matt Dunbabin and Feras Dayoub with the COTSbot

Other esteem measures garnered by Centre researchers include:

- **CI Ian Reid** was elected as a Fellow of the Australian Academy of Technology and Engineering (ATSE) in recognition of his many contributions to computer vision and robot navigation.
- **CI Richard Hartley** has been elevated to Fellow of AFCV (Asian Federation of Computer Vision) for his significant contributions to computer vision research especially in the areas of multiple view geometry and optimisation.
- **CI Rob Mahony** has been named IEEE Fellow in recognition of his contributions to control aspects of aerial robotics.
- **CI Michael Milford** has been named as an AMP Tomorrow Maker 2016 and awarded $25,379 to help his Math Thrills program reach thousands of students Australia-wide.
- **CI Hongdong Li** has been appointed as an Associate Editor for the IEEE Transactions on Pattern Analysis and Machine Intelligence (IEEE-TPAMI). CIs Ian Reid and Hongdong Li are currently the only AEs from Australian Universities.
- It has been announced that ACCV 2018 **will be held in Perth. Congratulations to Chief Investigator Ian Reid,** ACCV 2018
- **AI Tristan Perez** has been appointed the Chair of the Technical Committee on Marine Systems of the International Federation of Automatic Control (2018-2021).
- Three of our researchers have been promoted to Centre Chief Investigator, former Associate Investigators: **Dr Tat-Jun Chin (Adelaide)**, **Dr Jonathan Roberts (QUT)** and **Dr Matthew Dunbabin (QUT)**, and we look forward to welcoming them as CIs in 2017.

CASE STUDY: FELLOWSHIPS FOR CENTRE LEADERS

Two top scientific societies recognised members of the Centre's Executive as Fellows in 2016.

The Australian Academy of Technological Sciences and Engineering (ATSE) named Centre Deputy Director Ian Reid (Adelaide) a Fellow for his lifetime of research work. The academy said Ian has substantially shaped the world of robot navigation, using computer vision techniques to map a robot's environment so that it can operate autonomously in complex, urban situations.

"ATSE is Australia's peak body for engineering and technology," Ian said. "I am deeply honoured to be selected."

Ian is a Professor of Computer Science at The University of Adelaide, where he has been since 2012. He was formerly a Professor of Engineering Science at the University of Oxford. His research interests span a wide range of topics in computer vision and robotics.

In addition, Centre Chief Investigator Rob Mahony (ANU) was named a Fellow of the IEEE for his contributions to control aspects of aerial robotics.

"I'm pleased this brings recognition to the work being done in our Centre," Rob said.

The IEEE is the world's leading professional association for advancing technology for humanity. An IEEE fellow is the highest grade of membership and is recognised by the technical community as a prestigious honour and an important career achievement. Rob is a Professor in the Research School of Engineering at ANU, where he has been since 2001. His research interests are in non-linear control theory with applications in robotics, geometric optimisation techniques and systems theory. Rob's honour increases the Centre's tally of IEEE fellows amongst Chief and Partner Investigators to eight.

Centre Deputy Director Ian Reid (left) with Centre RF Yasir Latif

Centre Chief Investigator Rob Mahony

L to R: Research Fellow Feras Dayoub and Chief Investigators Gordon Wyeth, Matt Dunbabin and Michael Milford at the 2016 Australian Museum Eureka Prize awards dinner in Sydney.

CASE STUDY: Eureka Prize Finalists

A Centre Chief Investigator and a Centre research team were both named finalists for the 2016 Australian Museum Eureka Prizes. The Eureka Prizes are Australia's most prestigious science awards, and celebrate excellence in several areas, including research and innovation. Centre Chief Investigator Michael Milford was a finalist for the Eureka Prize for Outstanding Early Career Researcher. Our COTSbot team of Matt Dunbabin, Feras Dayoub, and Centre Director Peter Corke were finalists for the Eureka Prize for Environmental Science.

"I've been awed by some of the past winners and what they've done. So it was really an honour to be nominated and to be a finalist," said Michael Milford.

Michael was named a finalist for his work in robotic vision navigation systems for robotics of all types, and for his work where he blends the fields of neuroscience and robotics. He says his passion is to understand intelligence, and robotics and neuroscience are helping him do that.

"Robotics is very tangible. If a robot is doing something clumsy or dumb, you can see it," Michael said. "I think the things that robots do, like navigation, and helping us out with tasks, are a great way to leapfrog our way to really understanding what intelligence is, and how we can replicate it."

For Centre Director Peter Corke, the COTSbot project is one he considers to be a great example of a robotic vision system in action. The COTSbot robot was designed to help target a very real threat to the Great Barrier Reef – the crown-of-thorns starfish, or COTS. Using a sophisticated vision system designed by the team, the COTSbot identifies a COTS and injects it with a solution to kill it.

Matt Dunbabin designed the robot to be a first responder system to beef up the existing program that uses divers to hunt for COTS. Dr Feras Dayoub trained the COTSbot to pick out the pest from other sea life. In other words, it can think for itself using a visual recognition system made possible by algorithms to identify COTS in the visually challenging environment of the Great Barrier Reef.

"I'm very proud to be associated with this prestigious prize," said Dr Dayoub. "This is one example of where research can make a real difference in the world."

Dayoub says the COTSbot basically came together over the course of just one year.

"It was a rapid transformation between ideas, research, and then building this robot and testing it on the reef."

Both the COTSbot team and Michael they say they were just thrilled to be nominated and be included with the amazing group of people who were also nominated and honoured by the Australian Museum.

SENSING – Robust Vision (RV)

Program Leader
Michael Milford

Michael Milford (right) demonstrates his work to Minister for Education and Training Senator Simon Birmingham

Bio

Michael Milford is an Associate Professor at QUT, Australian Research Council Future Fellow and Microsoft Research Faculty Fellow. Michael conducts interdisciplinary research at the boundary between robotics, neuroscience and computer vision, modelling the neural mechanisms in the brain underlying tasks like navigation and perception to develop new technologies like all-weather, anytime positioning for autonomous vehicles. He was a national finalist for the 2016 Eureka Prize for Outstanding Early Career Researcher.

Michael is also an educational entrepreneur and recently launched the company Math Thrills, which combines mass market entertainment and STEM education. Math Thrills received pre-seed funding on Kickstarter ($2500), seed funding ($50,000) from QUT Bluebox, an AMP Tomorrow Fund award ($25,000) and is in school trials. The initiative has led to a 2015 TedXQUT talk and multiple awards including the 2015 Queensland Young Tall Poppy of the Year Award.

Program Description

The Robust Vision program develops robotic vision algorithms and novel vision hardware to enable robots to see and act in all viewing conditions. The program is developing a suite of algorithms that enable robots to perceive their environments and consequently act purposefully under the incredible range of environmental conditions possible, including low light, rain, snow, ice sleet, fog, smoke, dust, wind, glare and heat. In addition, it is further developing innovative sensing hardware to facilitate robot operation under challenging viewing conditions such as low light, or through partial obscuration cameras and hyperspectral cameras.

The key question we are addressing is, how can innovations in existing computer vision and robotic vision techniques and vision sensing hardware enable robots to perform well under the wide range of challenging conditions encountered by robots and applied computer vision technology in the real world?

RV1: Robust Robotic Visual Recognition

Project Leader
Michael Milford

Bio

(please see RV program leader biography)

RV1 Research Team

Michael Milford (CI), Hongdong Li (CI), Chunhua Shen (CI), Jonathan Roberts (AI), David Thompson (PI), Abigail Allwood (PI), Nigel Boswell (PI), Chuong Nguyen (RF), Niko Sünderhauf (RF), James Sergeant (PhD), Sean McMahon (PhD), Sourav Garg (PhD), James Mount (PhD), Medhani Menikdiwela (PhD), Juergen Leitner (RF)

RV1 Project Description

Robust robotic visual recognition for adverse conditions will develop algorithms that solve the fundamental robotic tasks of place recognition and object recognition under challenging environmental conditions including darkness, weather, adverse atmospheric conditions and seasonal change, and translate them into applications in industry. This project is relatively mature and hence is now pushing more heavily towards industry outcomes and engagement than some of the other projects.

The key question we are addressing is, how can existing computer vision and robotic vision techniques be innovated to enable them to perform well under the wide range of challenging conditions encountered by robots and applied computer vision technology in the real world?

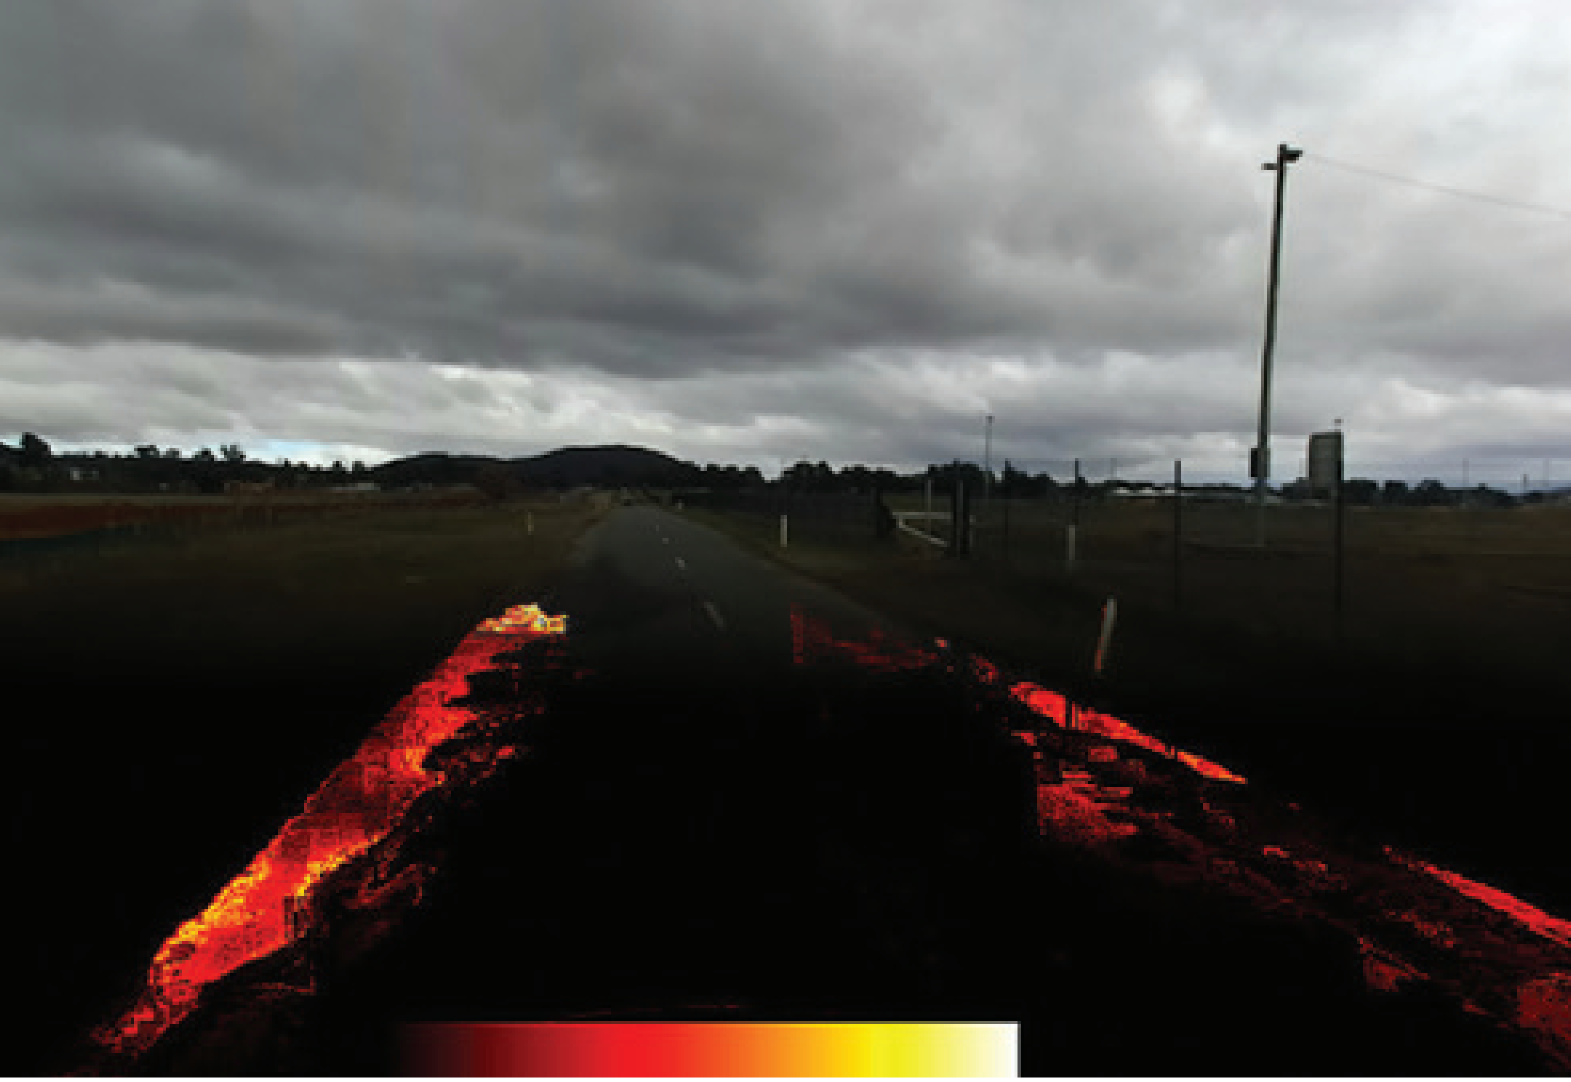

Modelling polarized sky light interaction with water surfaces enables robust detection of water hazards for autonomous robots and vehicles. From Chuong V. Nguyen, Michael Milford, Robert Mahony, "3D tracking of water hazards with polarized stereo cameras", accepted to the IEEE International Conference on Robotics and Automation, 2017

RV2: Novel visual sensing & hybrid hardware for robotic operation in adverse conditions & for difficult objects

Project Leader
Chuong Nguyen

Bio

Chuong Nguyen is a Research Fellow at the Australian National University (ANU). Chuong completed his PhD in Mechanical and Bio Engineering at Monash University in 2010 and joined the Centre in February 2015. Prior to that, he completed his Master of Engineering (M.Eng.), in Eco-Environmental Civil Engineering in 2003 at Ritsumeikan University, Kyoto, Japan. His research interests include autonomous navigation, robot arm manipulation, image-based 3D reconstruction, medical imaging, fluid mechanics and particle image velocimetry.

RV2 Research Team

Peter Corke (CI), Robert Mahony (CI), Michael Milford (CI), Donald Dansereau (PI), Chuong Nguyen (RF), Dan Richards (PhD), Dorian Tsai (PhD), James Sergeant (PhD), James Mount (PhD)

RV2 Project Description

Novel visual sensing for robotic operation in adverse conditions will advance the performance of robot vision algorithms by using new conventional low light cameras, as well as rotational filters, hyperspectral cameras and thermal cameras to improve robot autonomy under any viewing condition. We are looking at the development of new algorithms that exploit the particular advantages of these innovative cameras to enable new performance benchmarks in tasks such as scene understanding, place recognition and object recognition. The development of hardware-software solutions that can deal with these corner cases - reflections, transparency and low light conditions - will be applicable to all Centre projects involving visual sensing and can be used by those projects to robustify their systems for performance under adverse conditions.

UNDERSTANDING – Semantic Representations (SR)

Program Leader
Stephen Gould / Ian Reid

Bio

Stephen Gould is an Associate Professor in the Research School of Computer Science in the College of Engineering and Computer Science at the Australian National University. He received his BSc degree in mathematics and computer science and BE degree in electrical engineering (University of Sydney, 1994, 1996) and MS degree in electrical engineering and PhD from Stanford University. During his time in industry he co-founded Sensory Networks, which sold to Intel in 2013. His research interests include computer and robotic vision, machine learning, probabilistic graphical models, and optimisation. From November 2016, he is on leave from ANU and working on a project at Amazon.

Ian Reid took over leadership of the program following Steve Gould's secondment. Ian is a Professor of Computer Science at University of Adelaide and an ARC Laureate Fellow. He received a BSc (Hons) in Computer Science (1987) UWA, and a Rhodes Scholarship leading to a D.Phil, University of Oxford (1992). Prior to his move to Adelaide in 2012, he was a Professor of Engineering Science at the University of Oxford. His work in computer vision includes seminal papers in active vision, visual tracking, SLAM and human motion capture. He has co-authored prize-winning papers at CVPR (2008), the British Machine Vision conference (2005, 2009, 2010) and the International Conference on 3D Vision (2014). He serves on the editorial board of IEEE T-PAMI and Computer Vision and Image Understanding. In 2016 he was elected Fellow of the Academy of Technological Sciences and Engineering.

Program Description

The Semantic Representations (SR) 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. The representations explored in this project will enable the recognition of activities from observer and actor viewpoints, fine-grained understanding of human-object, object-object and robot-object interactions, and large-scale semantic maps of environments. The project also investigates the mapping of visual inputs to sequence outputs for achieving a given task (e.g., describing a scene textually or providing a symbolic representation for a robotic task based derived from human demonstration).

Program Leader Stephen Gould (l) with Project Leader Anton van den Hengel (r)

SR1: Video Representations

Project Leader
Basura Fernando

Bio

Basura obtained his PhD from the VISICS group of KU Leuven, Belgium, and a European Erasmus Mundus Masters in Color in Informatics and Media Technology from the University of Saint-Etienne France and the Gjovik University Norway (2011). He joined the Centre in 2015 and facilitated a session on "CNN and Architectures" at CVPR2016. His research interests are computer vision and machine learning.

SR1 Research Team

Stephen Gould (CI), Basura Fernando (RF), Anoop Cherian (RF), Sareh Shirazi (RF), Rodrigo Santa Cruz (PhD)

SR1 Project Description

Robots should not view the world as a series of static images—they need to understand movement and dynamics. However, representing videos or short video segments in a way that is useful for robots is poorly understood. This project addresses the question of learning video representations.

SR2: Representing Human-Object Interactions

Project Leader
Anoop Cherian

Bio

Anoop Cherian is a research follow based at the Australian National University (ANU). He joined the Centre as a Research Fellow in 2015. Prior to that, he was a postdoctoral researcher at INRIA, Grenoble. He received his PhD in Computer Science in 2013 and M.S. in Computer Science in 2010, both from the University of Minnesota. He obtained his Bachelors in Technology (B. Tech) degree from National Institute of Technology, Calicut India in 2002 after which he worked as a software engineer with Microsoft until 2007. His current research interests include human pose estimation and activity recognition, scene understanding and context based reasoning, and machine learning on matrix manifolds. He received the "Best Student Paper Award" at the International Conference on Image Processing (ICIP) in 2012 for his work on robust hashing techniques and an "Outstanding Reviewer Award" at the European Conference on Computer Vision 2016 (ECCV16).

SR2 Research Team

Stephen Gould (CI), Anoop Cherian (RF), Markus Eich (RF)

SR2 Project Description

Human-Object Interactions investigates representations and models for the recognition interactions of humans with objects in an activity. In this setup, actions are expressed as the change in physical state of objects, and new object instances may be discovered while observing the activities. The project deals with human pose estimation and forecasting (in 2D and 3D) and representations for fine-grained activity recognition.

SR3: Scene Understanding for Robotic Vision

Project Leader
Niko Sünderhauf

Bio

Niko Sünderhauf obtained his PhD from Chemnitz University of Technology in Germany in April 2012. After working as a postdoc in Chemnitz, he joined the Robotics and Autonomous Systems Group at the Queensland University of Technology, and became a Research Fellow of the Centre in August 2014. Niko's research interests are within the fields of mobile robotics, computer vision, and machine learning. In 2016, he was involved in the Centre's Amazon Picking Challenge (APC) Team which came in 6th place in the competition. He was a co-organiser for the very successful workshop on deep learning at the Robotics: Science and Systems 2016 (RSS) conference held at the University of Michigan titled ""Are the Skeptics right? Limits and Potentials of Deep Learning in Robotics" which attracted nearly 200 attendees.

SR3 Research Team

Ian Reid (CI), Tom Drummond (CI), Michael Milford (CI), Trung Thanh Pham (RF), Yasir Latif (RF), Niko Sünderhauf (RF), Lachlan Nicholson (PhD), Saroj Weerasekera (Affiliated PhD) and Mehdi Hosseinzadeh (Affiliated PhD)

SR3 Project Description

For many robotic applications we need models of the environment that enable reasoning about geometric, semantic concepts and "affordances" (i.e. action possibilities) jointly. This project aims to develop algorithms and representations to acquire and reason about the uncertain and dynamic environments in which robots must operate. Among other objectives the project will provide the semantic representations for ACRV SLAM. Initial work in this project aims to develop maps of the environment that are labelled semantically with a "standard" set of useful class labels. Work in VL1 is showing how such labels can be generated for single views, but here the aim is to ensure that such labels are consistent (and accurate) within the 3D structure of the scene. We also aim to leverage prior information from scenes that can be learnt via CNNs. We will investigate how the information from thousands or millions of exemplars can be used to improve scene structure without imposing hard constraints such as Manhattan world models. Subsequent work aims to develop representations that enable a scene to be decomposed into its constituent parts and thereby used for planning for robotic navigation or acquisition/manipulation. Representation of uncertainty is a key element here; this is well-understood in the context of geometry, but is a research question how to determine, represent and use uncertainty resulting from inference over semantic entities. SR3 will draw on advances in VL1 to bridge this gap. Later work aims to go beyond simple type labels to a deeper and more powerful set of labels such as affordances of objects.

SR4: Joint Vision & Language Reps for Robots

Project Leader

Anton van den Hengel

Bio

Anton van den Hengel is the Director of the Australian Centre for Visual Technologies, and a Professor of Computer Science at the University of Adelaide. He has been a CI on over $50m in research funding, and leads a group of over 60 researchers working in Computer Vision and Machine Learning. Prof van den Hengel has had research funding from Canon, BAES, Google, Bayer and Microsoft, and has received a number of awards, including the Pearcey Award for Innovation and the CVPR Best Paper Award. He has published over 250 papers, and has 8 patents being exploited.

SR4 Research Team

Anton van den Hengel (CI), Chunhua Shen (CI), Stephen Gould (CI), Basura Fernando (RF), Chao Ma (RF), Peter Anderson (PhD)

SR4 Project Description

This project aims to build joint vision and language representations for describing scenes (captioning), answering queries (VQA), and describing and defining robotic actions. Beyond natural language the project also considers more general image-to-sequence models where a sequence may be intended as commands for robot control.

Unlike many computer vision tasks which can be precisely circumscribed (eg image segmentation) robots in an open world must be capable of learning to respond to unforeseen questions about unforeseen images, or develop action strategies for previously unseen circumstances.

Centre methods have been on top of the two primary image segmentation leader boards in 2016.
Images courtesy of Anton van den Hengel

CASE STUDY: NEW DEEP LEARNING SUPERCOMPUTER

Deep learning is a term that has not been far from anyone's lips since the inception of the Centre, with its impact on computer vision and robotics in the last 2 years both immediate and profound.

A new deep learning supercomputer went online in September 2016 at The University of Adelaide and is now providing the computing power the Centre needs to train deeper and larger models and to handle large numbers of parameters and big data.

Three Centre nodes, Adelaide, QUT and ANU, along with Deakin and UWA were all a part of a successful LIEF (Linkage Infrastructure Equipment and Facilities) grant award by the Australian Research Council. The grant, combined with funding from each of these universities, is worth nearly $400-thousand.

Centre Deputy Director Ian Reid was the driving force in getting the LIEF grant.

"The fact that we now have access to a computer that can train models, do evaluations, and so forth, is very significant for us," Ian said, "because we would not be able to do much of what we're trying to do, particularly in semantic vision, but in all sorts of other areas of the centre, were we not have access to the best models and the best recognition schemes."

In addition to the deep learning supercomputer at Adelaide, the money is also being used to enhance existing supercomputing power at QUT and ANU.

ACTING – Vision & Action (VA)

Program Leader

Robert Mahony

Bio

Rob Mahony is a Professor in the Research School of Engineering at the Australian National University and has been a Chief Investigator with the Centre since its inauguration in 2014. His research interests are in non-linear systems theory with applications in robotics and computer vision.

He wrote the seminal paper providing a clear exposition of non-linear complementary filters on the special orthogonal group for attitude estimation; an enabling technology in the early development of quadrotor aerial robotic vehicles. He was the first to provide a principled analysis for using optical flow of control of aerial robotic vehicles and was a coauthor on the first experimental paper that demonstrated landing of a quadrotor vehicle on a textured but featureless moving surface.

In 2016, Rob was named a Fellow of the IEEE, recognising his contribution to the control aspects of aerial robotics.

Program Description

Vision control of robots for manipulating real objects, pushing the boundaries on speed, coordination & complexity.

Harvey, a prototype robotic platform for harvesting capsicums, being tested in a protected cropping facility. Harvey was developed as part of a three-year Strategic Investment in Farm Robotics (SIFR) funded by the Queensland Department of Agriculture and Fisheries (DAF).

The Centre used a Baxter Robot (pictured) for the vision and manipulation tasks required as part of the Amazon Picking Challenge

VA1: Learning Robust Hand-eye Coordination for Grasping in Novel Environments

Project Leader
Juxi Leitner

Bio

Jürgen "Juxi" Leitner is a Research Fellow based at the Queensland University of Technology (QUT) and project leader for VA1. His interests include robotic learning, computer vision and machine intelligence. His background includes a Joint European Master in Space Science and Technology (SpaceMaster) and a Bachelors degree in Computer Science from the Vienna University of Technology.

Juxi's interest is in making (long-term) robotic autonomy a reality. Currently this is focussing on building a framework enabling eye-hand coordination on armed robots, such as, Baxter or the iCub. He is also trying to put a hopping robot on the Moon and was interviewed by ABC National Radio on his LunaRoo project. He led the Centre's 2016 Amazon Picking Challenge (APC) Team.

VA1 Research Team

Peter Corke (CI), Rob Mahony (CI), Chunhua Shen (CI), Francois Chaumette (PI), Juxi Leitner (RF), Fangyi Zhang (PhD), Adam Tow (PhD), Chris Lehnert (Affiliated RF)

VA1 Project Description

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. This project aims to go 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 VA1 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.

In 2017 the Centre's work on enabling robots to pick a variety of objects in complex environments will continue. A highlight will be the participation during the 2017 Amazon Robotics Challenge in July. In addition, the aim is to go beyond the state-of-the-art in warehouse automation and shelf picking, facilitated through robotic learning through visual demonstrations by human operators.

VA2: High Performance Visual Control and Navigation in Complex Cluttered Environments

Project Leader
Rob Mahony

Bio

(please see VA program leader biography)

VA2 Research Team

Rob Mahony (CI), Peter Corke (CI), Jonathan Roberts (CI), Jochen Trumpf (AI), Jean-Luc Stevens (PhD), Sean O'Brien (PhD), Juan Adarve (PhD), Peter Kujala (PhD), Dorian Tsai (PhD), Alex Martin (Research Engineer)

VA2 Project Description

Efficient and effective manipulation in a complex cluttered environment involves planning ahead and the ability to move quickly and safely in a cluttered environment.

The project will consider real-world control scenarios where there are multiple options to achieve the desired task including motions that interact with the environment; for example, moving a glass out of the way in order to reach a bottle, or picking up a sequence of items from a cluttered workspace in an order learned by the algorithm rather than predetermined by the engineer. Achieving high-performance of such tasks involves two separate capabilities:

1. An integrated control and sensing system that allows the robot to move quickly and robustly through a cluttered environment; in particular, algorithms for high-speed obstacle detection and avoidance control strategies must be developed.
2. An integrated decision and planning capability that allows the robot determine (and execute) effective and efficient solution to a complex task involving multiple components:

The defining aspect of this project will be to develop solutions to these capabilities that exploit vision sensing in a fundamental manner.

Thus, the obstacle avoidance and motion control system is vision-based, with additional sensor modalities as appropriate for given applications. Similarly the decision and planning capability will also be vision-based and will exploit semantic information and other cues as appropriate.

Members of the Centre's Amazon Picking Challenge team (listed in table) with the Baxter robot and with University of Adelaide researcher Trung Pham making an appearance via the Beam robot.

CASE STUDY: Amazon Picking Challenge

Centre researchers had the chance to show off their research on the world stage at the 2016 Amazon Picking Challenge. A team from the Australian Centre for Robotic Vision was one of only 16 teams from around the world selected to compete.

Amazon is already leading the world in logistics robotics with their purchase of the warehousing automation company Kiva Systems for USD$775 in 2015. Amazon uses more than 30,000 robots in its global network of fulfilment (distribution) centres but its robots have limited capability. Despite improvements in vision and manipulation capability, robots are still no match for humans when it comes to identifying and picking things from shelves, hence the origin of the challenge. Can robots automate picking in an unstructured environment?

The 2016 challenge was held in conjunction with RoboCup 2016 in Leipzig, Germany. Centre Research Fellow Dr Juxi Leitner led the team.

"We saw the Amazon Picking Challenge as a very interesting problem for us. It allowed us to bring two areas of research that we think belong together – robotics and computer vision – to create something larger than the sum of those two parts," Juxi said.

Juxi and members of the Centre team saw the Amazon Challenge as a chance to really advance the Centre's mission of creating robots that see and understand a task in the real world environment of the warehouse.

"Our hope was to take what we learnt from this and start applying it to other areas we're trying to solve in the Centre, like agricultural, infrastructure and medical applications," Juxi said.

The team found out in early 2016 it was selected, and only had about five months to get a system up and running for the competition.

"We had to build a lot and come up with a system in that time frame that really worked," said Dr Niko Sünderhauf, a Research Fellow with the Centre.

There were a lot of challenges to overcome in a short amount of time. The team was using a Baxter robotic platform to compete.

"To run the Amazon Picking Challenge and the whole code on the system, we had multiple computers," Juxi says. "There is one in Baxter, there are two other computers, including one that just runs the visual perception. So, there was quite a lot of infrastructure and it took actually quite some time to make all those systems work with each other."

Eighteen members from the Centre were a part of the team, with six of them actually travelling to Germany for the competition. Leading up to the event, teams were given a list of items that they might be asked to pick during the competition.

"You've got no idea on the day what items you will actually be asked to pick out of the shelf, you have no idea what arrangement those items will be, and what the lighting conditions will be like when you get there," said Adam Tow, a Centre PhD student who was a member of the team.

The challenge itself was actually broken into two days of competitions.

The first day was the stowing task, where teams were asked to pick objects out of a tote, and place them in a shelf. The Centre team was able to pick three items out of the tote, successfully placing two in the shelf.

The second day involved the picking task, with the Centre team placing sixth overall with a robust solution to the vision problem, which could be applied despite lighting conditions, something the other teams found challenging.

"We picked four objects from the shelf and we put them into the tray. It was really exciting to see the robot do what it was supposed to do," Juxi said.

"I think for all of us, it was a giant learning curve," Adam said. "We were we absolutely stoked with how we went on the day of the picking task, finishing as high as we did with the calibre of teams that were there."

For team members, the strong showing validated not only their hard work, but their research as well.

"From a scientific point of view, the system that we built, we will continue to work on," Juxi said. "This will be the baseline for our future research, especially focusing on manipulation, agriculture, and cluttered environments like shelves."

That research and hard work is continuing in 2017, with the Centre being selected to compete in the Amazon Robotics Challenge, being held in conjunction with RoboCup 2017 in Nagoya, Japan at the end of July.

TEAM MEMBERS

(Listed alphabetically)

Matthew Cooper

Jake Dean

Markus Eich

Andrew English

Peter Kujala

Chris Lehnert

Jürgen "Juxi" Leitner

Ruben Mangels

Steve Martin

Chris McCool

Lachlan Nicholson

Trung Pham

James Sergeant

Niko Sünderhauf

Adam Tow

Liao Wu

Fangyi Zhang

Assoc Prof Ben Upcroft
Centre Chief Investigator

Prof Peter Corke
Centre Director

LEARNING - Visual Learning (VL)

Program Leader
Gustavo Carneiro

Bio

Gustavo Carneiro is an Associate Professor of the School of Computer Science at the University Adelaide, which he joined as a senior lecturer in 2011. Gustavo has received prestigious fellowships, such as the Humboldt Fellowship for Experienced Researchers (2014) and the Marie Curie IIF Fellowship (2010). Before joining the University of Adelaide, Gustavo worked as a visiting assistant professor at the Technical University of Lisbon (2008 to 2011) and as a senior research scientist at Siemens Corporate Research in Princeton, USA (2006 to 2008). Between 2004 and 2005, A/Prof Carneiro held postdoc positions at the University of California, San Diego (working with Prof. Vasconcelos) and at the University of British Columbia (working with Prof. Lowe). His main research interests are in the fields of computer vision, medical image analysis and machine learning.

Program Description

Visual learning has enormous potential to solve previously impossible problems in machine perception. The recent deep learning breakthrough from the machine learning community has allowed researchers not only to address new visual learning problems, but also to solve old problems. In general, 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 these recent developments, there is a lack of understanding of how deep learning works, which invites questions about convergence, stability and robustness of such models. 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, etc. 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.

VL1: Fundamental Deep Learning

Project Leader
Vijay Kumar

Bio

Vijay Kumar is a Research Fellow based at the University of Adelaide. Prior to this, he was a Lead Engineer with Samsung Research Institute, India (in the Advanced Research Group). He received his Ph.D. in Computer Science from Queen Mary, University of London in June 2013 and MS in Electrical Engineering from Indian Institute of Technology, Madras in 2009. In his thesis, Vijay focused on developing supervised dictionary learning methods for action detection and action recognition. His research interests are Machine Learning, Computer Vision with current focus on feature learning and zero-shot recognition. He was a co-presenter of the tutorial "Deep learning and its applications in computer vision" at the 2016 Digital Image Computing: Techniques and Applications (DICTA) conference.

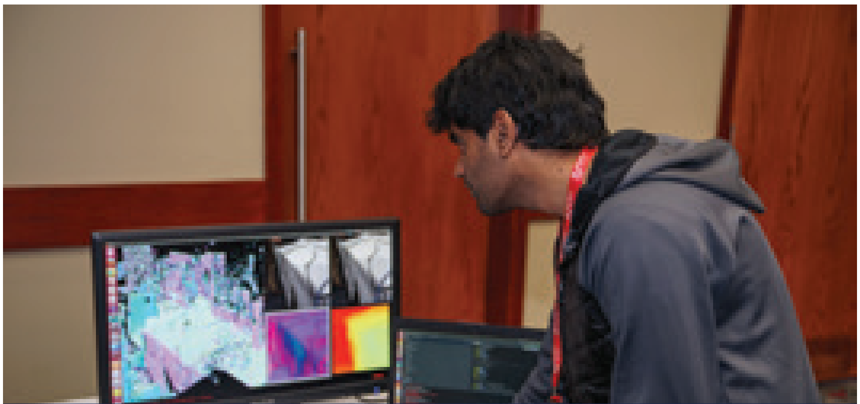

Demonstration of scene segmentation at RoboVis 2016 demo day

VL1 Research Team

Ian Reid (CI), Gustavo Carneiro (CI), Chunhua Shen (CI), Vijay Kumar (RF), Guosheng Lin (RF), Basura Fernando (RF), Jian (Edison) Guo (PhD), Ben Harwood (PhD), Yan Zuo (PhD), Zhibin Liao (PhD), Rafael Felix Alves (PhD)

VL1 Project Description

It is essential that the Centre be active at the forefront of current machine learning techniques. To explore, develop and exploit novel network architectures; to develop detection and instance level/pixel level annotations for 1000s classes and open sets of classes. To develop efficient and/or weakly supervised and/or online trained and/or unsupervised and/or zero-shot learning models. Active learning with and from temporal data.

VL2: Learning for Robotic Vision

Project Leader

Chunhua Shen

Bio

Chunhua Shen is an ARC Future Fellow (awarded in 2012) and Professor of Computer Science at University of Adelaide. He is part of the Australian Centre for Visual Technologies (ACVT). His research and teaching focus is on Statistical Machine Learning and Computer Vision.

Chunhua was awarded his PhD from the University of Adelaide; then worked at NICTA's computer vision program for six years. From 2006 to 2011, he held an adjunct position with the College of Engineering & Computer Science at the Australian National University. He moved back to Adelaide in 2011.

VL2 Research Team

Ian Reid (CI), Chunhua Shen (CI), Gustavo Carneiro (CI), Guosheng Lin (RF), Vijay Kumar (RF), Ben Meyer (PhD), Bohan Zhuang (PhD), Ben Harwood (PhD), Tong Shen (PhD)

VL2 Project Description:

Learning that is specific to robotic vision tasks where there are resource constraints (embedded vision system). Video segmentation (i.e. image segmentation for video, applied to static scenes / moving camera, and general scenes with unknown motion; DL suitable for deployment on storage and power constrained embedded systems (eg COTSbot); Fast, approximate, and asymmetrically computed inference; Robust inference (via understanding failure modes); Unsupervised learning; Online and lifelong learning for robotic vision; "Any-time" algorithms.

We have developed a unified image processing technique based on Deep Learning, which can be applied to various image enhancement tasks such as, but not limited to, image denoising, image inpainting, image super-resolution, image deblurring. These two images demonstrate an example of recovering clean photos from defocused photos.

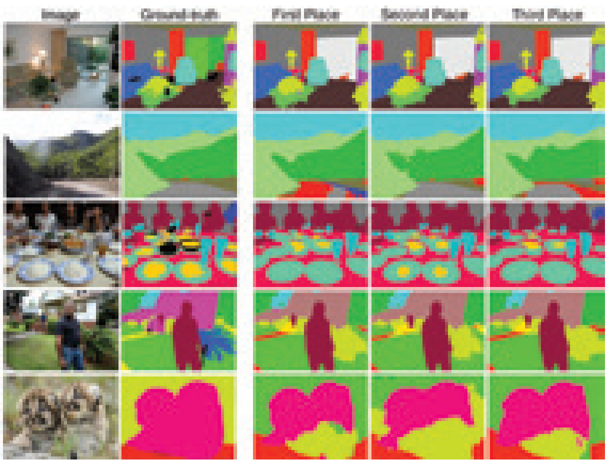

Top three image results for the ImageNet Scene challenge. Centre researchers placed second

CASE STUDY: ImageNet Scene Parsing Challenge

ImageNet is one of, it not the, most hotly contested challenge in Computer Vision each year. It attracts entries from most of the major groups in the field, both commercial and academic. Zifeng Wu, Chunhua Shen and Anton van den Hengel, placed second in this prestigious competition, a great result in this incredibly competitive challenge, particularly as it was achieved by developing a new architecture, which outperforms the very deep networks. That is, although our ensemble of networks was outperformed (slightly) by the team that won, our single model performance was the best. It is this single model performance which gives the best indication of the capacity of the architecture, which is why we're so enthusiastic about the result. The architecture we used was wide, rather than deep. This has a number of advantages, most of which reflect the fact that there is better parallelism in a wide network than a deep one.

TECHNOLOGY - Algorithms & Architecture (AA)

Program Leader
Hongdong Li

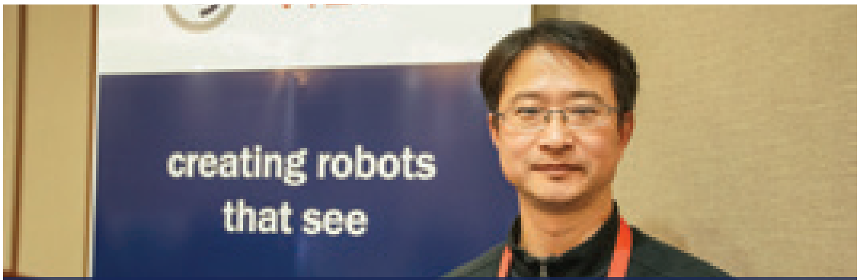

Program Leader Hongdong Li

Bio

Hongdong Li is a Chief Investigator based at ANU. His research interests include 3D computer vision, machine learning, robot visual navigation and SLAM, autonomous vehicles and self-driving cars. He has served as an Area Chair for most leading computer vision conferences such as CVPR, ICCV, ECCV and BMVC. He is an Associate Editor for IEEE Transactions on PAMI, IVC (Image and Vision Computing, journal), IET-Computer Vision, and IPSJ transactions on CVA (computer vision and applications). Hongdong has published over one hundred papers in major computer vision journals and conferences. He was a recipient of the prestigious IEEE CVPR Best Paper Award, the ICPR Best Student Paper Award, the ICIP Best Student Paper Award, and several other best paper awards. His research is funded by the Australia Research Council (ARC), ARC Centre of Excellence program and global ICT companies.

Program Description

This program aims to create advanced algorithms and techniques to allow computer vision to be run in real-time on robotic systems deployed in large-scale real-world applications, using distributed sensing and computation resources and to provide efficient and unified software platforms for real-time robot visual SLAM algorithms and techniques development and employment in real-world environment. There are three research projects (AA1, AA2, and AA3) under this program, each addressing a significant aspect of robotic vision research, development, and applications. AA1 (VOS) will provide a common, distributed computational platform that takes advantage of distributed sensing and computational capabilities to solve large complex robotic problems. AA2 (ACRV-SLAM) is focused on the development and integration of robot vision algorithms in robust vision, real-time vision and semantic vision areas, into a single SLAM-centred robot navigation framework. The framework will be demonstrated in real-world robot applications including AUV (autonomous underwater vehicle), UAV (unmanned aerial vehicle, or flying robot), and ground-based autonomous vehicles. AA3 (SIMUL) aims to provide a photorealistic graphics simulation environment to facilitate and accelerate the development of advanced robot vision algorithms and systems.

AA1: VOS-distributed robotic vision

Project Leader
Tom Drummond

Bio

Professor Drummond is a Chief Investigator based at Monash. He studied a BA in mathematics at the University of Cambridge. In 1989 he emigrated to Australia and worked for CSIRO in Melbourne for four years before moving to Perth for his PhD in Computer Science at Curtin University. In 1998 he returned to Cambridge as a post-doctoral Research Associate and in 1991 was appointed as a University Lecturer. In 2010 he returned to Melbourne and took up a Professorship at Monash University. His research is principally in the field of real-time computer vision (ie processing of information from a video camera in a computer in real-time typically at frame rate), machine learning and robust methods. These have applications in augmented reality, robotics, assistive technologies for visually impaired users as well as medical imaging.

AA1 Research Team:

Peter Corke (CI), Tom Drummond (CI), Vincent Lui (RF), Will Chamberlain (PhD), Steven Martin (Research Engineer)

AA1 Project Description

The goal of AA1 is to create a Vision Operating System that provides a framework for bringing together multiple sensing and computational resources to solve complex robotic vision problems. This will enable robots to make use of external sensing resources (e.g. CCTV cameras in the environment, or sensors mounted on other robots) as well as computation resources (either attached to those sensors, or provided as a large computing resource within the network). This kind of framework enables novel solutions to complex problems in which the various resources are combined collaboratively to solve complex localisation, navigation, understanding and planning problems.

AA2: ACRV SLAM Framework

Project Leader
Viorela Ila

Bio

Viorela Ila is a Research Fellow based at the Australian National University (ANU). Her research is in the field of simultaneous localization and mapping (SLAM) and 3D reconstruction.

Viorela received a Ph.D. degree in Information Technologies from the University of Girona, Spain, in 2005. After her PhD studies, she joined the Institut de Robótica i Informàtica Industrial, Barcelona, Spain where she led two Spanish national projects and participated in URUS EU FP7 project. In 2009 she received a MICINN/FULBRIGHT post-doctoral fellowship which allowed her to join the group of Prof. Frank Dellaert at College of Computing, Georgia Tech, Atlanta US. In 2010, she joined LAAS-CNRS, Toulouse, France to work in the ROSACE project founded by RTRA-STAE. Between 2012 and 2014 she was with Brno University of Technology in Czech Republic working in three EU and one national projects.

AA2 Research Team

Richard Hartley (CI), Hongdong Li (CI), Tom Drummond (CI), Ian Reid (CI), Laurent Kneip (AI), Andrew Davison (PI), Frank Dellaert (PI), Marc Pollefeys (PI), Viorela Ila (RF), Yasir Latif (RF), Vincent Lui (RF), Feras Dayoub (RF), Mina Henein (PhD), Thanuja Dharmasiri (PhD), Andrew Spek (PhD), Yi Zhou (PhD), Sean McMahon (Affiliated PhD Student)

AA2 Project Description

This project will develop novel SLAM algorithms which can perform in challenging environments (large-scale, dynamic, dense, non-rigid). ACRV-SLAM is a common framework that integrates efficient implementations of the proposed algorithms with the goal to facilitate their distribution to the robotics community and the industrial partners, and to produce high quality demonstrators.

AA3: Computer graphics simulation for robotic vision

Project Leader
Peter Corke

Bio

Peter Corke is a Professor of Robotic Vision at Queensland University of Technology, and Director of the ARC Centre of Excellence for Robotic Vision. His research is concerned with enabling robots to see, and the application of robots to mining, agriculture and environmental monitoring.

AA3 Research Team

Peter Corke (CI), John Skinner (PhD), Steven Martin (Research Engineer), Niko Sünderhauf (Affiliated RF), Trung Thanh Pham (Affiliated RF)

Project Leader Peter Corke (r) with PhD researcher Saroj Weerasekera

AA3 Project Description

The performance of a robotic vision system depends on the initial state of the robot and the world it perceives as well as the lighting conditions, unforseen distractors (transient moving objects) and unrepeatable sensor noise. A consequence is that no robotic vision experiment can ever be repeated and the performance of different algorithms cannot be rigorously and quantitatively compared. For machine learning applications a critical bottleneck is the limited amount of real world image data that can be captured and labeled for both training and testing purposes.

This project investigates the potential of photo-realistic graphical simulation based on state-of-the-art game-engine technology to address both these challenges.

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

Robots for...

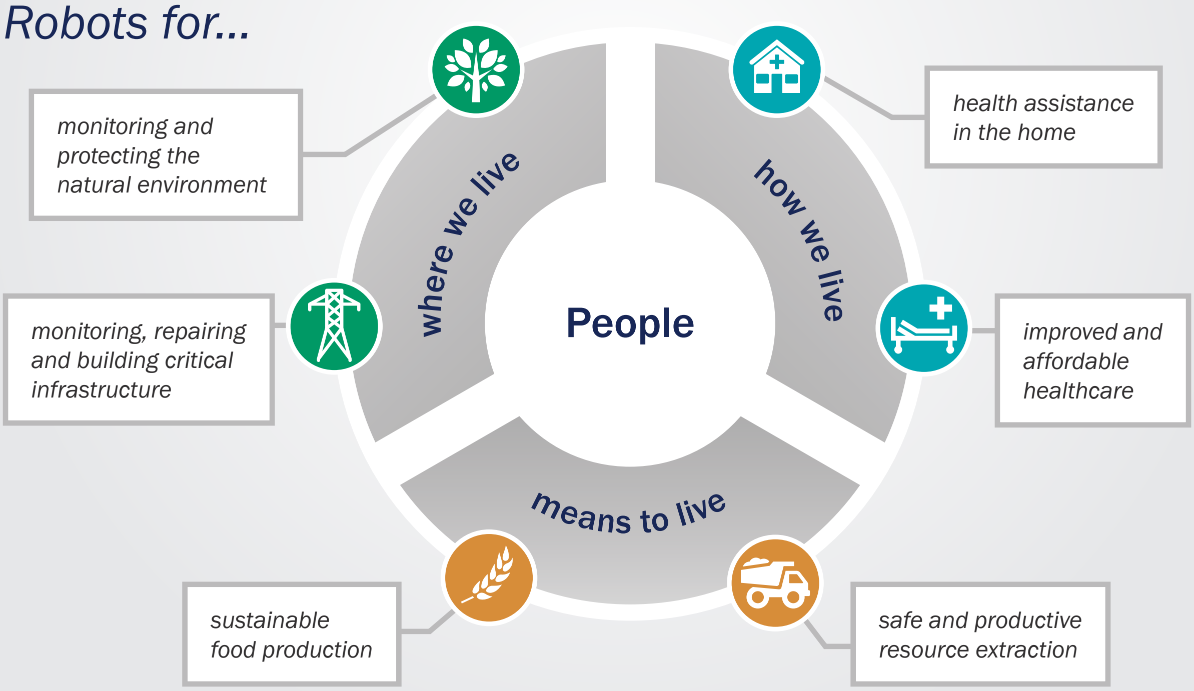

Where we live

The environment in which we live is critical to our wellbeing. This environment comprises both the natural environment and the built environment.

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 and 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 the cost of inspection, allowing more of the world to be inspected, more frequent inspections, or both. Ultimately, intelligent vision-capable robots would be capable of environmental remdiation as well as inspection.

The built environment, our cities, roads, dams, tunnels, pipelines and electrical distribution networks represent a massive investment over generations that needs to maintained and enhanced and passed on to future generations. Maintaining these assets in good condition requires continuous inspection and traditional methods are labour intensive and thus 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.

CASE STUDY: CRCSI PROJECT WITH ERGON ENERGY

It's a daunting task – having to maintain a network of hundreds of thousands of miles of power lines and the poles that support them. That's the challenge facing Queensland energy giant, Ergon.

Centre researchers are now part of an effort to help Ergon with its challenge. They are involved with a new CRCSI and QUT collaboration with Ergon to develop a vision and sensing platform for UAV's to fly over all the power poles and capture images. They are looking for things like rot in a pole's cross-arms, along with other possible structural issues.

Right now, the work is being done by UAV operators, who must keep the airborne vehicles within their line of sight. It's hoped this project will make their job much easier, by using on-board sensors to develop a force field, if you will, that will keep the UAV's from getting too close to the poles or wires.

Centre Associate Investigator Jason Ford is among those involved with the project, which is an example of the type of research being done to tackle issues with infrastructure.

"This project was empowered by the Centre existing and provides an example of the type of work the Centre can do," Jason said.

Other Centre members involved with the project include Centre Director Peter Corke, Research Fellow Feras Dayoub, and PhD Student Steven Martin. Other researchers involved include Andrew Keir, from QUT's Institute of Future Environments and QUT Research Fellow Aaron Mcfadyen.

In 2016, Ergon merged with Energex, with the new company called Energy Queensland.

Inspection of infrastructure is an ongoing challenge that can be partially addressed using robotic vision

Maintenance of solar arrays using robotic vision is the subject of an ARENA grant to Centre researchers.

CASE STUDY: HELPING MAINTAIN AND SERVICE AUSTRALIA'S SOLAR INDUSTRY

Researchers at the Australian National University (ANU) were successful in the 2016 Australian Renewable Energy Agency (ARENA) grant outcomes.

The project titled "A Robotic Vision System for Automatic Inspection and Evaluation of Solar Plant Infrastructure" received direct funding of $876,000 from ARENA with a total project value of $3 million with in-kind contributions from Fotowatio Renewable Ventures (FRV), Vast Solar, Data61, and 4D Surveying.

The project was initiated by Centre Associate Investigator Fatih Porikli (ANU) and Chief Investigator Robert Mahony (ANU) in collaboration with Dr Evan Franklin (ANU) and Dr Joe Coventry (ANU) who bring renewables specific subject expertise to the project.

"This project is the first in the world where advanced robotic vision technology will be deployed to provide deep and comprehensive diagnostic information to solar plant owners," Rob Mahony said.

Rob says the rapidly expanding solar industry in Australia will lead to millions of square kilometres of photovoltaic panels and solar concentrating mirrors that must be maintained and serviced.

"Robotic vision has a key role to play in autonomous inspection to identify and monitor faults and performance issues such as soiling levels, allowing renewable infrastructure to be cheaply and efficiently maintained at peak efficiency," Rob said.

Rob also says the project draws on three key research themes:

- **Robust vision** – to deal with the high exposure environments.
- **Visual Learning** - in visually identifying faults and soiling levels.
- **Vision and action** – in control of the aerial drone.

Others involved in the project include Dr Naveed Akhtar (ANU), who is lead postdoctoral fellow along with Dr Salman Khan (Data61), and Mr Ehab Salahat (ANU) who is a centre affiliated PhD student working on the project.

How we live

The health of the population is key concern for any society. However the populations in many societies are ageing. By 2050, a quarter of Australia's population will be over 65 and this has significant economic costs. Firstly, the dependency ratio increases, meaning fewer working people to carry the burden of work in society, including caring for those no longer working. Secondly, the cost of healthcare increases and will at some point become unsustainable for national budgets.

Healthy independent living describes how people can continue to live at home, even if they require some medical assistance. The key to independent living is access to affordable assistive technology, that is, devices or systems that provide practical solutions to everyday activities. This is an area where robotics is already having an impact with the development of exoskeletons, care assistants, communications devices, and devices to improve mobility. As well as reducing the costs of healthcare, caring for one's self preserves a sense of dignity and freedom.

Hospital based healthcare costs can be reduced by the introduction of various automation systems, bringing all the advantages that automation has brought to other sectors: increased productivity, higher quality and removing workers from injurious work. Hospital automation extends beyond surgery to cover areas such as internal logistics, patient handling, rehabilitation and care.

CASE STUDY: ROBOTS IN HOSPITALS

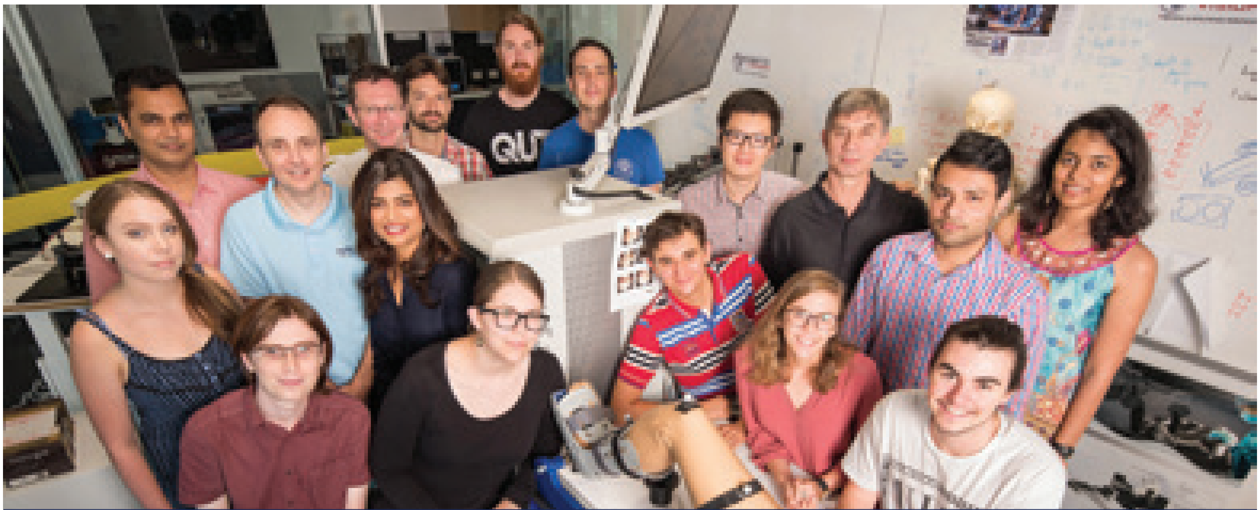

Medical Robotics team. Clockwise from far left: Lindsey Paul, Timothy Chant, Dr Anjali Jaiprakash, Kristine Davis, Dr Ajay Pandey, Professor Jonathan Robterts, Professor Ross Crawford, Dr Thierry Peynot, Douglas Palmer, Jeremy Opie, Dr Leo Wu, Andrew Razjigaev, Mario Strydom, Clare Villalba, Hardik Patel, Vibhavari Dasagi and Charles Grist. Not photographed: Dr Anders Eriksson, Thomas Coppin, Andres Felipe Marmol Velez and Jordan Laurie

Hospitals globally have been slow to adopt robotics and artificial intelligence into patient care, although both have been widely used and tested in other industries.

Medicine has traditionally been slow to change, as safety is at its core. Financial pressures will inevitably force industry and governments to recognise that when robots can do something better and for the same price as humans, the robot way will be the only way.

Medicine has long been segmented into many specialised fields but the doctor has been expected to travel with the patient through the full treatment pathway.

A surgeon, for example, is expected to be compassionate, and good at many tasks, such as diagnosing, interpreting tests, such as X-rays and MRIs, performing a procedure and post-operative care.

As in numerous other industries, new technology will be one of the drivers that will change this traditional method of delivery. We can see that one day, each of the stages of care through the hospital could be largely achieved by a computer, machine or robot.

There are programs to make diagnoses based on a series of questions, and algorithms inform many treatments used now by doctors.

Surgeons are already using robots in the operating theatre to assist with surgery. Currently, the surgeon remains in control with the machine being more of a slave than a master. As the machines improve, it will be possible for a trained technician to oversee the surgery and ultimately for the robot to be fully in charge.

Hospitals will be very different places in 20 years. Beds will be able to move autonomously transporting patients from the emergency room to the operating theatre, via X-ray if needed.

Triage will be done with the assistance of an AI device. Many decisions on treatment will be made with the assistance of, or by, intelligent machines.

Your medical information, including medications, will be read from a chip under your skin or in your phone. No more waiting for medical records or chasing information when an unconscious patient presents to the emergency room.

Robots will be able to dispense medication safely and rehabilitation will be robotically assisted. Only our imaginations can limit how health care will be delivered.

Source: Extracts from an article "Robots in health care could lead to a doctorless hospital" in The Conversation written by Anjali Jaiprakash, Jonathan Roberts and Ross Crawford.

Means to live

Modern life as we know it is resource intensive, requiring reliable and low-cost supply of food, energy and minerals. Australia is an abundant producer, and exporter, of food, energy and minerals. environment in which we live is critical to our wellbeing. This environment comprises both the natural environment and the built environment.

Food production today is ultimately reliant on human labour. The least intensive is extensive crop agriculture and dairy production which are highly mechanised and beginning to become automated. Animal agriculture, extensive grazing and feedlots, is not labour intensive but the work is physically hard and in remote locations. Horiculture, the production of fruit, vegetables, nuts and flowers, varies from highly mechanised to labour intensive, and again the work is physically hard and in remote locations. It is increasingly difficult to source workers to perform remote physical work which leads to 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, pollenate flowers and even muster cattle.

Mineral production, in particular mining, is highly mechanised and increasingly automated. The business drivers are increased productivity and capital utilisation, reduced machine damage and removing workers from hazardous environments. 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: Robotic Vision and Mining Automation

Researchers from the Queensland University of Technology (QUT) and the Australian Centre for Robotic Vision, in collaboration with Mining3, have been chosen to help Caterpillar take its mining equipment and automation technology to the next level.

The Queensland Government awarded a team led by Michael Milford, and including Affiliate Researcher Thierry Peynot, 428,000 in funding as part of its Advance Queensland Innovation Partnerships program.

The funding, combined with other funding from QUT, Caterpillar and Mining3, will help Michael, Thierry, and their team develop technologies that could ultimately enable the automation of underground mining vehicles.

Right now, lasers are being used to help with attempts to automate vehicles involved in underground mining operations. The aim of this project is to develop a camera-based positioning system on mining vehicles to help track them in these harsh, underground environments. That will make the work safer and more economic.

“If you know where everything is in a mine site at all times you will be able to optimise how the mine site operates and keep your industry competitive,” Michael said. “We hope that we can develop some next-generation positioning technologies which can be deployed throughout their fleet of vehicles that are all around the world.”

The Queensland Government funding was announced in September 2016 by State Development Minister Anthony Lynham and by State Minister for Innovation, Science and the Digital Economy Leeanne Encoch MP.

“The researchers will look to solutions based on developing a cost-effective, reliable camera-based positioning system for locating and tracking underground mining vehicles within one metre of accuracy as well as a sophisticated, multi-sensor system that provides centimetre-accurate positioning,” Ms Enoch said.

Milford says Caterpillar will give researchers access to some very good equipment. In collaboration with Mining3, they will also provide access to some exclusive mine sites.

“Regularly interacting with and talking with key companies and players in this space including Caterpillar, Mining3 and other companies has been invaluable in further shaping our research agenda moving forward to maximize its relevance to what industry needs now and into the future,” Michael said.

Michael says some of the systems his team will be working on are already automated. Their job will be to improve their reliability. Michael also thinks this type of research and technology could have wider implications outside of mining.

“It will be interesting to see how technology development in mining will interact with and complement that being developed in other major fields such as self-driving cars,” Michael said.

This was one of 15 projects to receive an Advance Queensland grant.

Harvey, a prototype robotic platform for harvesting capsicums, being tested at a Queensland Government protected cropping facility by Chris Lehnert

CASE STUDY: Harvey the Capsicum-picking robot

Sustainable food' production is an area where the Centre is already making an impact.

An example of this is a project that involves a robot nicknamed Harvey. Harvey harvests Capsicums, and he's doing it better than ever before.

Harvey was created in 2015, as part of the Queensland Department of Agriculture and Fisheries (DAF) three-year Strategic Investment in Farm Robotics (SIFR), which identified capsicum (bell or sweet peppers) as an important Queensland crop that could benefit from robotic harvesting.

The SIFR team led by Centre AI, Professor Tristan Perez, and Centre Research Affiliate Chris Lehnert, developed their agricultural robot prototype in 2015. In November 2015, the team conducted the first trials of 'Harvey' at a Queensland Government protected cropping facility in North Queensland. Harvey was tasked with identifying and picking red capsicums. Harvey achieved a 50 per cent success rate with unmodified crops (this is, with no leaves removed or fruit moved before harvesting).

In 2016, the team made several advancements in Harvey. It tested a new grasping approach, which dramatically boosted grasping performance. The team also trialled different varieties of capsicums, but that led to a new challenge for the team involving Harvey's vision.

"What we learnt from there was that detecting the peduncle is a big challenge. That's the part of the capsicum that attaches to the plant. We realised this was an issue that required more attention, so we developed a new method for peduncle detection," Chris said.

Centre Research Affliate Dr Chris McCool developed the peduncle detection method using a deep neural network.

The team also had two more field trials for Harvey in 2016. The first was again in North Queensland. The second was at a facility in Cleveland, in southeast Queensland, where the team actually grew capsicums from scratch. Chris says the latest field trial showed a success rate of 75 per cent.

Harvey is also getting faster. The time it takes for Harvey to pick a capsicum dropped from 35-40 seconds down to 20 seconds. Commercial viability is 10 seconds, so the team feels like it's getting closer to reaching that goal.

Moving forward, the team also has two other goals. To come up with a new design to make it more robust for commercial use. The other is to determine what other crops it should also consider harvesting.

How does Harvey work?

Harvey's robotic arm has a camera and a unique harvesting tool attached to it. Using data from the camera, the robot creates a 3D model of each fruit and its surroundings and plants and controls the robotic arm and cutting tool as they locate and detach the fruit. The combination of state-of-the-art robotic vision software and novel and crop-manipulation tools enables successful harvesting of the crop and promises significant benefits for horticulture growers, who export (export sounds like to international customers) more than $2b in products every year.

The Centre partnered with Robotics WPS and Scienceworks to deliver robotic vision workshops to high school teachers in Queensland and Victoria. Here Michele Miller from Robotics WPS is presenting to teachers in Brisbane. Our workshops were oversubscribed within days of launching.

National Benefit

National Research and Innovation Priorities

Robots are already widely used in Australia on factory floors (automotive, food manufacturing), and in the field (autonomous mining vehicles, autonomous port container vehicles). However, robots are notably absent from many other sectors. Where are all the robots? Without vision, a vast array of potential applications is closed to robots, for example: complex manual assembly, packing, manipulation, navigation, machine operation, fruit picking, crop spraying, remote assistance, smart homes, smarter appliances, autonomous driving, environmental surveying and monitoring.

Robotic vision technologies developed by our Centre will deliver important economic, environmental and social benefits for Australia.

Robots that see will be a transformative technology that offer solutions to important economic and social problems facing Australia in this century such as, a changing climate, and a growing, ageing population where there are half as many workers as those that have retired from the workforce. Productivity growth can improve people's lives, lifting living standards, and create wealth and wellbeing. We expect robots to help provide solutions to meet Australia's future challenges, in line with the nation's Science & Research Priorities including:

- labour shortages and low productivity growth in key industries;
- diminishing international competitiveness due to high-wage rates;
- rising OH&S compliance costs;
- the need for increased productivity due to an ageing population;
- rising healthcare costs and demands for customised healthcare solutions;
- ageing infrastructure and asset base;
- growing demands for minerals, energy and food; and
- the need to preserve the environment.

In April 2015 the Australian government defined a set of national Science and Research Priorities that we have mapped our research program against (See SRP breakout box), identifying overlap in the areas of Food; Soil & Water; Transport; Advanced Manufacturing, Environmental Change; and Health. Many of these priority areas are addressed, in varying degrees, by the Centre's work in robotic vision. Increased application of robotic technology facilitated by the enhanced capability of robots (through robotic vision) will help build and maintain resilient human and natural environments. Robots are increasingly being used to promote health and well-being and improve healthcare delivery. The application of robotics to increase agricultural productivity and food processing capabilities, and to also improve infrastructure management (e.g., water storage and transportation) is important to the long-term sustainability of Australia's precious soil and water assets. Novel robotic vision techniques developed by the Centre may be applied not only to robotics but will also help safeguard Australia from a range of security threats.

Finally, current trends suggest robots are key to lifting productivity in Australia's resources, services and manufacturing industries, and will support the development of new industries while fostering the development of an entrepreneurial and innovative knowledge economy that will benefit Australia.

Our Centre clearly addresses the National Innovation Priorities through: undertaking world-class research grounded by national challenges around productivity and competitiveness; training a generation of experts in robotics and vision who will work in industry, government and academia; translating research results in robotics and vision to important future industries and new companies through the transfer of trained people; creating awareness of the Centre's technologies through effective communications leading to a variety of engagement models with enterprises from small to large; active collaboration with others in the national innovation system including non-partner universities and organisations such as DST Group as well as with industry through collaborative or contract research; and strong international engagement with top international universities in the field that can be leveraged by Australian industry.

National Benefit KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Contribution to Australia's Strategic Research Priorities **The Centre will contribute to:** • Living in a changing environment • Promoting population health and wellbeing • Managing our food and water assets • Securing Australia's place in a changing world • Lifting productivity and economic growth **Percentage of publications relevant to SRPs**	Annually	80%	80%	80%	100%	100%	100%
Measure of expansion of Australia's capability in the priority area(s)	At review	24 postdoctoral fellows trained, 60 PhD students trained, 390 honours students involved, 10 additional staff attracted to partner universities.					
Measures of international reputation and competitiveness	At review						
• Number of awards		1	13	2	17	3	10
• Number of fellowships		2	3	3	3	3	1
• International grants		0	1	1	0	2	1

Science & Research Priorities (SRPs)

The Australian government identifies nine science and research priorities along with associated practical challenges:

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, 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 ecosystems
- 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 cybersecurity
- 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
- crosscutting technologies that will de-risk, scale up, and add value to Australian manufactured products
- specialised, high value-add areas such as high-performance materials, composites, alloys and polymer

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 health care 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 health care

End-User Engagement

"It is not enough for the Centre to do fantastic, internationally-impactful science if it fails to translate research results into tangible benefits to end-users." - Professor Peter Corke, Director.

Our Centre is committed to increasing our network of end-users and to conduct research that attracts investment in the form of additional research income (from research contracts, commercial contracts, licensing/selling Centre IP). We aim to build sustainable partnerships across the research sector and public and private enterprises. Linkage targets, our End-Users, cover a range of sectors, with our key focus aligned to our impact areas in Infrastructure Monitoring & Asset Inspection, Agriculture, Aquaculture and Bio-monitoring (Environment), Building & Construction, and Medical & Healthcare. These areas have been identified as economically important to Australia with high potential for application of robotic vision technologies. The Centre is on track to attract an additional $10 million in investment to develop prototypes that have impact in these sectors.

The Centre's End-User engagement strategy guides the Centre in taking a leadership role in developing the fledgling robotic vision industry in Australia. We seek to develop links and potential partnerships with those companies already in the robotic vision space in Australia.

In general, our impression is that the robotics and computer vision (RV) industries in Australia are fledgling and fragmented. It is not the Centre's role to compete with these commercial enterprises so we have conducted a first pass audit of Australian RV companies to help gauge those we can partner with to deliver technological solutions We are leading the development of a roadmap defining the future for a robotic vision industry in Australia. As robotic vision is a technical infrastructure innovation, it can impact on all industries. For this reason, our outreach to industry will have many facets. End-users and potential end-users will be invited to the Centre's annual symposium technology demonstrations, as well as local events such as demonstrations held at least annually at each of our nodes. On the advice of our End-User Advisory Board we are also hosting sector-specific industry workshops to identify the challenges faced by many businesses, which could be addressed using robotic vision.

Centre COO Sue Keay explains the Centre's research capability to an audience from the manufacturing sector

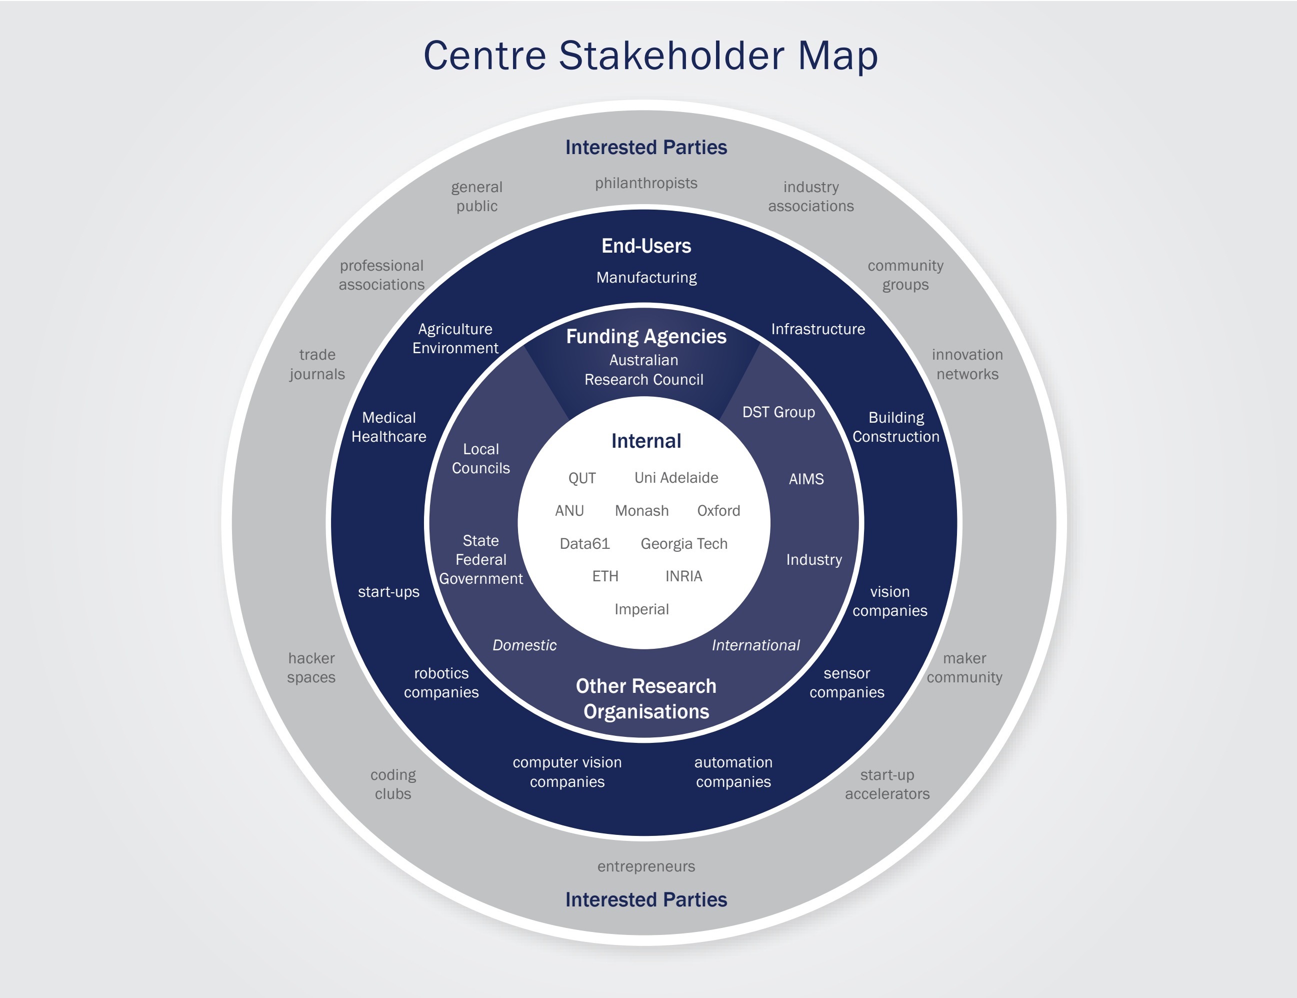

Identification of end-users mapped versus intimacy with the Centre where those closest to our influence (internal) are employees and legally bound partners, closely linked with research collaborators and funders, then we are cultivating relationships with potential end-users, particularly from our target industries and in the outer circle, interested parties – a diverse group with interests aligned with the Centre.

We are partnering with PwC on their innovation challenge platform (see CASE STUDY p. 48) to work with providers to find appropriate robotic vision solutions to these sector-specific challenges. Further engagement occurs through the general media, articles in trade journals, as well as general community engagement channels (media coverage, website, social media, printed material). In 2017 a targeted campaign to raise the profile of the Centre will commence, with Centre articles feeding in to trade journals and Centre researchers speaking at industry-specific conferences. Where significant industry demand is identified the Centre will run workshops in robotics, vision and/or robotic vision. Industry engagement is critical to the success of our Centre. Members of the Centre will continue to host one-on-one meetings with members of various industry segments related to our impact areas (See Research Impact p. 36). Information about these meetings is tracked using the Centre's customer-relationship management (CRM) system, which forms the backbone to our Engagement strategy recording all touchpoints with industry representatives.

We have delivered 64 government, industry and business community briefings and 17 public awareness and outreach programs. In 2016 we hosted visits from Australia's Chief Scientist, Dr Alan Finkel, the Chair of CSIRO's Board, former Telstra CEO David Thodey, Queensland Chief Entrepreneur Mark Sowerby, former Queensland Premier Peter Beattie, and top executives and representatives from more than 20 different organisations such and, amongst many visits by industry representatives, including from as Suncorp, Stryker, The US Navy, Queensland Health and Brisbane City Council (see our KPIs).,We also hosted visits from then federal MP, Ian Macfarlane (now CEO of the Queensland Resources Council), Senator Deborah O'Neill (Shadow Assistant Minister for Innovation & Shadow Assistant Minster for Mental Health) and the Queensland Minister for Agriculture and Fisheries The Hon Leanne Donaldson MP.

End-User Links KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Number of government, industry and business community briefings	Annually	10	79	10	54	10	64
Number and nature of public awareness/outreach programs Public lectures, media pieces	Annually	4	40	4	12	12	17
Currency of information on the Centre's website *Number of website updates annually	Annually	12 updates	60 updates	12 updates	116 updates	12 updates	36 updates
Number of website hits	Annually	50K	34K	200K	98K	300K	110K
Number of talks given by Centre staff open to the public	Annually	5	10	10	12	10	19

Research Fellow Guosheng Lin demonstrates his research to PhD researcher Dorian Tsai at RoboVis2016

AgBotII and Harvey, are prototype robotic platform for agriculture developed as part of a three-year Strategic Investment in Farm Robotics (SIFR) funded by the Queensland Department of Agriculture and Fisheries (DAF).

CASE STUDY: INNOVATION CHALLENGES AND A ROBOTIC VISION ROAD MAP FOR AUSTRALIA

PwC's Open Innovation Platform is an online marketplace designed to connect governments and businesses to a network of thousands of innovators from small to medium enterprises to solve important problems. With the support of PwC (a valued member of our End-User Advisory Board), we're running a series of innovation challenges aimed to identify people and companies in the Australian robotics and vision communities.

We have two aims:

1. to identify companies we can collaborate with on solving challenges, partners who we can work with or already have solutions to some of the problems we are asked to address, and companies who might be interested in working with us to commercialise our technologies in the future.
2. to engage with companies who broadly operate in the automation/robotics/sensing/vision space to collectively develop a robotic and vision roadmap for Australia and raise awareness of the important economic benefits of the industry to Australia.

We see the innovation challenges as an opportunity to partner on robotic vision solutions for different sectors and to raise the profile of the robotic vision industry.

Our first two challenges are focussed on robotic vision solution providers in the manufacturing and food production sectors.

The Centre is looking to engage with all companies in Australia broadly operating in the robotics, computer vision, automation and sensing space as well as manufacturers of robotics, vision and sensing equipment, tech companies and industry bodies.

Entrepreneurship and Commercialisation

We encourage entrepreneurship within the Centre.

Some examples include: research engineer, Steve Martin attended QUT's Ubercamp, Centre COO Sue Keay interviewed Sphero founder Ian Bernstein at a live event which was webcast, PhD student James Mount started shipping out his Kickstarter robots to Australian Schools.

Two of the Centre's nodes were visited by Ben Swan from Muru-D (Telstra's start-up accelerator), two special seminars on start-ups were held at QUT and the Centre also hosted a visit from Queensland Chief Entrepreneur, Mark Sowerby.

Protecting and developing Centre IP is a fundamental objective of the Centre. We have guidelines/policies on intellectual property and open-sourcing, and are developing a register of intellectual property developed as part of the Centre's Activities, as well as any background IP dependencies. All personnel involved in Centre Activities assign their IP to their host university (project party). Ownership of Intellectual Property amongst project parties will be decided based on the terms of the Centre Agreement. Project parties consult with the Centre Executive over protection and commercialisation of Project Intellectual Property.

So far only one piece of commercialisable technology has been identified and is the subject of a patent application being led by ANU. Further impact from our research is encouraged by our open sourcing policy. Open sourcing in one way to gain impact from Centre Intellectual Property, and the Centre recommends the BSD 3 clause license unless the IP has commercial potential, in which case, an alternative license and/or IP protection strategy will be employed. The Centre newsletter promotes the different open source code and code repositories created by Centre researchers.

CASE STUDY: GAME-CHANGING TECHNOLOGY WITH NVIDIA

If we want robots to be able to do things in the real world, they need to be able to react to what they're seeing in real time.

Thanks to a new collaboration with NVIDIA, the Australian Centre for Robotic Vision now has the computing power it needs for its research to help robots learn to see.

The central point of this collaboration is taking place at one of our four partner universities - Monash University in Melbourne. NVIDIA in 2016 made Monash a GPU Research Centre.

NVIDIA, of course, is best known for designing Graphics Processing Units (GPUs) for the gaming market, and has a significant share of the GPU market. About 10 years ago, scientists worked out there was more computing power inside the graphics cards they were buying from NVIDIA, than the computers they were putting them into.

"The cards are about 10 to 20 times faster than a whole computer," says Prof Tom Drummond, a Chief Investigator with the Australian Centre for Robotic Vision at Monash.

It's that type of computing power that's needed to help robots see and react in real time.

"With a single-threaded code going from a Conventional Processing Unit (CPU) over to GPU, we may see a 150-fold speedup in computing. That's a game changer," says Prof Drummond. "It means we can do things in robotics, which has to be real-time, that we couldn't otherwise."

In addition, the NVIDIA collaboration provides some key hardware for the Centre. For instance, there is a collaboration between the Monash and QUT nodes to develop a Vision Operating System. The system will enable large sets of robots and cameras to contribute resources together to solve big problems. NVIDIA is helping make that happen with the donation of its TX1 developer boards. TX1 has the power to process image feeds, and the software tools to instantly analyse and provide context to what's being seen.

Finally, NVIDIA is also providing access to its education program so Centre researchers can learn how to use the hardware. That's beginning at Monash, but is expected to move to the other Centre nodes at QUT, Adelaide and the Australian National University (ANU).

Communication & Media

Our communication activities are wide-ranging. We conduct a number of engagement activities including forums, workshops and presentations. We actively promote our Centre through traditional media channels that target local, national and international news, current affairs programs, and other media outlets.

Our Centre also promotes its activities online and through social media channels including Twitter, Facebook, Google+, LinkedIn, Flickr, Instagram and YouTube. Our website (roboticvision.org) targets a cross-section of audiences, providing information about the Centre, access to resources, research services, and downloads of research, teaching and educational tools. As it evolves our website will also host general interest material such as videos, interactive technology demonstrators, open online educational material suitable for self-directed study by interested students; as well as information for external researchers such as published papers, reports and open-source software. Our annual report highlights the achievements of our researchers and is available in electronic form and distributed in hardcopy to selected stakeholders.

Our Centre's researchers and research garnered a lot of media attention in 2016 as we continued to engage with people about the potential of robotic vision technologies to transform the world. Centre Associate Investigator Matt Dunbabin and Research Fellow Feras Dayoub were the focus of numerous articles and interviews about their COTSbot project, and for later being a part of the winning Google Impact Challenge team with their RangerBot platform. Our agricultural robot, Agbot II, received a lot of attention, along with the researchers involved with the project. The Rangerbot, AgBot II and research affiliate Felipe Gonzalez's work tracking koalas with AI-equipped drones, were all featured in Gizmodo's "The Best Australian Science Discoveries of 2016". Research Fellow Juxi Leitner was interviewed about LunaRoo by Ron Vanderkley (listen to the Podcast on Robohub) and Juxi also appeared on Channel 7's Scope program along with Chris Lehnert and Owen Bawden talking agricultural robotics.

Our medical robotics team of Associate Investigators Jonathan Roberts and Ross Crawford, along with Research Fellow Anjali Jaiprakash, were in the news for Australia's first robotic help in a hip replacement operation. The team had three articles in The Conversation talking about robotics and healthcare (see CASE STUDY p. 39). Jon Roberts had an additional five articles in The Conversation examining the future of robotics in society. NVIDIA introduced the world's first Deep Learning Supercomputer in Australia, and that brought attention to Centre Chief Investigator Tom Drummond and our node at Monash University. Other issues that brought the Centre attention included self-driving cars, driverless mining trucks, and drones. In all, the number of radio and TV interviews combined with online articles exceeded 80 this year. Many of those articles were picked up by multiple online platforms.

Our Centre has also been active in the debate about robots, jobs and society with Centre researchers invited to give public lectures, participate in panel discussions and Director Peter Corke was interviewed about robots and jobs for the Australian Financial Review. A number of our Chief and Associate Investigators are in high demand for public appearances as evidenced by our key performance indicators, which show the high level of interest generated by our Centre in terms of the number of meetings, public lectures, outreach and media activities.

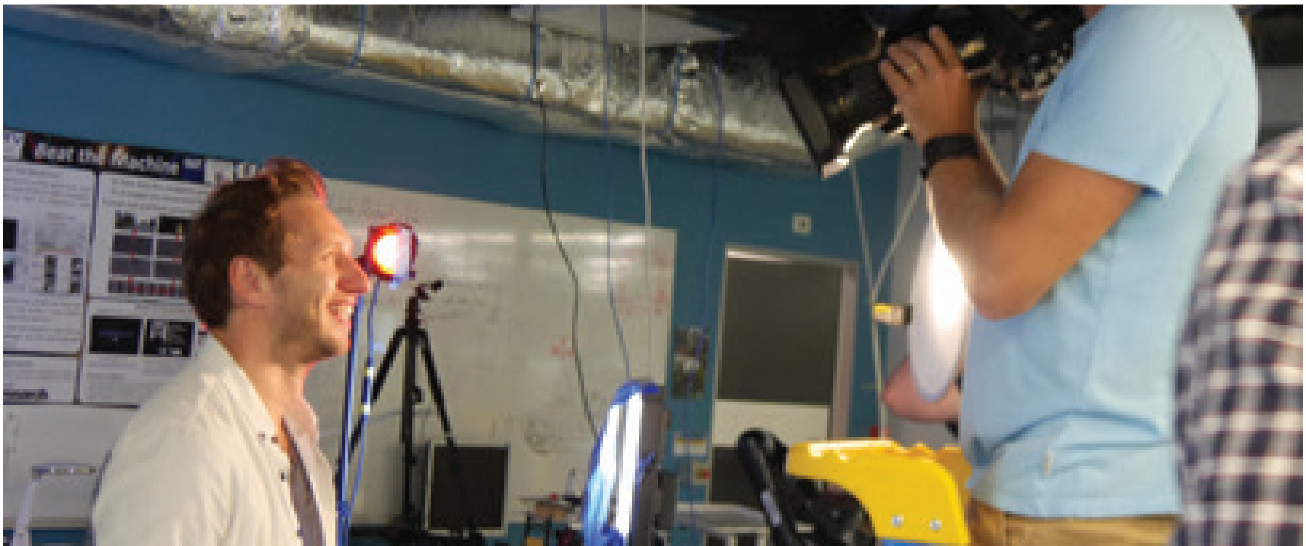

Research Fellow Juxi Leitner is interviewed about his plans to land a robot on the moon, LunaRoo, by Scope TV.

CASE STUDY: INTRODUCTION TO ROBOTICS 3-PART COURSE OPENED ON UK-BASED FUTURELEARN PLATFORM

In 2016, Centre Director Peter Corke, working with QUT Learning Decisions, repackaged some of his signature Introduction to Robotics courses into a number of shorter courses for release on the UK-based FutureLearn platform. The courses will be accessible to a wider audience with reduced technical depth, due to their shorter course length. The first course in the Introducing Robotics program tackles the basics of what robots are, what they are not, why we need them and the implications of robots in society.

Centre Director Peter Corke with Mobie, has made his online learning material available on FutureLearn and is looking to develop a fully fledged Robot Academy.

Outreach

Our Centre is actively building robotic vision capacity at the high school, undergraduate and graduate level, as well as making robotic vision accessible to the general public. Numerous outreach activities have been hosted by the Centre including involvement in the Moreton Bay STEMfest, collaboration with Coding for Kids, CoderDojo and Queensland's Chief Scientist, numerous school visits, tours by schools, future leaders and business groups, MIT Innovation. The Centre also ran oversubscribed professional development courses in coding and robotic vision for school teachers in association with RoboticsWPS (Brisbane) and Science Works (Melbourne).

Our community outreach program includes an innovative undergraduate curriculum, as well as the organisation of an annual intensive residential summer school for students from across the world at the graduate level to connect the largely disjoint fields of robotics and computer vision and undertake hands-on project work. Our second Robotic Vision Summer School (RVSS 2016) was again held at ANU's Kialoa campus in 2016 (see CASE STUDY p. 53).

To make robotics accessible to everyone, Centre Director, Peter Corke, developed two 6-week robotics MOOCs (massive, open, online courses).

The first of these, "Introduction to Robotics", which covers the fundamental principles of robotics with a focus on robotic manipulator arms, has been converted into a three-part course and made available on the UK-based FutureLearn platform with thousands of enrolments by students from over 150 countries.The courses were developed at QUT with the support of MathWorks and Springer. The courses will continue to be refined and delivered in the future, and the development of a robotic vision resources hub is underway. Other individual efforts in outreach include CIs Stephen Gould and Hongdong Li's outreach programs at ANU, including STEM workshops and robotics competitions, while CI Michael Milford has launched an educational mathematics venture "Maths Thrills" to encourage STEM engagement in high schools (see CASE STUDY p. 52).

Michael Milford AMP Tomorrow Maker (l) receives his award from Adam Spencer (Image credit: Michael Milford)

CASE STUDY: SAVING THE WORLD WITH MATHS

Centre Chief Investigator Michael Milford is using Hollywood-style blockbuster entertainment to excite and engage students in mathematical and scientific learning.

Michael created a program called Math Thrills, which uses young adult fiction, movies, and games as well as hands-on workshops to engage students in grades seven to nine.

"What makes these workshops different is that the students are living out the action-packed scenarios they see in their favourite movies, books and games," Michael said.

One example of a workshop game is students using mathematics to crack codes on a nuclear generator to keep it from melting down and destroying a city. The students also solve action scenarios using things like nerf guns and hot wheels cars.

"It's really great to see kids and teachers so excited about maths and the feedback we get from these workshops is incredible," Michael said.

The program also employs an action novel co-written by Michael and Jemma Pollari, called 'Code Bravo,' which is stealthily filled with mathematical concepts.

Michael created Math Thrills created about five years ago, and started trialling it in schools in 2015. He received $2,500 in pre-seed funding on Kickstarter, and $50,000 in seed funding from QUT Bluebox.

The program was given another big boost in 2016, when Michael received a $25,000 AMP Tomorrow Fund award. Michael will use that money on the next phase of the MathThrills program which will draw on the "invisible maths" in movies.

"Receiving a grant from the AMP Tomorrow Fund will act as a fantastic force multiplier – taking what has already been proven to be successful on a small scale in schools with students and teachers and scaling it up," Michael said.

Michael has written and produced innovative textbooks for high school students for fifteen years. He started Math Thrills, though, after he realised his efforts in teaching maths were still not having an impact on those who most needed it – disengaged students.

"My long-term dream is to reach students around Australia, and even the world, who don't have access to education or don't feel like they're interested in mathematics," Michael said.

The initiative has led to awards and honours including the 2015 Queensland Young Tall Poppy of the Year Award and a 2015 TedXQUT talk.

Find out more: MathThrills.com

CASE STUDY: ROBOTIC VISION SUMMER SCHOOL (RVSS)

How do you attract the next generation of roboticists? It's a challenge we decided to tackle by starting a series of Robotic Vision Summer Schools.

We held our second RVSS in April 2016 at Kioloa (New South Wales), at the Australian National University's Coastal Campus. The Centre invited late stage undergraduate and early stage PhD researchers to attend.

We invited high profile international guest speakers to our summer school, including our PI Paul Newman from the Mobile Robotics Group, Oxford University, PI Frank Dellaert from Georgia Tech (now at Facebook), Stefan Williams from the Australian Centre for Field Robotics, University of Sydney, and Jana Košecká from George Mason University.

There were a series of practical sessions for participants to choose from, including a unique opportunity to experiment with computer vision algorithms on actual robotic hardware, the turtle bot. A number of short tutorials were presented on ROS, OpenCV and deep learning, which the students required to process images and control a robot for a task of their own choosing. A demonstration of the robots at the end of the week illustrated that this was a very successful course with a lot of positive feedback.

Participants at one of the Centre's Robotic Vision Summer Schools held at ANU's coastal Kioloa campus.

Networks (International, National, Regional)

We encourage our chief investigators, research fellows, and PhD researchers to travel to our overseas partner universities, and to host visits by researchers from our overseas partners.

In 2016 Centre Partner Investigators Paul Newman (Oxford University) and Frank Dellaert (Georgia Tech, US) visited our nodes in Australia. Other international visitors to the Centre included Professor Fredrik Kahl (Chalmers University of Technology and Lund University), Associate Professor Jana Košecká (George Mason University), Professor Inaba Masayuki (University of Tokyo), Assistant Professor Jacinto Carlos Nasainento (Instituto Superior Tecnico, Lisbon) and Professor Tarek Hamel (University of Nice Sophia Antipolis).

In total we welcomed 22 international visitors to the Centre while our researchers visited 41 international universities and companies.

We support our researchers to visit conferences to tell the world about the great research we are doing, to learn what others are doing, to maintain and extend our networks and to recruit new researchers and students. In 2017 Centre researchers visited 30 international labs and 8 national labs as well as presenting papers at 10 - RSS, ECCV, IROS, CVPR, IJCAI, ICRA, ICCP, NIPS, ACCV, DICTA conferences. Our international reputation is enhanced through publications and the research profile generated through the Centre's critical mass in this important field, which in turn improves our ability to attract top researchers to Australia. We will continue to enhance our international linkages through the effective and ongoing research collaborations that exist between our Australian and international partners. Meaningful visits between the laboratories for our investigators, research fellows and

PhD researchers are encouraged, creating a transnational research community around robotic vision.

To create relationships and networks, as well as promoting Australian research, the Centre has organized a number of successful and high profile workshops at top international conferences in 2016 including:

- **Robotics Science and Systems (RSS) Workshop** Centre RFs Niko Sünderhauf and Jurgen (Juxi) Leitner ran a highly successful workshop at the Robotics Science and Systems (RSS) Conference held at the University of Michigan, USA. The workshop titled "Are the Skeptics right? Limits and Potentials of Deep Learning in Robotics" attracted a crowd of nearly 200 and received a great response on social media.
- **ECCV Workshop** RF Basura Fernando co-organised a very successful workshop at the European Conference of Computer Vision (ECCV), held in the Netherlands, called "Brave new ideas for motion representations in videos". The event was a success with more than 140 people attending the workshop and leaders in motion representations, including Professors Max Welling and Ivan Laptev, giving talks.

We will continue to promote robotic vision in major international conferences and conventions in 2017 and beyond. Australia has won the right to host ICRA, IEEE's International Conference on Robotics and Automation, in Brisbane in 2018.

The bid was led by Alex Zelinsky, the Chair of the Centre's Advisory Committee and Centre Director Peter Corke. It will be the first time ICRA has been hosted in Australia, in fact the first time it has moved to the southern hemisphere. Established in 1984 and held annually, ICRA is the IEEE Robotics and Automation Society's flagship conference and brings a premier international forum for robotics researchers to present their work. The conference joins experts in the fields of robotics and automation together for technical communications through presentations and discussions. In addition to contributed paper sessions, ICRA conferences also include plenary sessions, workshops and tutorials, forums, exhibits, and robot challenges as well as technical tours. Typically ICRA attracts over 2000 delegates from around the world.

From 2017, our summer school will transition from an internal PhD bootcamp to an international summer school in order to establish our leading position in the new field, and also help us build local capacity.

The Centre encourages, and has funding to support, visits between the laboratories of our local and international investigators, research fellows and PhD researchers, creating a transnational research community around robotic vision. We have hosted more than 20 international researchers and visited more than 40 international labs and facilities. (Full details are given in our KPI tables).

Centre researchers at the International Conference on Intelligent Robots and Systems. From left to right, Trung Pham, Thanuja Dharmasiri, Cesar Cadena, Peter Corke and Markus Eich.

Oxbotica envisions a world where cars think for themselves
www.oxbotica.com

CASE STUDY

One of the Centre's original Chief Investigators has taken on a new role as Director of Projects with Oxford start-up, Oxbotica.

Ben Upcroft left the Centre at the end of 2016 to move to the UK. Ben says his new role with Oxbotica has a bit of everything. That includes scoping projects, interacting with clients, working with the technical teams, code development and management.

"It's great being able to get my hands 'dirty' with code and then be able to step back and take a leadership role in the team all within the span of a day," Ben said.

Oxbotica was co-founded by Centre Partner Investigator Paul Newman. The company is working to develop autonomous car software it says will power a variety of vehicles.

"I love that we all have a common goal. I also love the feel of a start-up where things are always changing and evolving," Ben said. "It's never a boring day."

Ben was one of the Centre's original Chief Investigators when the Centre formed in 2014. He joined QUT in 2011 as an Associate Professor, and was leader of the school's Robotics and Autonomous Systems (RAS) Discipline.

Before leaving the Centre, Ben was in charge of its 'Vision and Action' research program. The program focussed on vision control of robots for manipulating objects.

Ben says his time at QUT and with the Centre really prepared him for his current role at Oxbotica.

"The skills that I have developed within research, leadership, collaboration, and education have been invaluable for my role here at Oxbotica." Ben said.

"I loved my time with the Centre. I think of myself as a very lucky person to have been even invited to be part of the Centre. To have seen how it has grown and evolved as we learnt how to work together with so many amazing researchers, has been an opportunity that only a few receive and I feel privileged that I was a part of it."

While he loves his new role with Oxbotica, Ben says he misses working with everyone at the Centre.

International and National Links & Networks KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Number of international visitors and visiting fellows	Annually	2	8	8	11	8	22
Number of national and international workshops held/organised by the Centre	Annually	1	1	2	2	3	10
Number of visits to overseas laboratories and facilities	Annually	2	11	8	25	8	56
Examples of relevant interdisciplinary research supported by the Centre *Number of interdisciplinary projects within the Centre	Annually	0	0	1	3	2	5

Centre Partner Investigator Paul Newman from Oxford University presenting to Centre researchers at QUT

People

Over the life of the Centre, our funding will be used to support 24 Research Fellows and 80 PhD researchers at our collaborating organisations. In 2016 we accepted the nominations of eight new affiliate members.

Two additional Associate Investigators (AIs), Nick Barnes and Felipe Gonzalez, have been recognised, while three more are becoming chief investigators and two have changed roles but remain connected to the Centre as research affiliates (Laurent Kneip, Ahmet Sekercioglu).

Our Centre brings together a critical mass of outstanding researchers with expertise in robotics and computer vision, as well as track records in research leadership and research training track. The original team of investigators comprises a blend of experienced and early career researchers with world-class skills across all the key areas including machine learning (Shen, Gould, Reid, van den Hengel, Carneiro, Drummond, Newman, Torr), mapping and navigation (Reid, Drummond, Wyeth, Milford, Upcroft, Davison, Newman, Dellaert), visual servo control (Corke, Mahony, Chaumette), three-dimensional reconstruction (Hartley, Li, Upcroft, van den Hengel, Torr, Davison, Pollefeys), low-level and high-speed vision (Corke, Davison, Drummond, Pollefeys) and distributed systems (Drummond, Dellaert, Corke). To this group of talented individuals, the Centre has actively recruited five research fellows at QUT, five research fellows at The University of Adelaide, two research fellows at Monash University and four research fellows at ANU. We also have recruited an additional 17 PhD researchers bringing our total PhD cohort numbers to 44. Three new Associate Investigators joined us in 2015 including environmental roboticist Dr Matt Dunbabin (see CASE STUDY p. 21), orthopaedic surgeon and medical robotic enthusiast Prof Ross Crawford (see CASE STUDY p. 39) and mobile computer vision specialist Dr Laurent Kneip (ANU).

The research committee, comprising the executive and the program leaders, was established in May 2016 and has met every month. This has provided better visibility of progress at project level and the group is empowered to resolve project-level challenges.

We are building new capacity by training the next generation of research talent. Already we have grown Australia's research capacity in this area with more than 56 early career researchers being part of our Centre, with a strong international influence, 100% of our Research Fellows are from overseas while 37% of our PhD researchers are international.

As we build our Centre we are also developing a Centre culture, with a set of pervasive values, beliefs and attitudes that characterise who we are. As people join the Centre, we introduce the Centre's Vision and Values, our history, and our plans for the future via an onboarding process. We are a nimble enterprise, with the flexibility and capacity to handle diversity and a range of new ideas. Our Centre has an active intranet (Atlassian's Confluence), an internal Google+ community (robotvision) with 81 members, as well as a public-facing identity on Google+, Facebook, Twitter, LinkedIn, Instagram and Flickr. To promote Centre achievements we produce an online newsletter for both internal and external stakeholders.

The quality of our Centre's personnel is noticed by others and internationally regarded researchers are being lured away. CI Stephen Gould is taking a leave of absence from ANU to work at Amazon, and CI Ben Upcroft is taking a leave of absence from QUT to work at Oxbotica, an Oxford University spinout (see CASE STUDY p. 55). A number of research fellows have left well before the end of their terms such as Donald Dansereau to Stanford and Cesar Cadena to ETHZ.

Retention or replacement of key personnel is a significant threat to the Centre. The competition for PhD and postdoctoral research fellows is fierce. To mitigate this risk, apart from implementing the Centre's succession plan, we encourage recruitment and retention by ensuring that the Centre is a rewarding place to conduct research, and by encouraging our alumni to stay in close contact. Our alums, as well as being a good source of talent, may also have key industry experience and contacts that help further our networks. The advantage of leveraging our alumni network in this way is that our researchers can also benefit from these industry links and potential career opportunities.

Diversity of Staff

Position	#	FTE	Gender Ratio female:male
Chief Investigator	13	3.9	0 : 13
Partner Investigator	6	0.3	0 : 6
Associate Investigator	18	-	0 : 18
Professional Staff*	11	7.68	7 : 4
Research Fellows	17	17	2: 15
PhD Researchers	44	44	5 : 39

*Professional Staff - includes full-time and part-time operations staff and 2 Centre Research Engineers

Centre researchers enjoy dinner at RoboVis2016, Lorne, Victoria

CASE STUDY: ROBOVIS2016

To encourage communication and collaboration between researchers at different universities, we hold an annual symposium called "RoboVis," attended by all of the Centre's Chief Investigators, Associate Investigators, Research Fellows, PhD researchers and friends of the Centre. The term, "RoboVis," symbolises what we're all about, bringing researchers from the fields of robotics and computer vision together.

Our third annual Centre conference, RoboVis, was held in Lorne, Victoria at the end of September. This year we were joined by members of our Centre Advisory Committee and End-User Advisory Board. The symposium included technical talks spanning the range of our research programs, as well as the ever-popular demo session where we saw some Centre technologies in action. Innovations this year included a three-minute thesis competition where PhD researchers gave pitches on how their research would change the world and a Centre awards ceremony. The awards recognised exceptional performance in raising our Centre's profile and collaborative endeavours.

Incorporated in the symposium was the first in a series of customised knowledge leadership training modules to help support the career development of our research fellows and PhD researchers.

RoboVis will continue to be an annual showcase and celebration of the Centre's research with RoboVis2017 scheduled to be hosted by our QUT node at Moreton Island in October 2017.

Research Training

Seventeen new postgraduate researchers started with our Centre in 2016 as well as four new research fellows.

Over the life of our Centre we will train 24 postdoctoral fellows and 80 PhD students. We provide a rich training experience for young researchers. Development opportunities include: travelling to other Centre nodes or our overseas partner universities to experience different research environments; attending top international conferences; regular seminars at all nodes; and interacting with eminent visiting researchers from our local and overseas partners or other universities. Our annual symposium,RoboVis, includes presentations from Centre researchers, invited guests as well as PhD student thesis pitches and formal training sessions. To support the development of our early career researchers, a range of research training opportunities are being provided including a mixture of face-to-face and online opportunities, supplemented by a Research Training Toolbox hosted on our wiki intranet, Confluence.

At our annual symposium, RoboVis, a day was set aside exclusively for research training for PhD researchers and research fellows focussed on knowledge leadership and career development (see CASE STUDY above).

Additional courses through 2016 were delivered by COO Sue Keay, including a start-up course using based on Lean Launchpad method and using Steve Blank's Udacity MOOC with Eric Ries book on The Lean Start-Up used as reference material. Another course delivered as a half-day workshop has been around career planning, underpinned by Peter Fiske's Career Guide for Scientists (and Engineers).

The Centre Executive members have demonstrated track records in research training and mentor their PhD researchers, postdoctoral fellows and younger CIs. In addition we encourage all Centre researchers to have a mentor outside their direct supervisory team, to act as a trusted advisor. Mentoring relationships are flexible and intended to support (not replace) existing supervisory relationships. Participation is voluntary. We are currently looking at how best to support our early career researchers to make the most of such mentoring opportunities and covered this explicitly as part of our Knowledge Leadership training at our annual symposium in September 2016.

Research Training and Professional Development KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Number of professional training courses for staff and postgraduate students attended * Number of courses for early career researchers, staff and PhD students	Annually	0	0	4	7	4	5
Number of Centre attendees at all professional training/development courses offered by the Centre (include courses offered for external stakeholders and clients)	Annually	0	0	30	120	50	83
Number of new postgraduate students working on core Centre research and supervised by Centre staff (include PhD, Masters by research and Masters by coursework)	Annually	10	6	20	23	10	17
Number of new postdoctoral researchers recruited to the Centre working on core Centre research	Annually	7	6	9	11	1	3
Number of new Honours students working on core Centre research and supervised by Centre staff	Annually	4	10	8	7	8	9
Number of postgraduate completions and completion times, by students working on core Centre research and supervised by Centre staff	Annually	0	0	0	0	0	5
Number of Early Career Researchers (within five years of completing PhD) working on core Centre research	Annually	7	6	16	14	16	16
Number of students mentored	Annually	45	86	90	125	90	157
Number of mentoring programs offered by the Centre (include programs for students, new staff, external stakeholders and clients)	Annually	0	1	2	2	2	2

Gender Diversity

Our Chief Investigators are all male, and this reflects gender balance issues which are acute in computer science, and only slightly better in engineering.

The Centre is developing strategies to remediate this as much as possible, for example looking at ways to address unconscious bias in the advertising and recruitment of research fellows and PhD researchers using Textio (an online program that identifies bias in job recruitment advertisements), and will work with existing programs such as the Anita Borg Foundation to use inspirational topics in robotics and computer vision to attract more women into IT and engineering at the undergraduate level. In 2016 Centre COO Sue Keay hosted a visit by QUT's Girls in Engineering Making Statements (GEMS) student club, while Centre CI Stephen Gould hosted a Girls in Engineering and Technology (GET Set) program & other outreach schemes at ANU. In 2017 we will engage with external consultants, Gender Matters, to continue to explore opportunities to encourage diversity.

Sarah Allen, Thuy Mai and Kate Aldridge with Centre Director Peter Corke

Leadership Development

As part of the human capital development and planning the Executive Committee has identified successors for the key roles of Director and Node representatives.

These individuals will be mentored in areas of leadership and take on responsibilities such as participating in Executive Committee meetings when the incumbent is unable to attend.

We are building new capacity by training the next generation of research talent. Already we have grown Australia's research capacity in this area with more than 56 early career researchers being part of our Centre, with a strong international influence, 100% of our Research Fellows are from overseas while 37% of our PhD researchers are international.

The quality of our Centre's personnel is noticed by others and internationally regarded researchers are being lured away. CI Stephen Gould is taking a leave of absence from ANU to work at Amazon, and CI Ben Upcroft is taking a leave of absence from QUT to work at Oxbotica, an Oxford University spinout. A number of research fellows have left well before the end of their terms: Donald Dansereau to Stanford, Cesar Cadena to ETHZ and others from Adelaide.

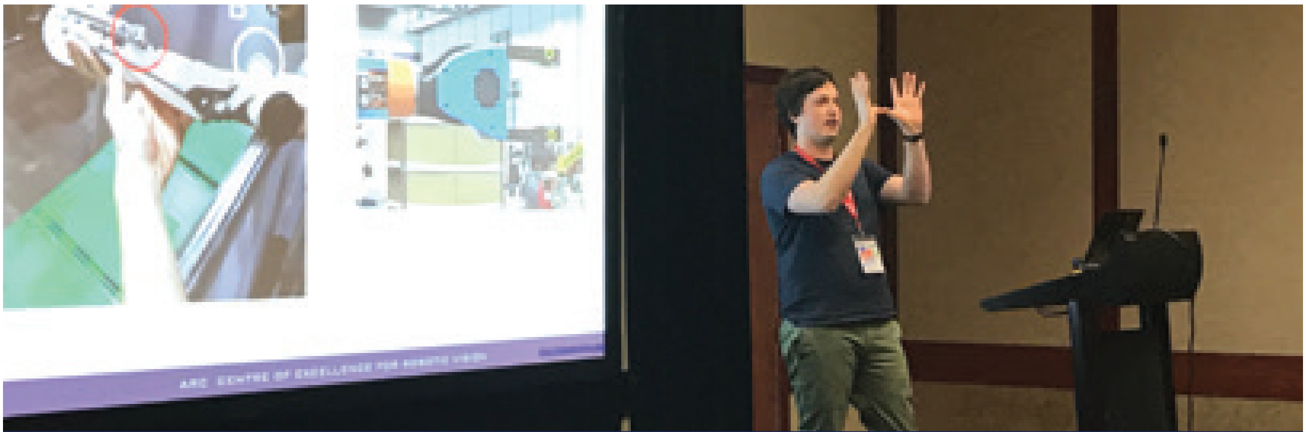

Adam Tow, a Centre PhD Researcher, presents his research at the Centre's RoboVis annual symposium.

Early Career Researchers at the Centre's annual symposium RoboVis, also enjoyed the first session of our Knowledge Leadership program run by Evexia.

CASE STUDY: WORLD FIRST KNOWLEDGE LEADERSHIP PROGRAM FOR EARLY CAREER RESEARCHERS

We have an important responsibility to nurture the next generation of knowledge leaders, innovating robotic vision experts who will go on to work in industry, government and academia. Our research fellows and PhD researchers will help to translate robotic vision research to adapt existing industries and to convert robotic vision research into new products, services and enterprises so as to forge whole new industries.

To support our early career researchers we have commissioned a tailor-made knowledge leadership program from workplace psychologists, Evexia. The program, which commenced last year, focusses on the skills our researchers need for successful careers. Unlike the plethora of other leadership training programs which focus on either the research or corporate worlds, our bespoke program covers some of the specialist skills required by researchers but also encompasses entrepreneurship, publicity, engagement and resilience. We hope to equip our researchers with skills that will help them to not only succeed, but to also lead in any sphere of influence.

The program started last year with the introduction of career development planning and has been supported by the Centre COO, Sue Keay, with workshops at each node on careers. The future focus of the program is on helping researchers with developing purpose and a unique identity, learning to influence, translating research and bringing entrepreneurial resolve to bear, and also on further career planning and resilience.

We have set the bar high and believe we have developed a unique program combining elements of traditional leadership training with a specific focus on research and entrepreneurship.

CASE STUDY: KICKSTART FOR SCHOOL ROBOTS

The dream for one PhD student from the Australian Centre for Robotic Vision to provide robotics for schools around the country becomes a reality.

James Mount created BBot robots to send to schools and individuals around Australia. BBots are easy-to-use, inexpensive robotic vehicles with several modes to help introduce kids to robotics and the scientific process.

The BBots were created as part of a Kickstarter campaign, where James had to reach a goal of $500 to get the project up and going. He reached that goal in July 2016, a week before the Kickstarter deadline. He was then able to begin the design and assembly process. When he was finished, James had created 17 BBots to send out.

James' mother is a school teacher, so he wanted to figure out how he could give back and help schools.

"This allowed me to combine my family's background in education with my expertise in robotics," said James. "I think robotics is a great teaching tool that can engage and excite students."

"It's been a bit of a roller coaster. It's been challenging at times, but in the end it's been worthwhile. I hope that the people who supported the project are happy with the BBot, whether they want to use it to teach or just learn from."

The goal of the BBot project is to engage kids and help them understand the scientific process. For example one of the worksheets uses BBot's differential drive system to help kids explore the scientific concepts of hypothesising and experimental design. The workshop challenges students to think scientifically by asking them what they think BBot will do if both wheels are running at the same speed, if one wheel is running faster than the other, or if one wheel isn't running at all.

James is working toward his PhD at the Queensland University of Technology (QUT). His PhD supervisor is Associate Professor Michael Milford, a Chief Investigator at the Australian Centre for Robotic Vision.

Centre PhD researcher James Mount with two of his BBot robots

Michael saw this project as a way to help schools deal with a world where technology is constantly evolving.

"I hope this is the start of an inclusive movement to give everyone the basic scientific and technological awareness they need to prosper in this rapidly changing world," said Milford.

The project started with a very basic design, and then evolved to its finished product. James said he couldn't be happier with how it turned out.

Mentoring

Our goal is for all Centre researchers to have a mentor who will act as a trusted advisor. Mentors are selected by the mentee and mentoring relationships are intentionally flexible with a view to support (not replace) existing supervisory relationships. Participation is voluntary.

Benefits of having a mentor include access to an informed second opinion, gaining insight into one's own performance through a 'critical friend', identifying personal development needs and opportunities as well as learning from the experience of the mentor.

The role of a mentor may include:

- Sharing expertise and experience to help mentees develop their talents;
- Listening, clarifying, reflecting back and, when called for, challenging mentees to view issues from a variety of perspectives;
- Opening doors, helping the mentee to network and develop their careers;
- Providing a safe sounding board for mentees to raise and talk about issues.

The role of a mentee may include:

- Taking responsibility for identifying and achieving development and career goals;
- Initiating meetings with their mentor, managing meeting dates and times and setting the agenda;
- Being open to and appreciating different perspectives as well as constructive and honest feedback;
- Being considerate of the demands placed on their mentor's time.

Early career researchers select their own mentors in a structured way supported by the Centre's Administrative Coordinator and we are instituting regular mechanisms to follow-up and ensure that mentoring relationships are satisfactory to all concerned.

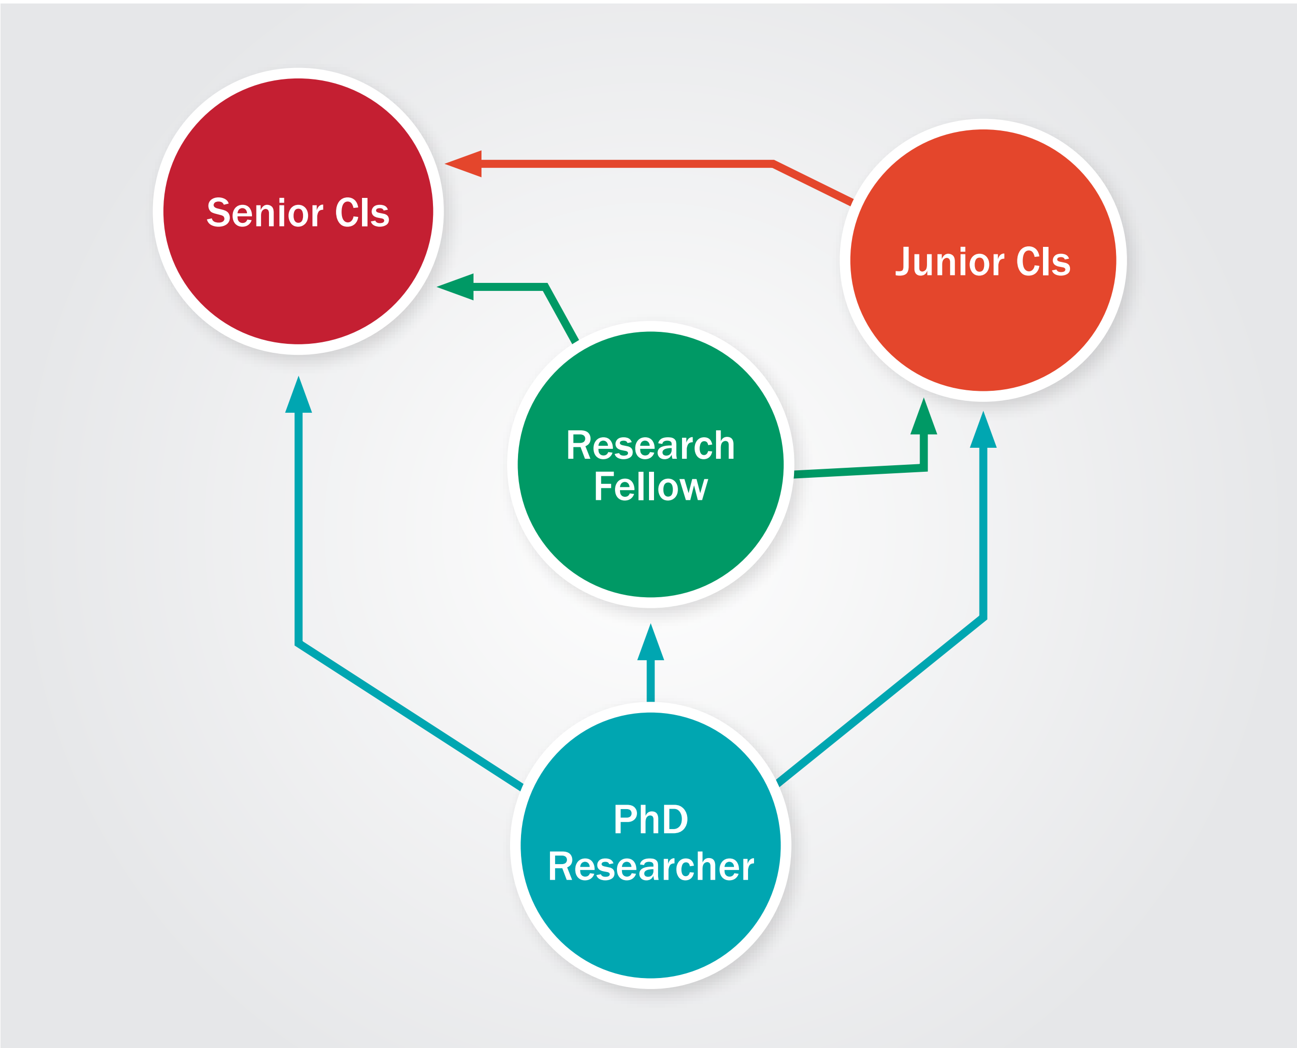

Chief Investigators

*2017 roles are indicated in parentheses

Peter Corke

SENSING, ACTING, TECHNOLOGY

Rob Mahony

SENSING, ACTING

Chunhua Shen

SENSING, UNDERSTANDING, ACTING, LEARNING

Ian Reid

UNDERSTANDING, LEARNING, TECHNOLOGY

Hongdong Li

SENSING, TECHNOLOGY

Gustavo Carneiro

LEARNING

Tom Drummond

UNDERSTANDING, TECHNOLOGY

Michael Milford

SENSING, UNDERSTANDING

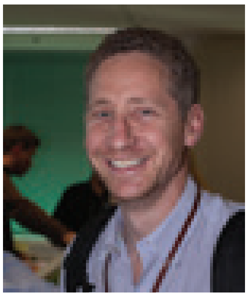

Stephen Gould (Research Affiliate)

UNDERSTANDING

Richard Hartley

TECHNOLOGY

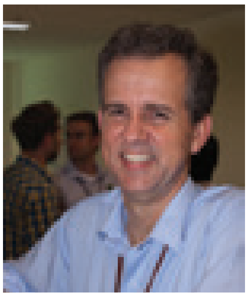

Gordon Wyeth

UNDERSTANDING

Ben Upcroft (Research Affiliate)

ACTING

Anton van den Hengel

UNDERSTANDING

Associate Investigators

*2017 roles are indicated in parentheses

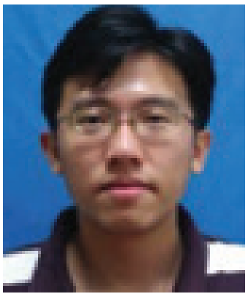
Tat-Jun Chin (Chief Investigator)


Matt Dunbabin (Chief Investigator)

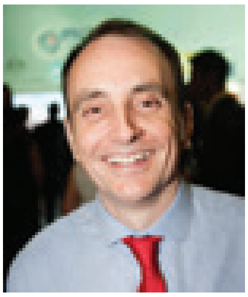
Jonathan Roberts (Chief Investigator)


Nick Barnes


Ross Crawford


Anthony Dick

Anders Eriksson


Clinton Fookes


Jason Ford


Jonghyuk Kim


Laurent Kneip (Research Affiliate)


Luis Mejias Alvarez


Tristan Perez

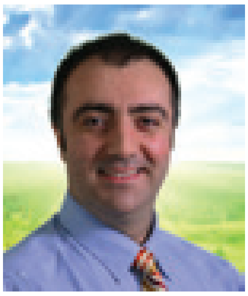
Fatih Porikli


Ahmet Sekercioglu (Research Affiliate)


Qinfeng (Javen) Shi


David Suter


Jochen Trumpf

Partner Investigators


François Chaumette


Paul Newman


Andrew Davison


Frank Dellaert


Phillip Torr


Marc Pollefeys

Research Fellows

Khurrum Aftab

*Left the Centre

Feras Dayoub

ACTING

Basura Fernando

UNDERSTANDING, LEARNING

Anoop Cherian

UNDERSTANDING

Thanh-Toan Do

UNDERSTANDING, LEARNING

Viorela Ila

ACTING, TECHNOLOGY

Donald Dansereau (Research Affiliate)

SENSING

Markus Eich

UNDERSTANDING

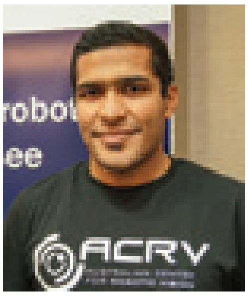

Vijay Kumar

LEARNING

Yasir Latif

UNDERSTANDING, TECHNOLOGY

Vincent Lui

TECHNOLOGY

Trung Than Pham

UNDERSTANDING, TECHNOLOGY

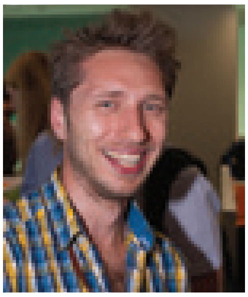

Jürgen "Juxi" Leitner

SENSING, ACTING

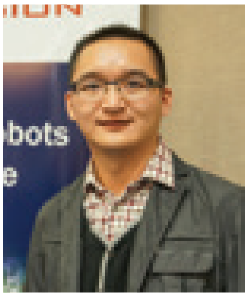

Chao Ma

UNDERSTANDING

Sareh Shirazi

UNDERSTANDING

Research Fellows (cont.)

*Associated

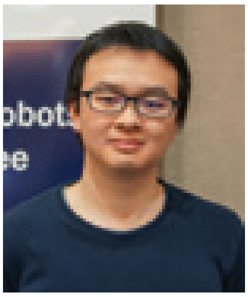

Guosheng Lin

LEARNING

Chuong Nguyen

SENSING

Niko Sünderhauf

SENSING,
UNDERSTANDING,
TECHNOLOGY

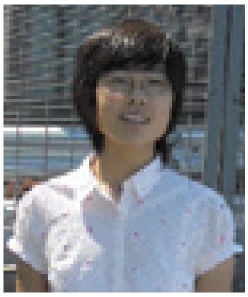

Lin Wu

UNDERSTANDING

*Left the Centre

Alireza Khosravian*

Pan Ji*

Ravi Garg*

PhD researchers

Juan Adarve
ACTING 2015-

Ming Cai
UNDERSTANDING 2016-

Jeffrey Devaraj
ACTING 2016-

Rafael Felix Alves
LEARNING 2016-

William Chamberlain
TECHNOLOGY 2015-

Thanuja Dharmasiri
TECHNOLOGY 2015-

Peter Anderson
UNDERSTANDING 2015-

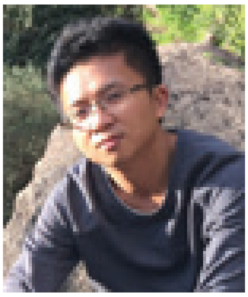

Zetao "Jason" Chen
SENSING 2014-2016

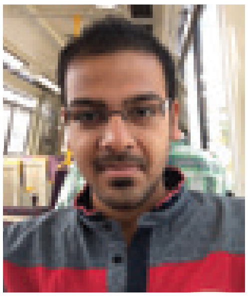

Sourav Garg
SENSING 2015-

Zongyuan Ge
LEARNING 2014-2017

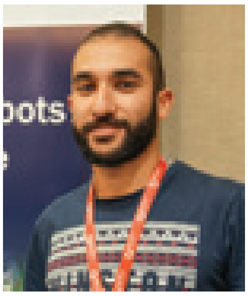

Mina Henein
TECHNOLOGY 2016-

Chris Jeffery
MEDICAL 2015-

Jian "Edison" Guo
LEARNING 2015-

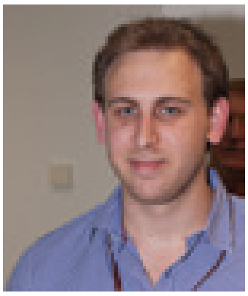

Adam Jacobson
SENSING 2014-

Yuchao Jiang
LEARNING 2015-

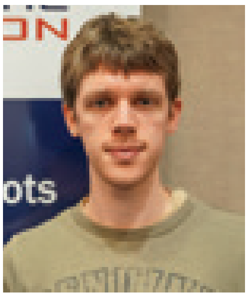

Ben Harwood
LEARNING 2015-

Jasmin James
SENSING 2016-

Peter Kujala
ACTING 2016-

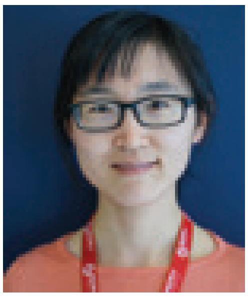

Hui Li
LEARNING 2015-

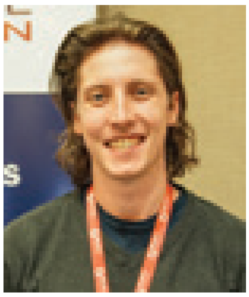

Sean McMahon
SENSING, TECHNOLOGY 2015-

PhD researchers (cont.)

*Associated

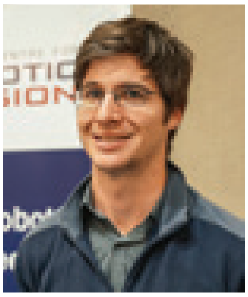

James Mount
SENSING 2016-

Zhibin Liao
LEARNING 2015-

Mehdani Menikdiwela
SENSING 2016-

Sean O’Brien
ACTING 2016-

Andres Felipe Marmol Velez
MEDICAL 2015-

Ben Meyer
LEARNING 2015-

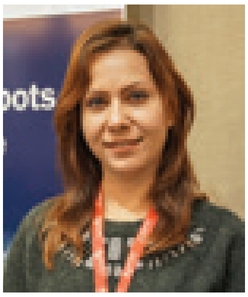

Fahimeh Rezazadegan
UNDERSTANDING 2014-

Dan Richards
SENSING 2016-

James Sergeant
SENSING 2016-

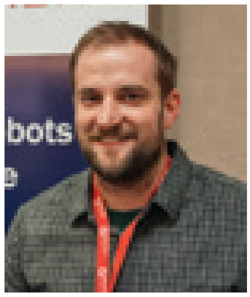

Andrew Spek
TECHNOLOGY 2015-

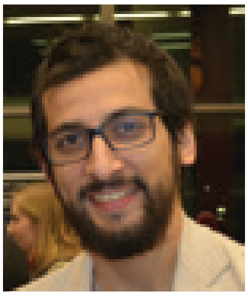

Cristian Rodreguez
TECHNOLOGY 2016-

Tong Shen
LEARNING 2016-

Jean-Luc Stevens
ACTING 2016-

Rodrigo Santa Cruz
UNDERSTANDING 2015-

John Skinner
TECHNOLOGY 2016-

Adam Tow
ACTING 2015

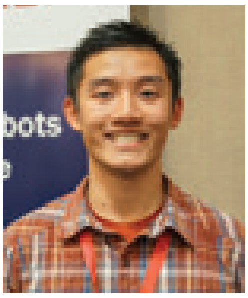

Dorian Tsai
ACTING, SENSING 2016-

Jun Zhang
TECHNOLOGY 2016-

Bohan Zhuang
LEARNING 2015-

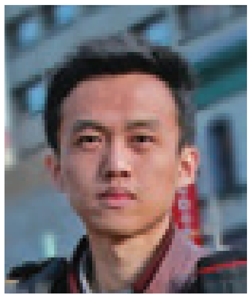

Xiaoqin Wang
TECHNOLOGY 2015-2016

Fangyi Zhang

ACTING 2015-

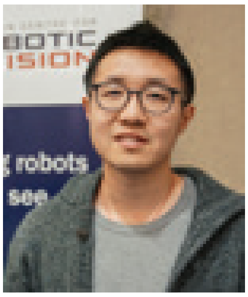

Yan Zuo

LEARNING 2015-

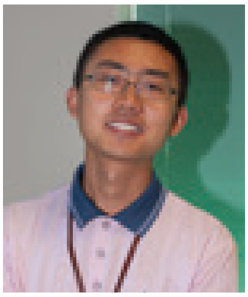

Yi “Joey” Zhou

TECHNOLOGY 2014-

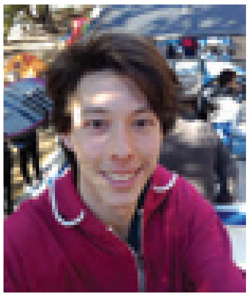

Cedric Scheerlinck*

Josh Boys*

Mehdi Hosseinzadeh*

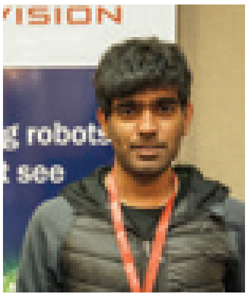

Saroj Weerasekera*

Josh Weberruss

TECHNOLOGY 2015-

Research Affiliates

*2017 roles are indicated in parentheses

Cesear Cadena

Felipe Gonzalez (Associate Investigator)

Mark McDonnell

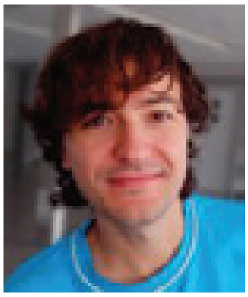

Anton Milan

Thierry Peynot

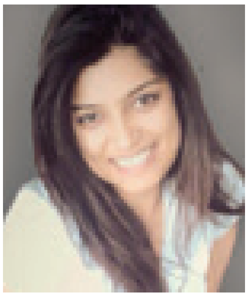

Anjali Jaiprakash

Ajay Pandey

Leo Wu

Chris McCool

Chris Lehnert

Research Engineers

Alex Martin

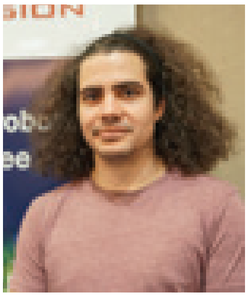

Steven Martin

Professional Staff


Kate Aldridge


Tracy Kelly


Thuy Mai


Sarah Allen


Greg Lee


Sandra Pedersen


Sue Keay

Tim Macuga

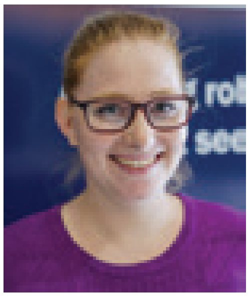
Rachel Sinnott

Governance

Our Centre is an unincorporated collaborative venture given carriage of $25.6m funding over seven years to meld the disciplines of robotics and computer vision together and pursue an ambitious research agenda in the new field of robotic vision.

Good governance is the responsibility of our Centre Executive Committee, with oversight provided by the Centre's Advisory Committee. The Centre Executive are accountable to the Australian Research Council, a statutory agency responsible for Australia's National Competitive Grants Program, which contributes $19m in public funding to our Centre. The Executive also represents our four domestic research partners; Queensland University of Technology (QUT), The University of Adelaide (Adelaide), The Australian National University (ANU), and Monash University.

We are fortunate to be able to call on the extensive research management experience of the leadership team formed by the Centre Director (Professor Peter Corke), Deputy Director, ARC Laureate Professor Ian Reid (representing University of Adelaide), Professor Robert Mahony (representing ANU), Professor Tom Drummond (representing Monash), and Chief Operating Officer Dr Sue Keay. Together they comprise the Centre's Executive Committee, providing the leadership and direction critical to the management of our ambitious research program. The Executive currently meets fortnightly (or more regularly) via videoconference with F2F meetings rotated between nodes every quarter. Once a month the Executive also meet with all Research Program leaders forming the Centre's Research Committee, to review research progress, to make decisions on the direction of the Centre's research, and to annually review and modify the strategic research direction to ensure the Centre can meet its many objectives.

Our Centre Advisory Committee, oversees the strategic direction of the Centre and comprises an independent Chair (Dr Alex Zelinsky, CEO of DST Group) and up to six committee members representing academia and industry with expertise in science and technology transfer. Our Advisory Committee meets at least annually, and is also invited to the Centre's annual symposium (RoboVis) to meet with our researchers and staff. Meetings of our Advisory Committee were held in January 2016, June 2016, and as part of RoboVis in September 2016. The Centre has also constituted its End-User Advisory Board (EUAB), chaired by Dr Russel Rankin (CEO of Food Innovation Partners) and including up to six representatives of industry. Meetings have been held in June 2016, and as part of RoboVis in September 2016.

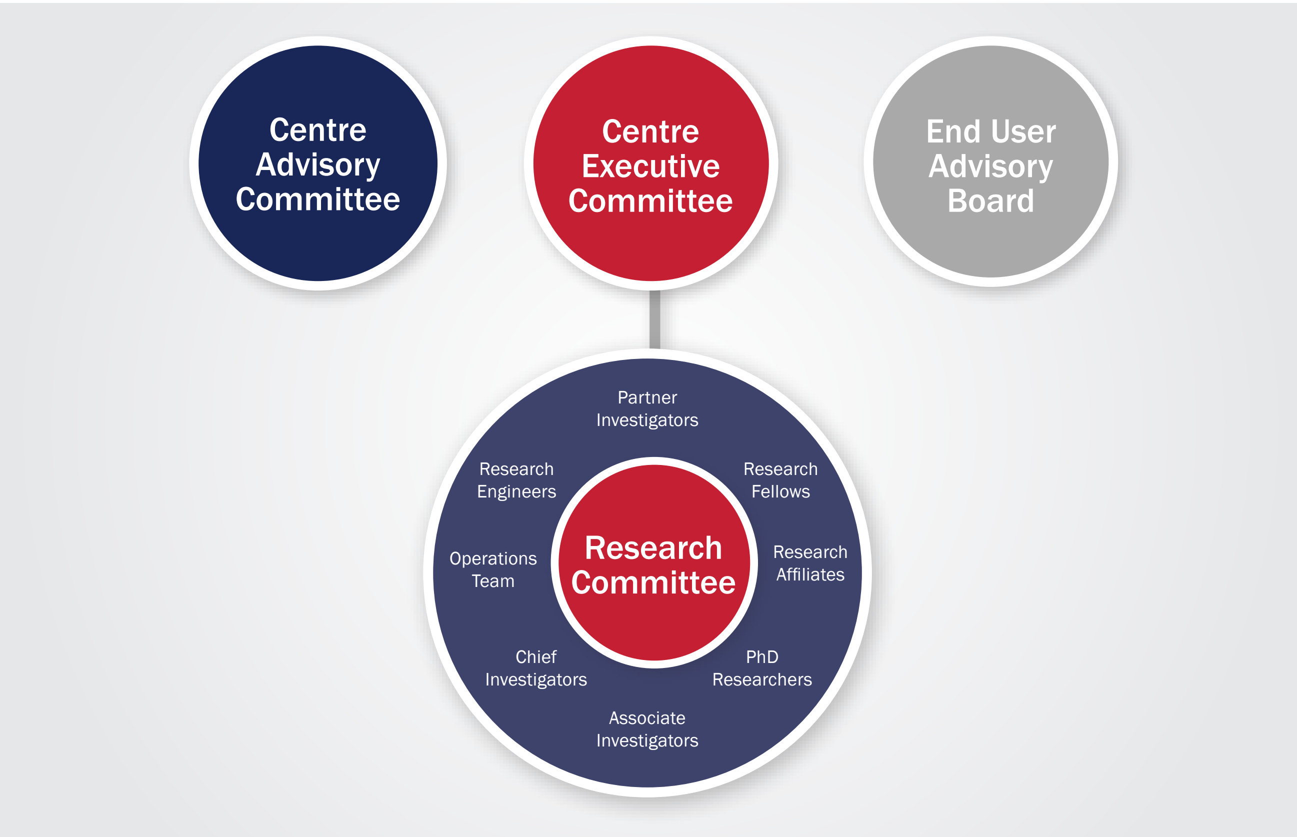

A team including Centre researchers gained 2nd place in the International RobotX Maritime challenge held in Hawaii.

Our Committees

Centre Advisory Committee (CAC)

Our Centre Advisory Committee comprises six people with expertise in the relevant science and have a track record in technology commercialisation. Our CAC meets annually with the Executive Committee, and is also invited to the Centre's annual symposium to afford the chance for the Advisory Committee members to meet with our researchers and staff. Our CAC provides frank advice on the research direction of the Centre, evaluates our progress toward goals and KPI achievement, and provides introductions and opportunities from their own extensive networks.

Centre Executive Committee (CEC)

The CEC governs the Centre. The committee comprises our Centre Director, three CIs representing Monash, University of Adelaide and ANU, and the Chief Operating Officer. The Committee meets fortnightly via Citrix Go-To-Meeting: and quarterly F2F, to ensure the effective operation of the Centre towards its goals; to develop the annual operational plan; to track performance against the agreed measures; to share project status, opportunities and upcoming events; to resolve problems; to identify opportunities for collaboration between themes and locations; and to identify protectable intellectual property. Agendas, actions, resolutions, and notes are recorded using the Centre's Confluence-based intranet.

End-User Advisory Board (EUAB)

Our End-User Advisory Board includes up to six representatives of industry including the Chair, particularly in the Centre's main areas of application: agriculture; medical & healthcare; infrastructure & asset management; building & construction; and smart manufacturing. Members of the EUAB are appointed for a period of two years. The role of the EUAB is to provide advice/ experience regarding industry needs and how the work of the Centre can align with those needs; use their understanding of the Centre to act as advocates for the Centre; assist in reviewing progress towards achieving strategic end-user objectives and suggest improvements; and provide a short report to the Centre Director and the Executive Committee on an annual basis with comment and advice on the Centre's End-User engagement.

Meetings involve our board members engaging with our executive, participating in one of our industry breakfasts, and assessing how our industry engagement strategy is working, and offering suggestions for impact and engagement.

Research Committee

All research programs report to the Centre Research Committee. The committee comprises all five research program leaders and the Centre Executive. Their role is to review research progress, to make decisions on the direction of the Centre's research, and to annually review and modify the strategic research direction to ensure the Centre can meet its many objectives.

Members of the Centre Advisory Committee

Dr Alex Zelinsky (Chair)

Alex is the Chief Defence Scientist and head of the Defence Science and Technology Group. Before joining Defence he 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. Alex 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 Alex was Professor of Systems Engineering. Previously Alex researched in 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. Alex has extensive experience in advising Federal and State governments in Australia, including as a member of the Australian Government's Defence Industry Innovation Board. Alex 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.

Scientia Professor Michelle Simmons

Michelle is an Australian Research Council Laureate Fellow & Director of the highly successful Centre of Excellence for Quantum Computation and Communication Technology. She has pioneered unique technologies internationally to build electronic devices in silicon at the atomic scale, including the world's smallest transistor, the narrowest conducting wires and the first transistor where a single atom controls its operation. This work opens up the prospect of developing a silicon-based quantum computer: a powerful new form of computing with the potential to transform information processing. Professor Simmons is one of a handful of researchers in Australia to have twice received a Federation Fellowship and now a Laureate Fellowship, the Australian Research Council's most prestigious award of this kind. She has won both the Pawsey Medal (2006) and Lyle Medal (2015) from the Australian Academy of Science for outstanding research in physics and was, upon her appointment, one of the youngest fellows of this Academy. She was named Scientist of the Year by the New South Wales Government in 2012 and in 2014 became one of only a few Australians inducted into the American Academy of Arts and Sciences. In 2015 she was awarded the CSIRO Eureka Prize for Leadership in Science and in 2016 the Foresight Institute Feynman Prize in Nanotechnology for her work in 'the new field of atomic-electronics, which she created'. She is Editor-in-Chief of Nature Quantum Information and was recently named the 2017 L'ORÉAL-UNESCO Asia-Pacific Laureate in the Physical Sciences.

Professor Mandyam Srinivasan

Srini's research focuses on the principles of visual processing, perception and cognition in simple natural systems, and on the application of these principles to machine vision and robotics. He is presently Professor of Visual Neuroscience at the Queensland Brain Institute and the School of Information Technology and Electrical Engineering of the University of Queensland.

Srini holds an undergraduate degree in Electrical Engineering from Bangalore University, a Master's degree in Electronics from the Indian Institute of Science, a Ph.D. in Engineering and Applied Science from Yale University, a D.Sc. in Neuroethology from the Australian National University, and an Honorary Doctorate from the University of Zurich.

Among his awards and honours are Fellowships of the Australian Academy of Science, of the Royal Society of London, of the Royal Institute of Navigation, and of the Academy of Sciences for the Developing World, an Inaugural Federation Fellowship, the 2006 Australia Prime Minister's Science Prize, the 2008 U.K. Rank Prize for Optoelectronics, and the 2009 Distinguished Alumni Award of the Indian Institute of Science, the Membership of the Order of Australia (AM) in 2012, the Queensland Science Championship in 2014, and the Harold Spencer-Jones Gold Medal of the Royal Institute of Navigation in 2014.

Professor Hugh Durrant-Whyte

Hugh is a Professor and ARC Federation Fellow at The University of Sydney. His research work focuses on robotics and distributed sensor networks and he has published over 350 papers. His work with industry includes major robotics and automation projects in cargo handling, surface and underground mining, defence, unmanned flight vehicles and autonomous sub-sea vehicles. He has won numerous awards and prizes for his work including the ATSE Clunies Ross Award, IFR/IEEE Invention and Entrepreneurship Award, the NSW Pearcey Award, and four IEEE Best Paper prizes. He was named Professional Engineer of the year (2008) by the Institute of Engineers Australia Sydney Division,and NSW Scientist of the Year (2010). He was an IEEE Robotics and Automation Society Distinguished Lecturer (2006-10). He is a Fellow of the Academy of Technological Sciences and Engineering (FTSE), a Fellow of the Institute of Electrical and Electronic Engineers (FIEEE), a Fellow of the Australian Academy of Science (FAA), a Fellow of the Royal Society (FRS). He served as the Chief Executive Officer of National ICT Australia Limited (NICTA) from December 2010 to November 2014.

Professor Michael Brady

Mike is the Professor of Oncological Imaging in the Department of Oncology at the University of Oxford, having retired from his Professorship in Information Engineering after 25 years (1985-2010). Prior to joining Oxford, he was Senior Research Scientist in the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology (MIT), where he was one of the founders of the Robotics Laboratory. Mike has been elected a Fellow of the Royal Society, Fellow of the Royal Academy of Engineering, Fellow of the Academy of Medical Sciences, Honorary Fellow of the Institution of Engineering and Technology, Fellow of the Institute of Physics, and Fellow of the British Computer Society. He was awarded the Institution of Engineering and Technology (IET) Faraday Medal for 2000, the IEEE Third Millennium Medal for the UK, the Henry Dale Prize (for "outstanding work on a biological topic by means of an original multidisciplinary approach") by the Royal Institution in 2005, and the Whittle Medal by the Royal Academy of Engineering 2010. He was knighted in the New Year's honours list for 2003.

Members of the Centre's End-User Advisory Board

Russel Rankin (Chair)

With more than thirty years' experience in the food and beverage industry in various senior commercial and research positions. Russel has an inherent ability to identify and drive innovation opportunities that deliver a competitive advantage. He has extensive experience in commercialisation of ideas from within commercial businesses and research organisations. Currently Russel is Director and Founder of Food Innovation Partners Pty Ltd, a company that makes the connections between commercial companies, research organisations, Governments, Finance providers, marketing and Industry Bodies. Food Innovation Partners provides business, innovation and commercialisation services to the food industry, along with business development services for companies and research organisations. Through his ability to translate research outcomes into a commercial competitive advantage and back translate commercial opportunities into a research strategy, he is able to match food companies with research providers; develop pre-competitive, syndicate projects with multiple commercial partners to; assist companies to access Government support programs; assist companies to assess equity and make acquisitions; and helping food businesses innovate and commercialise new ideas.

Russel is entrepreneurial in his thinking, having recently established with his partners a number of new businesses to take a range of new innovative food and beverage products to market.

Prior to starting Food Innovation Partners, Russel 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.

Russel is currently Chair of Queensland Department of Agriculture & Queensland University of Technology's Agricultural Robotics program Advisory Board; Member of SA Government's Advisory and Assessment Board for their 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 (TFMC); Director of Freshly & Co and Director of Beauty Drink Pty Ltd. Previously Russel has been a Board member of the Clean Technology Food & Foundries Investment Program, an initiative of AusIndustry in the Department of Industry, Innovation, Climate Change, Science, Research and Tertiary Education. He has also been an Advisory Board Member for Oski Drinks Pty Ltd.

Alan Davie

Alan has worked across a range of industries throughout the world including Australia US, UK, SE Asia, and the Pacific. He has over forty years experience in the planning, assessment, engineering and construction management of major resources, urban development and infrastructure redevelopment projects, tourism feasibility and planning studies. He has qualifications in Project Management, Town Planning, Environmental Engineering and Civil Engineering and is experienced in People Management, Company Management and Stakeholder Communications.

Alan has held long term Board positions on company and advisory boards including Sinclair Knight Merz P/L, Project Dynamics P/L, ANU Enterprise P/L, Australian Scientific Instruments P/L, Very Small Particle Company P/L, Griffith University School of Environmental Engineering Advisory Board. He is now Managing Director of his own consultancy which undertakes assessment and advice on strategic infrastructure projects.

While employed for most of his career in a global engineering company, Alan held management positions including Queensland and Pacific Business Development Manager, Queensland Operations Manager, General Manager Water and Environment, General Manager International Development Assistance.

Andrew Harris

Andrew 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, with local turnover of ~$3 bn p.a.

Andrew 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.

Andrew 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.

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

Trent Lund

Trent is the Lead Partner for Innovation & Digital Ventures at PwC Australia. He helps organisations leverage emerging technologies to transform ideas into customer-centered, commercial outcomes. With two decades of industry knowledge, Trent has worked across the globe - including in the Asia-Pacific, 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.

Peter Katsos

Peter is employed by ABB, a global leader in power and automation technologies and currently is the General Manager for ABB Robotics Australia. He has a technical background and holds a Bachelor Degree in Electrical Engineering & Computing.

Peter has worked in industrial automation and robotics for over 25 years and has gained a wide range of experience ranging from service, design installation and project management of turnkey systems through to sales, business development & management.

Rob Wood

Rob is the Director, Research and Development at Stryker Australia, one of the world's leading medical technology companies. Stryker offers a diverse range of innovative medical technologies, including reconstructive, medical and surgical, and neurotechnology and spine products. He has a technical background holding a Master of Science in Mechanical Engineering from Stanford University, which focused on Mechanical Design and Orthopedic Biomechanics.

Governance KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Breadth, balance and experience of the members of the Advisory committee	At review						
Frequency, attendance and value added by Advisory Committee meetings *a quorum is defined as 75% of Advisory Board members attending.	At review	1 meeting of Advisory Board	0	1	1	1	2
Vision and usefulness of the Centre strategic plan	At review	Annual Review by Board			at review		at review
The adequacy of the Centre's performance measure targets	At review	As above			at review		at review
Effectiveness of the Centre in bringing researchers together to form an interactive and effective research team	Annually						
• Weeks spent by Centre researchers in other nodes		10	1	45	51	45	31
• Number of new joint research projects		0	0	1	1	2	2
Capacity building of the Centre through scale and outcomes	At review	Early career staff will be developed through supervisor mentoring, a Centre-run mentoring program, and opportunities for supervision (of PhD and undergraduate students), teaching and grant writing. PhD students will be developed through supervisor mentoring, a Centre-run mentoring program, and opportunities for supervision of undergraduate students.					

Finance + Operations

Our Financial Statement provides a summary of our financial performance for the 2016 calendar year.

We receive funding from two main sources, the Australian Research Council, and the Centre's Collaborating Organisations. Details of these contributions are given below.

As well as cash contributions, the Centre's Collaborating and Partner Organisations also provide significant resources, mainly in the form of time commitments from our researchers, as "in-kind" contributions. As the Centre progresses, additional funding will also be sourced via industry engagement and relevant industry projects. Already the halo of industry-funded projects related to the Centre's work is almost $21m, although these projects are not reported as part of these financial statements (see Figure 3).

Our expenditure is dominantly on personnel, with smaller amounts allocated towards travel, equipment and operating expenses (see Figure 4). A significant carry-forward was accumulated in early in the life of the Centre as ARC funds were released more than six months before legal agreements governing the operation of the Centre could be put in place. Expenditure has been below target until we were able to recruit the people required to deliver the Centre's objectives. This carry-forward has continued up until 2016 but has decreased to <$0.75m. This carry-forward is fully allocated in our budget and we project the carry-forward will be fully expended with annual expenditure projected to exceed "new" annual income for 2017 and beyond (where "new" income is the ~$3.7m cash contracted to be received by the Centre each year until 2021).

The Centre Executive ensures the effective operation of the Centre to achieve its goals. Budget allocations are determined based on the original bid submission with all nodes contributing to operating expenses. Some expenditure has been centralised at QUT and budgets are developed by mutual consent of all members of the Executive.

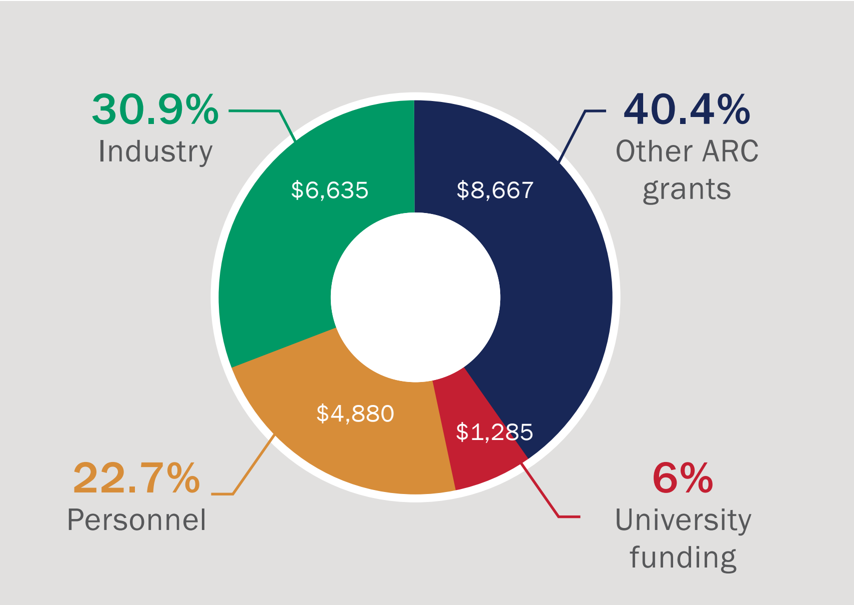

Figure 3: 2016 Halo of industry funding

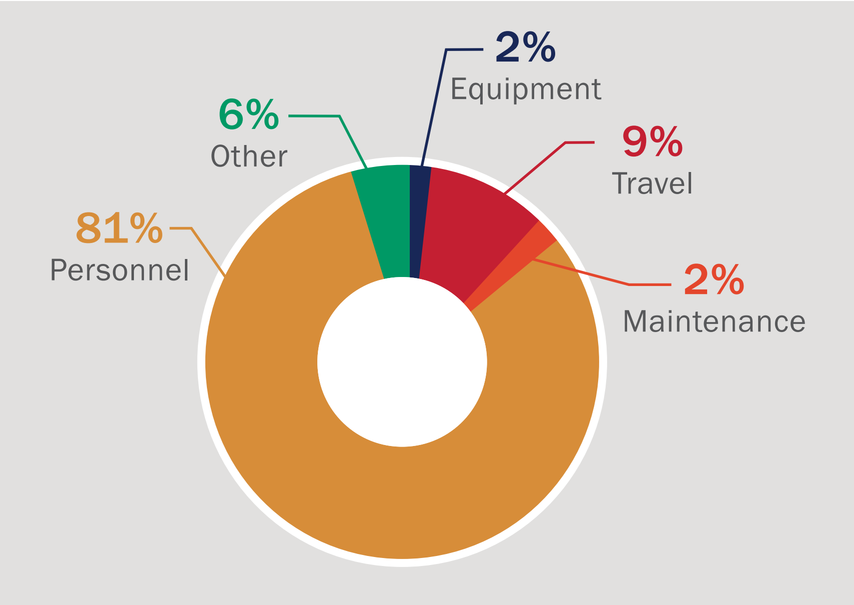

Figure 4: 2016 Allocation of Centre Resources

We have created a pool of untied cash to support strategic initiatives by using the (not guaranteed) indexation funds sent to QUT by the ARC. The primary objective of the fund is to be used as leverage to encourage industry investment in the Centre. We encourage and support travel to other Centre nodes, including overseas partner universities, and to host visits by researchers from our overseas partners.

Centre COO Sue Keay shows a group of students interested in coding around the QUT robotic vision labs.

Summary of Contributions from All Parties

The administering and collaborating organisations are contributing $940k per annum in cash, which amounts to $6.58m in cash over the life of the Centre, and nearly $1m per annum in-kind totalling $7m over the life of the Centre. Our international partner organisations are contributing $139k per annum of in-kind totalling $973k of in-kind over the seven-year life of the Centre. The collaborating organisations (where most Centre researchers are based) will also provide access to a broad range of robotic vision equipment conservatively valued at over $1m per annum ($7m in total). The table below summarises the total cash and in-kind contributions over seven years.

	QUT	ANU	Adelaide	Monash	Oxford	ICL	GT	INRIA	ETHZ	NICTA
Cash	$2.45m	$1.61m	$1.71m	$815k	$0	$0	$0	$0	$0	$0
In-kind	$2.16m	$2.47m	$1.23m	$1.14m	$224k	$127k	$126k	$101k	$185k	$210k

Income and Expenditure Report 2016

Income	2014	2015	2016
ARC Centre Grant distributed as follows	2,714,290	2,714,284	2,714,285
• Monash University	116,929	116,929	116,929
• University of Adelaide	244,400	244,400	244,400
• Australian National University	230,000	230,000	230,000
• Queensland University of Technology	360,000	350,000	350,000
• ARC Indexation	82,027	132,110	180,499
• Interest from QUT (commenced 2016)			6,359
Total Income	3,747,646	3,787,723	3,842,472

Expenditure	2014	2015	2016
Purchased Equipment	42,256	51,829	57,325
Shared Equipment/Facilities	-	815	-
Travel and Professional Development	24,275	253,180	321,076
Maintenance (IT and lab)	8,350	19,639	71,769
Salaries/Personnel exp	373,061	2,337,787	2,813,795
Other	38,002	180,876	192,309
Total Expenditure	485,946	2,844,126	3,456,274
Surplus/Deficit	3,261,700	943,597	386,198
Previous year carry forward	-	3,261,700	4,205,297
Total carry forward surplus	3,261,700	4,205,297	4,591,495

Governance KPIs

Organisational support

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Annual cash contributions from Administering and Collaborating Organisations	Annually						
• QUT		$350,000	$360,000	$350,000	$350,000	$350,000	$350,000
• Monash		$116,400	$116,929	$116,400	$116,929	$116,400	$116,929
• ANU		$230,000	$230,000	$230,000	$230,000	$230,000	$230,000
• Adelaide		$244,400	$244,400	$244,400	$244,400	$244,400	$244,400
Annual in-kind contributions from Administering and Collaborating Organisations	Annually						
• QUT		$309,000	$309,000	$309,000	$250,000	$309,000	$490,759
• Monash		$151,000	$151,000	$151,000	$229,000	$159,000	$180,275
• ANU		$332,000	$114,000	$332,000	$150,000	$316,000	$411,926
• Adelaide		$153,000	$200,000	$153,000	$205,000	$184,000	$211,172
Annual cash contributions from Partner Organisations	Annually	Nil					Nil

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016
Annual in-kind contributions from partner Organisations	Annually						
• Georgia Tech		$18,000	$3,600	$18,000	$18,000	$18,000	-
• INRIA		$14,500	$0	$14,500	$14,500	$14,500	$14,500
• Imperial College		$18,100	$3,600	$18,100	$18,100	$18,100	$18,100
• aData61 / CSIRO		$30,000	$5,800	$30,000	$30,000	$30,000	$30,000
• Swiss Federal Institute		$26,400	$0	$26,400	$26,400	$26,400	$26,400
• Oxford		$32,000	$3,200	$32,000	$32,000	$32,000	$32,000
Other research income sourced by Centre ***End User (industry, public sector, ARC Linkage and Discovery in non-core areas, CRC**	Annually	$0	$1M	$1M	$4.1M	$1.5M	$8M
Number of new organisations collaborating with, or involved in, the Centre	Annually	2	2	2	2	2	4
Level and quality of infrastructure provided to the Centre	At review						

Outputs

2016 Publications

***denotes core Centre Outputs**

Book Chapter (3)

*Chaumette, F., Hutchinson, S., & Corke, P. (2016). Visual Servoing. In Springer Handbook of Robotics pp. 841–866. Cham: Springer International Publishing.

Mahony, R., Beard, R. W., & Kumar, V. (2016). Modeling and Control of Aerial Robots. In Springer Handbook of Robotics pp. 1307–1334. Cham: Springer International Publishing.

Nardi, D., Roberts, J., Veloso, M., & Fletcher, L. (2016). Robotics Competitions and Challenges. In Springer Handbook of Robotics pp. 1759–1788. Cham: Springer International Publishing.

Journal Articles (52)

*Adarve, J. D., & Mahony, R. (2016). A Filter Formulation for Computing Real Time Optical Flow. IEEE Robotics and Automation Letters, 1(2), pp.1192–1199.

*Ball, D., Upcroft, B., Wyeth, G., Corke, P., English, A., Ross, P., et.al. (2016). Vision-based Obstacle Detection and Navigation for an Agricultural Robot. Journal of Field Robotics, 33(8), pp.1107–1130.

*Bewley, A., & Upcroft, B. (2016). From imagenet to mining: Adapting visual object detection with minimal supervision. In 10th International Conference on Field and Service Robotics, FSR 2015; Vol. 113, pp. 501–514. Toronto, Canada: Springer Verlag.

Bialkowski, A., Lucey, P., Carr, P., Matthews, I., Sridharan, S., & Fookes, C. (2016). Discovering Team Structures in Soccer from Spatiotemporal Data. IEEE Transactions on Knowledge and Data Engineering, 28(10), pp.2596–2605.

*Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., et.al. Leonard, J. J. (2016). Past, Present, and Future of Simultaneous Localization And Mapping: Towards the Robust-Perception Age. IEEE Transactions on Robotics, 32(6), pp.1309–1332.

*Cadena, C., Latif, Y., & Reid, I. D. (2016). Measuring the performance of single image depth estimation methods. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) pp. 4150–4157. Daejeon, Korea: IEEE.

Cherian, A., Morellas, V., & Papanikolopoulos, N. (2016). Bayesian Nonparametric Clustering for Positive Definite Matrices. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(5), pp.862–74.

*Dansereau, D. G., Williams, S. B., & Corke, P. I. (2016). Simple change detection from mobile light field cameras. Computer Vision and Image Understanding, 145(April), pp.160–171.

*Ferrari, R., McKinnon, D., He, H., Smith, R., Corke, P., González-Rivero, M., et.al. (2016). Quantifying Multiscale Habitat Structural Complexity: A Cost-Effective Framework for Underwater 3D Modelling. Remote Sensing, 8(2), pp.113.

Grubman, T., Şekercioğlu, Y. A., & Wood, D. R. (2016). Partitioning de Bruijn graphs into fixed-length cycles for robot identification and tracking. Discrete Applied Mathematics, 213, pp.101–113.

Harandi, M. T., Hartley, R., Lovell, B., & Sanderson, C. (2016). Sparse Coding on Symmetric Positive Definite Manifolds Using Bregman Divergences. IEEE Transactions on Neural Networks and Learning Systems, 27(6), pp.1294–1306.

*He, X., Nguyen, C. V., Pratap, M., Zheng, Y., Wang, Y., Nisbet, D. R., et.al. (2016). Automated Fourier space region-recognition filtering for off-axis digital holographic microscopy. Biomedical Optics Express, 7(8), pp.3111.

Hou, X., & Mahony, R. (2016). Dynamic Kinesthetic Boundary for Haptic Teleoperation of VTOL Aerial Robots in Complex Environments. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 46(5), pp.694–705.

*Hu, Q., Paisitkriangkrai, S., Shen, C., van den Hengel, A., & Porikli, F. (2016). Fast Detection of Multiple Objects in Traffic Scenes With a Common Detection Framework. IEEE Transactions on Intelligent Transportation Systems, 17(4), pp.1002–1014.

Khosravian, A., Trumpf, J., Mahony, R., & Hamel, T. (2016). State estimation for invariant systems on Lie groups with delayed output measurements. Automatica, 68, pp.254–265.

*Kristan, M., Matas, J., Leonardis, A., Vojir, T., Pflugfelder, R., Fernandez, G., et.al. (2016). A Novel Performance Evaluation Methodology for Single-Target Trackers. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(11), pp.2137–2155.

*Lai, T., Wang, H., Yan, Y., Chin, T.-J., & Zhao, W.-L. (2016). Motion Segmentation Via a Sparsity Constraint. IEEE Transactions on Intelligent Transportation Systems, PP (99), 1–11. Article in press

*Leitner, J., Harding, S., Förster, A., & Corke, P. (2016). A Modular Software Framework for Eye–Hand Coordination in Humanoid Robots. Frontiers in Robotics and AI, 3, 26.

*Li, H., Li, Y., & Porikli, F. (2016). Convolutional neural net bagging for online visual tracking. Computer Vision and Image Understanding, 153 (December 2016), pp.120–129.

*Li, X., Shen, C., Dick, A., Zhang, Z. M., & Zhuang, Y. (2016). Online Metric-Weighted Linear Representations for Robust Visual Tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(5), pp.931–950.

*Liu, F., Shen, C., Lin, G., & Reid, I. (2016). Learning Depth from Single Monocular Images Using Deep Convolutional Neural Fields. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(10), pp.2024–2039.

*Liu, F., Shen, C., Reid, I., & van den Hengel, A. (2016). Online unsupervised feature learning for visual tracking. Image and Vision Computing, 51(July), pp.84–94.

*Liu, L., Wang, L., & Shen, C. (2016). A Generalized Probabilistic Framework for Compact Codebook Creation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(2), pp.224–37.

*Lowry, S., & Milford, M. J. (2016). Supervised and Unsupervised Linear Learning Techniques for Visual Place Recognition in Changing Environments. IEEE Transactions on Robotics, 32(3), pp.600–613.

*Lowry, S., Sunderhauf, N., Newman, P., Leonard, J. J., Cox, D., Corke, P., & Milford, M. J. (2016). Visual Place Recognition: A Survey. IEEE Transactions on Robotics, 32(1), pp.1–19.

*Ma, B., Huang, L., Shen, J., Shao, L., Yang, M.-H., & Porikli, F. (2016). Visual Tracking Under Motion Blur. IEEE Transactions on Image Processing, 25(12), pp.5867–5876.

*Milford, M., Jacobson, A., Chen, Z., & Wyeth, G. (2016). RatSLAM: Using Models of Rodent Hippocampus for Robot Navigation and Beyond. In Robotics Research pp. 467–485.

*Ming, Y., Li, H., & He, X. (2016). Contour Completion Without Region Segmentation. IEEE Transactions on Image Processing, 25(8), pp.3597–3611.

*Paisitkriangkrai, S., Shen, C., & Hengel, A. van den. (2016). Pedestrian Detection with Spatially Pooled Features and Structured Ensemble Learning. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(6), pp.1243–1257.

*Parra Bustos, A., Chin, T.-J., Eriksson, A., Li, H., & Suter, D. (2016). Fast Rotation Search with Stereographic Projections for 3D Registration. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(11), pp.2227–2240.

*Pepperell, E., Corke, P., & Milford, M. (2016). Routed roads: Probabilistic vision-based place recognition for changing conditions, split streets and varied viewpoints. The International Journal of Robotics Research, 35(9), pp.1057–1079.

Pontes, J. K., Britto, A. S., Fookes, C., & Koerich, A. L. (2016). A flexible hierarchical approach for facial age estimation based on multiple features. Pattern Recognition, 54, pp.34–51.

*Sa, I., Ge, Z., Dayoub, F., Upcroft, B., Perez, T., & McCool, C. (2016). DeepFruits: A Fruit Detection System Using Deep Neural Networks. Sensors, 16(8), pp.1222.

Saha, S. K., Fernando, B., Xiao, D., Tay-Kearney, M.-L., & Kanagasingam, Y. (2016). Deep Learning for Automatic Detection and Classification of Microaneurysms, Hard and Soft Exudates, and Hemorrhages for Diabetic Retinopathy Diagnosis. Investigative Ophthalmology & Visual Science, 57(12), pp.5962–5962.

Schiffner, I., Perez, T., Srinivasan, M. V., Angelov, P., Padian, K., Chiappe, L., et.al. (2016). Strategies for Pre-Emptive Mid-Air Collision Avoidance in Budgerigars. PLOS ONE, 11(9), e0162435.

Seng, W. L., Barca, J. C., Şekercioğlu, Y. A., & Ahmet Ekercio˘ Glu, Y. (2016). Distributed formation control of networked mobile robots in environments with obstacles. Robotica Robotica Robotica, 34(34), pp.1403–1415.

Stacey, G., & Mahony, R. (2016). A Passivity-Based Approach to Formation Control Using Partial Measurements of Relative Position. IEEE Transactions on Automatic Control, 61(2), pp.538–543.

Stratton, P., Hasselmo, M., & Milford, M. (2016). Unlocking neural complexity with a robotic key. The Journal of Physiology, 594(22), pp.6559–6567.

Suddrey, G., Lehnert, C., Eich, M., Maire, F., & Roberts, J. (2016). Teaching Robots Generalisable Hierarchical Tasks Through Natural Language Instruction. IEEE Robotics and Automation Letters, 2(1), pp.201–208.

Tennakoon, R. B., Bab-Hadiashar, A., Cao, Z., Hoseinnezhad, R., & Suter, D. (2016). Robust Model Fitting Using Higher Than Minimal Subset Sampling. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(2), pp.350–362.

*Wang, P., Shen, C., van den Hengel, A., & Torr, P. H. S. (2015). Efficient Semidefinite Branch-and-Cut for MAP-MRF Inference. International Journal of Computer Vision, 117(3), pp.269–289.

*Wang, W., Shen, J., Shao, L., & Porikli, F. (2016). Correspondence Driven Saliency Transfer. IEEE Transactions on Image Processing, 25(11), pp.5025–5034.

*Wang, X., Sekercioglu, Y. A., Drummond, T., Natalizio, E., Fantoni, I., & Fremont, V. (2016). Fast Depth Video Compression for Mobile RGB-D Sensors. IEEE Transactions on Circuits and Systems for Video Technology, 26(4), pp.673–686.

*Warren, M., Corke, P., & Upcroft,

B. (2016). Long-range stereo visual odometry for extended altitude flight of unmanned aerial vehicles. The International Journal of Robotics Research, 35(4), pp.381–403.

Wu, L., Crawford, R., & Roberts, J. (2017). Dexterity Analysis of Three 6-DOF Continuum Robots Combining Concentric Tube Mechanisms and Cable-Driven Mechanisms. IEEE Robotics and Automation Letters, 2(2), pp.514–521.

Wu, L., Wang, J., Qi, L., Wu, K., Ren, H., & Meng, M. Q.-H. (2016). Simultaneous Hand–Eye, Tool–Flange, and Robot–Robot Calibration for Comanipulation by Solving the Problem. IEEE Transactions on Robotics, 32(2), pp.413–428.

Xiao, G., Wang, H., Lai, T., & Suter, D. (2016). Hypergraph modelling for geometric model fitting. Pattern Recognition, 60, pp.748–760.

Xu, J., Denman, S., Fookes, C., & Sridharan, S. (2016). Detecting rare events using Kullback–Leibler divergence: A weakly supervised approach. Expert Systems with Applications, 54, pp.13–28.

*Yang, J., Li, H., Campbell, D., & Jia, Y. (2016). Go-ICP: A Globally Optimal Solution to 3D ICP Point-Set Registration. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(11), pp.2241–2254.

Yu, H., Wu, L., Wu, K., & Ren, H. (2016). Development of a Multi-Channel Concentric Tube Robotic System With Active Vision for Transnasal Nasopharyngeal Carcinoma Procedures. IEEE Robotics and Automation Letters, 1(2), pp.1172–1178.

Zhang, C., Shen, C., & Shen, T. (2015). Unsupervised Feature Learning for Dense Correspondences Across Scenes. International Journal of Computer Vision, 116(1), pp.90–107.

*Zhang, L., Wei, W., Zhang, Y., & Shen, C. (2016). Dictionary Learning for Promoting Structured Sparsity in Hyprspectral Compressive Sensing. IEEE Transactions on GeoScience and Remote Sensing, 54(12), pp.7223–7235.

Conference Papers (85)

*Ajanthan, T., Hartley, R., & Salzmann, M. (2016). Memory efficient max flow for multi-label submodular MRFs. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 5867–5876). Las Vegas, Nevada: IEEE Computer Society.

Abbasnejad, I., Sridharan, S., Denman, S., Fookes, C., & Lucey, S. (2016). Complex event detection using joint max margin and semantic features. In Digital Image Computing: Techniques and Applications (DICTA). Gold Coast, Queensland.

Allibert, G., Mahony, R., & Bangura, M. (2016). Velocity Aided Attitude Estimation for Aerial Robotic Vehicles Using Latent Rotation Scaling. In IEEE International Conference on Robotics and Automation (ICRA 2016).

*Anderson, P., Fernando, B., Johnson, M., & Gould, S. (2016). SPICE: Semantic Propositional Image Caption Evaluation. In Computer Vision - ECCV 2016 (pp. 382–398). Springer International Publishing.

Behendi, S. K., Morgan, S., & Fookes, C. B. (2016). Non-Invasive Performance Measurement in Combat Sports. In Proceedings of the 10th International Symposium on Computer Science in Sports (ISCSS) (pp. 3–10). Springer, Cham.

*Bewley, A., Ge, Z., Ott, L., Ramos, F., & Upcroft, B. (2016). Simple online and realtime tracking. In 2016 IEEE International Conference on Image Processing (ICIP) (pp. 3464–3468).

*Bewley, A., Ott, L., Ramos, F., & Upcroft, B. (2016). Alextrac: Affinity learning by exploring temporal reinforcement within association chains. In 2016 IEEE International Conference on Robotics and Automation (ICRA) (pp. 2212–2218). Stockholm, Sweden: IEEE.

*Bilen, H., Fernando, B., Gavves, E., & Vedaldi, A. (2016). Dynamic image networks for action recognition. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016.

*Chamberlain, W., Leitner, J., Drummond, T., & Corke, P. (2016). A Distributed Robotic Vision Service. In IEEE International Conference on Robotics and Automation (ICRA 2016). Stockholm, Sweden.

*Chin, T.-J., Heng Kee, Y., Eriksson, A., & Neumann, F. (2016). Guaranteed Outlier Removal With Mixed Integer Linear Programs. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016 (pp. 5858–5866). Las Vegas, Nevada.

*Dai, Y., & Li, H. (2016). Simultaneous Correspondences Estimation and Non-Rigid Structure Reconstruction. In 2016 International Conference on Digital Image Computing: Techniques and Applications, DICTA 2016. Gold Coast, Queensland: Institute of Electrical and Electronics Engineers Inc.

*Dai, Y., Li, H., & Kneip, L. (2016). Rolling shutter camera relative pose: Generalized epipolar geometry. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 4132–4140). Las Vegas, United States: IEEE Computer Society.

de Marco, S., Marconi, L., Hamel, T., & Mahony, R. (2016). Output regulation on the Special Euclidean Group SE(3). In 2016 IEEE 55th Conference on Decision and Control (CDC) (pp. 4734–4739).

*Dharmasiri, T., Lui, V., & Drummond, T. (2016). MO-SLAM: Multi object SLAM with run-time object discovery through duplicates - IEEE Xplore Document. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2016. Daejeon, Korea.

*Dingfu Zhou, Dai, Y., & Hongdong Li. (2016). Reliable scale estimation and correction for monocular Visual Odometry. In 2016 IEEE Intelligent Vehicles Symposium (IV) (pp. 490–495). Gothenburg, Sweden: IEEE.

Dunbabin, M. (2016). Autonomous greenhouse gas sampling using multiple robotic boats. In 10th International Conference on Field and Service Robotics, FSR 2015 (Vol. 113, pp. 17–30). Toronto, Canada: Springer Verlag.

Eriksson, A., Bastian, J., Chin, T., & Isaksson, M. (2016). A Consensus-Based Framework for Distributed Bundle Adjustment. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016. Las Vegas, Nevada.

*Felsberg, M., Berg, A., Häger, G., Ahlberg, J., Kristan, M., Matas, J., et.al. (2016). The Thermal Infrared Visual Object Tracking VOT-TIR2015 Challenge Results. In 15th IEEE International Conference on Computer Vision Workshops (pp. 639–651). Santiago, Chile: Institute of Electrical and Electronics Engineers Inc.

*Fernando, B., Anderson, P., Hutter, M., & Gould, S. (2016). Discriminative hierarchical rank pooling for activity recognition. Proc. CVPR.

*Frolov, K., Fripp, J., Nguyen, C. V., Furbank, R., Bull, G., Kuffner, P., et.al. (2016). Automated Plant and Leaf Separation: Application in 3D Meshes of Wheat Plants. In 2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–7). Gold Coast, Queensland: IEEE.

*Garg, R., Vijay Kumar, B. G., Carneiro, G., & Reid, I. (2016). Unsupervised CNN for single view depth estimation: Geometry to the rescue. In Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (Vol. 9912 LNCS, pp. 740–756). Springer Verlag.

*Ge, Z., Bewley, A., McCool, C., Corke, P., Upcroft, B., & Sanderson, C. (2016). Fine-grained classification via mixture of deep convolutional neural networks. In 2016 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 1–6).

*Ge, Z., McCool, C., Sanderson, C., Wang, P., Liu, L., Reid, I., & Corke, P. (2016). Exploiting Temporal Information for DCNN-based Fine-Grained Object Classification. In Digital Image Computing: Techniques and Applications (DICTA). Gold Coast, Queensland.

Ghasemi, A., Denman, S., Sridharan, S., & Fookes, C. (2016). Discovery of facial motions using deep machine perception. In 2016 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 1–7).

*Harwood, B., & Drummond, T. (2016). FANNG: Fast Approximate Nearest Neighbour Graphs. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016 (pp. 5713–5722). Las Vegas, Nevada.

*Hua, M.-D., Trumpf, J., Hamel, T., Mahony, R., & Morin, P. (2016). Feature-based Recursive Observer Design for Homography Estimation. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016.

He, X., Nguyen, C. V., Pratap, M., Zheng, Y., Wang, Y., Nisbet, D. R., et.al. (2016). Adaptive spatial filtering for off-axis digital holographic microscopy based on region recognition approach with iterative thresholding. In SPIE BioPhotonics Australasia (Vol. 10013). Adelaide, Australia: SPIE.

Hou, X., Wang, X., & Mahony, R. (2016). Haptics-aided path planning and virtual fixture based dynamic kinesthetic boundary for bilateral teleoperation of VTOL aerial robots. In 2016 35th Chinese Control Conference (CCC) (pp. 4705–4710).

*Ji, P., Li, H., Salzmann, M., & Zhong, Y. (2016). Robust multi-body feature tracker: A segmentation-free approach. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 3843–3851). Las Vegas, Nevada: IEEE Computer Society.

Johansen, T. A., & Perez, T. (2016). Unmanned aerial surveillance system for hazard collision avoidance in autonomous shipping. In 2016 International Conference on Unmanned Aircraft Systems (ICUAS) (pp. 1056–1065).

*Khan, A., Gould, S., & Salzmann, M. (2016). Deep Convolutional Neural Networks for Human Embryonic Cell Counting. In Computer Vision - ECCV 2016 Workshops (pp. 339–348). Springer International Publishing.

*Kim, J.-H., Cadena, C., & Reid, I. (2016). Direct semi-dense SLAM for rolling shutter cameras. In 2016 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1308–1315). Stockholm, Sweden: IEEE.

*Kneip, L., Sweeney, C., & Hartley, R. (2016). The generalized relative pose and scale problem: View-graph fusion via 2D-2D registration. In IEEE Winter Conference on Applications of Computer Vision, WACV 2016. Lake Placid, United States: Institute of Electrical and Electronics Engineers Inc.

*Koniusz, P., & Cherian, A. (2016). Sparse Coding for Third-order Super-symmetric Tensor Descriptors with Application to Texture Recognition. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016 (pp. 5395–5403). Las Vegas, Nevada.

*Koniusz, P., Cherian, A., & Porikli, F. (2016). Tensor Representations via Kernel Linearization for Action Recognition from 3D Skeletons. In Computer Vision – ECCV 2016 (pp. 37–53). Springer, Cham.

*Kumar, B. G. V., Carneiro, G., & Reid, I. (2015). Learning Local Image Descriptors with Deep Siamese and Triplet Convolutional Networks by Minimising Global Loss Functions. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016 (pp. 5385–5394). Las Vegas, Nevada.

*Kumar, S., Dai, Y., & Li, H. (2016). Multi-body non-rigid structure-from-motion. In Proceedings - 2016 4th International Conference on 3D Vision, 3DV 2016 (pp. 148–156). Stanford, United States: Institute of Electrical and Electronics Engineers Inc.

*Le, H., Chin, T.-J., & Suter, D. (2016). Conformal Surface Alignment With Optimal Mobius Search. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016 (pp. 2507–2516). Las Vegas, Nevada.

*Lehnert, C., Sa, I., McCool, C., Upcroft, B., & Perez, T. (2016). Sweet Pepper Pose Detection and Grasping for Automated Crop Harvesting. In IEEE International Conference on Robotics and Automation (ICRA 2016). Stockholm, Sweden.

*Leitner, J., Chamberlain, W., Dansereau, D. G., Dunbabin, M., Eich, M., Peynot, T., et.al. (2016). LunaRoo: Designing a hopping lunar science payload. In 2016 IEEE Aerospace Conference (pp. 1–12).

Letheren, B., Montes, G., Villa, T., & Gonzalez, F. (2016). Design and flight testing of a bio-inspired plume tracking algorithm for unmanned aerial vehicles. In 2016 IEEE Aerospace Conference (pp. 1–9).

*Liao, Z., & Carneiro, G. (2016). On the importance of normalisation layers in deep learning with piecewise linear activation units. In 2016 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 1–8). IEEE.

Liu, J., Sridharan, S., & Fookes, C. (2016). Recent Advances in Camera Planning for Large Area Surveillance. ACM Computing Surveys, 49(1), 1–37.

*Lin, G., Shen, C., Hengel, A. van dan, & Reid, I. (2015). Efficient piecewise training of deep structured models for semantic segmentation. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016 (pp. 3194–3203). Las Vegas, Nevada.

*Long, G., Kneip, L., Alvarez, J. M., Li, H., Zhang, X., & Yu, Q. (2016). Learning Image Matching by Simply Watching Video. In Computer Vision - ECCV 2016 (pp. 434–450). Springer International Publishing.

*Lu, H., Li, Y., Xu, X., He, L., Li, Y., Dansereau, D., & Serikawa, S. (2016). Underwater image descattering and quality assessment. In 2016 IEEE International Conference on Image Processing (ICIP) (pp. 1998–2002). IEEE.

*McCool, C., Sa, I., Dayoub, F., Lehnert, C., Perez, T., & Upcroft, B. (2016). Visual Detection of Occluded Crop: for automated harvesting. In IEEE International Conference on Robotics and Automation (ICRA 2016). Stockholm, Sweden.

*Mount, J., & Milford, M. (2016). 2D Visual Place Recognition for Domestic Service Robots at Night. In IEEE International Conference on Robotics and Automation (ICRA 2016) (pp. 4822–4829). Stockholm, Sweden.

*Namin, S. R., Alvarez, J. M., Kneip, L., & Petersson, L. (2016). Latent structural SVM with marginal probabilities for weakly labeled structured learning. In 23rd IEEE International Conference on Image Processing, ICIP 2016 (pp. 3733–3737). Phoenix, United States: IEEE Computer Society.

*Ng, Y., Jiang, B., Yu, C., & Li, H. (2016). Non-iterative, fast SE(3) path smoothing. In IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2016 (pp. 3172–3179). Daejeon, Korea: Institute of Electrical and Electronics Engineers Inc.

*Nguyen, C. V., Fripp, J., Lovell, D. R., Furbank, R., Kuffner, P., Daily, H., & Sirault, X. (2016). 3D Scanning System for Automatic High-Resolution Plant Phenotyping. In 2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–8). Gold Coast, Queensland: IEEE.

Nguyen, K., Fookes, C., & Sridharan, S. (2016). Deeper and wider fully convolutional network coupled with conditional random fields for scene labeling. In 2016 IEEE International Conference on Image Processing (ICIP) (pp. 1344–1348).

*Pham, T. T., Eich, M., Reid, I., & Wyeth, G. (2016). Geometrically Consistent Plane Extraction for Dense Indoor 3D Maps Segmentation. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2016. Daejeon, Korea: IEEE.

*Pham, T. T., Hamid Rezatofighi, S., Reid, I., & Chin, T.-J. (2016). Efficient Point Process Inference for Large-Scale Object Detection. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016 (pp. 2837–2845). Las Vegas, Nevada.

*Polok, L., Ila, V., & Smrz, P. (2016). 3D reconstruction quality analysis and its acceleration on GPU clusters. In European Signal Processing Conference (EUSIPCO) (Vol. 2016–November, pp. 1108–1112). Budapest, Hungary.

*Prasai, A. B., Jaiprakash, A., Pandey, A. K., Crawford, R., Roberts, J., & Wu, L. (2016). Design and fabrication of a disposable micro end effector for concentric tube robots Design and Fabrication of a Disposable Micro End Effector for Concentric Tube Robots. In 14th International Conference on Control, Automation, Robotics and Vision (pp. 13–15).

*Qiao, R., Liu, L., Shen, C., & Van Den Hengel, A. (2016). Less is more: Zero-shot learning from online textual documents with noise suppression. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 2249–2257). Las Vegas, Nevada: IEEE Computer Society.

*Rezatofighi, S. H., Milan, A., Zhang, Z., Shi, Q., Dick, A., & Reid, I. (2016). Joint probabilistic matching using m-best solutions. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 136–145). Las Vegas, Nevada: IEEE Computer Society.

*Saleh, F., Ali Akbarian, M. S., Salzmann, M., Petersson, L., Gould, S., & Alvarez, J. M. (2016). Built-in Foreground/Background Prior for Weakly-Supervised Semantic Segmentation. In Computer Vision - ECCV 2016 (pp. 413–432). Springer International Publishing.

*Shirazi, S., Sanderson, C., McCool, C., & Harandi, M. T. (2015). Bags of Affine Subspaces for Robust Object Tracking. In 2015 International Conference on Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–7).

*Skinner, J., Garg, S., Sunderhauf, N., Corke, P., Upcroft, B., & Milford, M. (2016). High-fidelity simulation for evaluating robotic vision performance. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2016. Daejeon, Korea

*Stone, T., Differt, D., Milford, M., & Webb, B. (2016). Skyline-based localisation for aggressively manoeuvring robots using UV sensors and spherical harmonics. In 2016 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5615–5622). Stockholm: IEEE.

*Strydom, M., Jaiprakash, A., Crawford, R., Peynot, T., & Roberts, J. M. (2016). Towards robotic arthroscopy: “Instrument gap” segmentation. In Australasian Conference on Robotics and Automation (ACRA) 2016. Brisbane, Queensland: Australian Robotics & Automation Association.

Suddrey, G., Eich, M., Maire, F., & Roberts, J. (2016). Learning Functional Argument Mappings for Hierarchical Tasks from Situation Specific Explanations. In AI 2016: Advances in Artificial Intelligence (pp. 345–352). Springer, Cham.

*Sunderhauf, N., Dayoub, F., McMahon, S., Talbot, B., Schulz, R., Corke, P., et.al. (2016). Place categorization and semantic mapping on a mobile robot. In 2016 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5729–5736).

Talbot, B., Lam, O., Schulz, R., Dayoub, F., Upcroft, B., & Wyeth, G. (2016). Find my office: Navigating real space from semantic descriptions. In 2016 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5782–5787).

*Vanegas, F., Campbell, D., Eich, M., & Gonzalez, F. (2016). UAV based target finding and tracking in GPS-denied and cluttered environments. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2016 (pp. 2307–2313). Daejeon, South Korea: Institute of Electrical and Electronics Engineers Inc.

*Wang, X., Sekercioglu, A., Drummond, T., Natalizio, E., Fantoni, I., & Fremont, V. (2016). Collaborative Multi-Sensor Image Transmission and Data Fusion in Mobile Visual Sensor Networks Equipped with RGB-D Cameras. In IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2016). Baden-Baden, Germany.

*Wang, Y., Wenjie, Z., Wu, L., Lin, X., Fang, M., & Pan, S. (2016). Iterative views agreement: An iterative low-rank based structured optimization method to multi-view spectral clustering. In IJCAI International Joint Conference on Artificial Intelligence (Vol. 2016–January, pp. 2153–2159). New York, United States.

Wu, L., Wu, K., & Ren, H. (2016). Towards hybrid control of a flexible curvilinear surgical robot with visual/haptic guidance. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2016 (pp. 501–507). Daejeon, Korea: Institute of Electrical and Electronics Engineers Inc.

*Wu, Q., Shen, C., Liu, L., Dick, A., & Van Den Hengel, A. (2016). What value do explicit high level concepts have in vision to language problems? In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 203–212). Las Vegas, Nevada: IEEE Computer Society.

*Wu, Q., Wang, P., Shen, C., Dick, A., & Van Den Hengel, A. (2016). Ask me anything: Free-form visual question answering based on knowledge from external sources. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 4622–4630). Las Vegas, Nevada: IEEE Computer Society.

Xiao, G., Wang, H., Yan, Y., & Suter, D. (2016). Superpixel-Based Two-View Deterministic Fitting for Multiple-Structure Data. In Computer Vision – ECCV 2016 (pp. 517–533). Springer, Cham.

*Yang, J., Li, H., Dai, Y., & Tan, R. T. (2016). Robust optical flow estimation of double-layer images under transparency or reflection. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016 (pp. 1410–1419). Las Vegas, Nevada: IEEE Computer Society.

*Yang, W., Zhong, N., Yang, X., & Cherian, A. (2016). Riemannian sparse coding for classification of PolSAR images. In International Geoscience and Remote Sensing Symposium (IGARSS) (pp. 5698–5701). Beijing, China: Institute of Electrical and Electronics Engineers Inc.

*Yu, X., & Porikli, F. (2016). Ultra-Resolving Face Images by Discriminative Generative Networks. In Computer Vision – ECCV 2016 (pp. 318–333). Springer, Cham.

Zeng, R., Lakemond, R., Denman, S., Sridharan, S., Fookes, C., & Morgan, S. (2016). Vertical Axis Detection for Sport Video Analytics. In 2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–7).

Zhang, J., Zhang, J., Lu, J., Shen, C., Curr, K., Phua, R., et al. (2016). SLNSW-UTS: A Historical Image Dataset for Image Multi-Labeling and Retrieval. In 2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–6). Gold Coast, Australia: IEEE.

*Zhang, L., Wei, W., Zhang, Y., Shen, C., van den Hengel, A., & Shi, Q. (2016). Cluster Sparsity Field for Hyperspectral Imagery Denoising. In Computer Vision - ECCV 2016 (pp. 631–647). Springer International Publishing.

*Zhou, H., Alvarez, J. M., & Porikli, F. (2016). Less Is More: Towards Compact CNNs. In Computer Vision – ECCV 2016 (pp. 662–677). Springer, Cham.

*Zhou, Y., Kneip, L., & Li, H. (2015). A Revisit of Methods for Determining the Fundamental Matrix with Planes. In 2015 International Conference on Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–7). IEEE.

*Zhou, Y., Kneip, L., & Li, H. (2016). Real-time Rotation Estimation for Dense Depth Sensors in Piece-wise Planar Environments. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2016. Daejeon, Korea.

*Zhu, G., Porikli, F., & Li, H. (2016). Beyond Local Search: Tracking Objects Everywhere with Instance-Specific Proposals. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016 (pp. 943–951). Las Vegas, Nevada.

*Zhu, G., Porikli, F., & Li, H. (2016). Robust Visual Tracking with Deep Convolutional Neural Network Based Object Proposals on PETS. In IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPRW 2016 (pp. 1265–1272). Las Vegas, Nevada: IEEE Computer Society.

*Zhuang, B., Lin, G., Shen, C., & Reid, I. (2016). Fast Training of Triplet-based Deep Binary Embedding Networks. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2016. Las Vegas, Nevada.

Edited Collection (2)

Inaba, M., & Corke, P. (2016). Robotics research: The 16th international symposium ISRR. In 16th International Symposium of Robotics Research, ISRR 2013 (Vol. 114). Singapore: Springer Verlag.

Reid, I. (2016). 12th Asian Conference on Computer Vision. Computer Vision and Image Understanding (Vol. 146). Elsevier Science Inc., New York.

Submitted (34)

*Abbasnejad, E., Dick, A., & Hengel, A. van den. (2016). Infinite Variational Autoencoder for Semi-Supervised Learning. Retrieved from http://arxiv.org/abs/1611.07800

*Adarve, J. D., & Mahony, R. (2016). Spherepix: a Data Structure for Spherical Image Processing. IEEE Robotics and Automation Letters, PP (99), 1–1. Retrieved from http://doi.org/10.1109/LRA.2016.2645119

*Aliakbarian, M. S., Saleh, F., Fernando, B., Salzmann, M., Petersson, L., & Andersson, L. (2016). Deep Action- and Context-Aware Sequence Learning for Activity Recognition and Anticipation. Retrieved from http://arxiv.org/abs/1611.05520

*Anderson, P., Fernando, B., Johnson, M., & Gould, S. (2016). Guided Open Vocabulary Image Captioning with Constrained Beam Search. Retrieved from http://arxiv.org/abs/1612.00576

Bangura, M., Melega, M., Naldi, R., & Mahony, R. (2016). Aerodynamics of Rotor Blades for Quadrotors. Retrieved from http://arxiv.org/abs/1601.00733

* Bebis, G., Boyle, R., Parvin, et al. (2016). Advances in Visual Computing. (Vol. 10072). Cham: Springer International Publishing. Retrieved from http://doi.org/10.1007/978-3-319-50835-1

*Cao, Y., Shen, C., & Shen, H. T. (2016). Exploiting Depth from Single Monocular Images for Object Detection and Semantic Segmentation. Retrieved from http://arxiv.org/abs/1610.01706

*Drory, A., Zhu, G., Li, H., & Hartley, R. (2016). Automated detection and tracking of slalom paddlers from broadcast image sequences using cascade classifiers and discriminative correlation filters. Computer Vision and Image Understanding. Retrieved from http://doi.org/10.1016/j.cviu.2016.12.002

* Eich, M., Shirazi, S., & Wyeth, G. (2016). ARTiS: Appearance-based Action Recognition in Task Space. Retrieved from https://arxiv.org/abs/1610.05432

*Fernando, B., Bilen, H., Gavves, E., & Gould, S. (2016). Self-Supervised Video Representation Learning With Odd-One-Out Networks. Retrieved from http://arxiv.org/abs/1611.06646

*Fernando, B., Gavves, E., Oramas, J., Ghodrati, A., & Tuytelaars, T. (2016). Rank Pooling for Action Recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, (99), pp. 1–1. Retrieved from http://doi.org/10.1109/TPAMI.2016.2558148

*Fernando, B., Shirazi, S., & Gould, S. (2016). Unsupervised Human Action Detection by Action Matching. Retrieved from http://arxiv.org/abs/1612.00558

Gould, S., Fernando, B., Cherian, A., Anderson, P., Cruz, R. S., & Guo, E. (2016). On Differentiating Parameterized Argmin and Argmax Problems with Application to Bi-level Optimization. Retrieved from http://arxiv.org/abs/1607.05447

*Harandi, M., & Fernando, B. (2016). Generalized BackPropagation, '{E}tude De Cas: Orthogonality. Retrieved from http://arxiv.org/abs/1611.05927

Harandi, M., Salzmann, M., & Hartley, R. (2016). Dimensionality Reduction on SPD Manifolds: The Emergence of Geometry-Aware Methods. Retrieved from http://arxiv.org/abs/1605.06182

*Leitner, J., Tow, A. W., Dean, J. E., Suenderhauf, N., Durham, J. W., Cooper, M., et al. (2016). The ACRV Picking Benchmark (APB): A Robotic Shelf Picking Benchmark to Foster Reproducible Research. Retrieved from http://arxiv.org/abs/1609.05258

*Li, Y., Lin, G., Zhuang, B., Liu, L., Shen, C., & Hengel, A. van den. (2016). Sequential Person Recognition in Photo Albums with a Recurrent Network. Retrieved from http://arxiv.org/abs/1611.09967

*Lin, G., Milan, A., Shen, C., & Reid, I. (2016). RefineNet: Multi-Path Refinement Networks for High-Resolution Semantic Segmentation. Retrieved from http://arxiv.org/abs/1611.06612

*Lin, G., Shen, C., Hengel, A. van den, & Reid, I. (2016). Exploring Context with Deep Structured models for Semantic Segmentation. Retrieved from http://arxiv.org/abs/1603.03183

*Liu, L., Li, H., & Dai, Y. (2016). Robust and Efficient Relative Pose with a Multi-camera System for Autonomous Vehicle in Highly Dynamic Environments. Retrieved from http://arxiv.org/abs/1605.03689

*Milan, A., Leal-Taixe, L., Reid, I., Roth, S., & Schindler, K. (2016). MOT16: A Benchmark for Multi-Object Tracking. Retrieved from http://arxiv.org/abs/1603.00831

*Purkait, P., Chin, T.-J., Sadri, A., & Suter, D. (2016). Clustering with Hypergraphs: The Case for Large Hyperedges. IEEE Transactions on Pattern Analysis and Machine Intelligence, PP (99), 1–1. Retrieved from http://doi.org/10.1109/TPAMI.2016.2614980

*Rezatofighi, S. H., G, V. K. B., Milan, A., Abbasnejad, E., Dick, A., & Reid, I. (2016). DeepSetNet: Predicting Sets with Deep Neural Networks. Retrieved from http://arxiv.org/abs/1611.08998

*Sünderhauf, N., Pham, T. T., Latif, Y., Milford, M., & Reid, I. (2016). Meaningful Maps $-$ Object-Oriented Semantic Mapping. Retrieved from http://arxiv.org/abs/1609.07849

*Teney, D., Liu, L., & Hengel, A. van den. (2016). Graph-Structured Representations for Visual Question Answering. Retrieved from http://arxiv.org/abs/1609.05600

*Tsai, D., Dansereau, D. G., Martin, S., & Corke, P. (2016). Mirrored Light Field Video Camera Adapter. Retrieved from http://arxiv.org/abs/1612.05335

*Wang, P., Liu, L., Shen, C., Hengel, A. van den, & Shen, H. T. (2016). Hi Detector, What's Wrong with that Object? Identifying Irregular Object From Images by Modelling the Detection Score Distribution. Retrieved from http://arxiv.org/abs/1602.04422

*Wang, P., Wu, Q., Shen, C., & Van Den Hengel, A. (n.d.). The VQA-Machine: Learning How to Use Existing Vision Algorithms to Answer New Questions. Retrieved from https://arxiv.org/abs/1612.05386

*Wu, Q., Shen, C., Hengel, A. van den, Wang, P., & Dick, A. (2016). Image Captioning and Visual Question Answering Based on Attributes and External Knowledge. Retrieved from http://arxiv.org/abs/1603.02814

*Wu, Z., Shen, C., & Hengel, A. van den. (2016). Wider or Deeper: Revisiting the ResNet Model for Visual Recognition. Retrieved from http://arxiv.org/abs/1611.10080

Xu, J., Fookes, C., & Sridharan, S. (2016). Automatic Event Detection for Signal-based Surveillance. Retrieved from http://arxiv.org/abs/1612.01611

*Yang, J., Ren, P., Chen, D., Wen, F., Li, H., & Hua, G. (2016). Neural Aggregation Network for Video Face Recognition. Retrieved from http://arxiv.org/abs/1603.05474

*Zhu, G., Porikli, F., & Li, H. (2016). Not All Negatives are Equal: Learning to Track with Multiple Background Clusters. IEEE Transactions on Circuits and Systems for Video Technology, PP (99), 1–1. Retrieved from http://doi.org/10.1109/TCSVT.2016.2615518

*Zhuang, B., Liu, L., Li, Y., Shen, C., & Reid, I. (2016). Attend in groups: a weakly-supervised deep learning framework for learning from web data. Retrieved from https://arxiv.org/abs/1611.09960

Our Centre's Robotic Vision Summer School is held annually at ANU's coastal Kioloa campus.

Glossary

3DV - International Conference on 3D Vision

ACCV - Asian Conference on Computer Vision

ACRA - Australiasian Conference on Robotics and Automation (run by ARAA Australian Robotics and Automation Association)

AI - Artificial Intelligence

AI - Associate Investigator

ANU - Australian National University

ARAA - Australian Robotics and Automation Association

ARC - Australian Reseacrh Council

bn - billion

CAC - Centre Advisory Committee

CEC - Centre Executive Committee

CI - Chief Investigator

COTSBOT - Crown of Thorns Starfish Robot

CVPR - IEEE Conference on Computer Vision and Pattern Recognition

DICTA - International Conference on Digital Image Computing: Techniques and Applications (premier conference of the Australian Pattern Recognition Society (APRS)

DSTO - Defence Science and Technology Organisation

ECCV - European Conference on Computer Vision

EOI - Expression of Interest

EUAB - End-User Advisory Board

FPGA - Field-Programmable Gate Array

IBVS - Image-Based Visual Servo

ICCP - International Conference on Computational Photography

ICIP - IEEE International Conferenc eonf Image Processing

ICRA - IEEE International Conference on Robotics and Automation

IEEE - Institute of Electrical and Electronics Engineering

IJCAI - International Joint Conference on Artificial Intelligence

IET - Institution of Engineering and Technology

IROS - International Conference on Intelligent Robots and Systems

ISWC - International Semantic Web Conference

KPIs - Key Performance Indicators

m - million

MOOC - Massive Open Online Course

MVG - Multi View Geometry

NICTA - National ICT Australia Limited

NIPS - Conference on Neural Information Processing Systems

NRP - National Research Priority

PhD - Doctor of Philosophy

PI - Partner Investigator

QUT - Queensland University of Technology

RHD - Research Higher Degree

RF - Research Fellow

SLAM - Simultaneous Localisation and Mapping

SRPs - Science and Research Priorities

Adelaide - University of Adelaide

VOS - Vision Operating System

Key Terms

algorithm - is a procedure or formula for solving a problem, typically implemented by computer software. For example, there are algorithms to help robots determine their location in the world, to navigate safely, to process images or recognise objects.

artificial intelligence (AI) - the simulation of intelligent behaviour in machines.

autonomous - without human intervention.

Bayesian (Bayes) nets (networks) - are graphical representations for probabilistic relationships among a set of random variables.

computer vision - methods for acquiring, processing, analysing and understanding images using a computer.

deep learning - a method of machine learning based on neural networks with many and varied layers that are able to form representations of data based on large amounts of training data.

homography - the relationship between any two images of the same planar surface in space.

machine learning - a type of artificial intelligence providing computers with the ability to learn based on large amounts of training data without needing to be explicitly programmed.

neural network - a computer system very loosely modeled on neurons and synaptic connections found in biological brains.

semantics - automatically applying human meaningful terms like ‘kitchen’ or ‘coffee cup’ to places or objects in the robotic vision system’s environment. Important to help robots understand their environment by recognising different features and labelling/ classifying them.

servo - a system that uses negative feedback to automatically correct its error.

SLAM (Simultaneous Localization and Mapping) - a robotics algorithm that allows a robot to determine its position in an environment while at the same time constructing a map of its environment.

support vector machine - an SVM classifies data by finding the best hyperplane that separates all data points of one.

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.

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

Robotic Vision Industry Partners

Australian Centre for Robotic Vision Head Office Contact Details

Queensland University of Technology	p	+61 7 3138 7549
2 George Street	e	info@roboticvision.org
Brisbane QLD 4000 AUSTRALIA	w	roboticvision.org

Design by Studio 55

creating robots that see...

roboticvision.org