AUSTRALIAN CENTRE FOR ROBOTIC VISION

ANNUAL REPORT 2015

creating robots that see...

...and will protect the environment

Vision

creating robots that see and understand their environment and work safely with us to help meet the challenges of the future

Values

Cool research – Impact – Creativity – Entrepreneurship – Community – Enabling – Fun – Looking Outwards

Objectives

SCIENCE

Create robots that see and understand their environment

INTEGRATE

Integrate robotics and computer vision together to create critical mass in robotic vision and to build enabling technology platforms

CULTURE

Establish and maintain an exciting, high-energy collaborative atmosphere that supports world class research

ENGAGE

Engage with stakeholders about robotics and robotic vision technologies and the impact these will have on society and the way we work

TRANSFORM

Transform industry, generate wealth for the community, prepare for the challenges of the future

Front cover image: Composite image showing underwater Crown-of-Thorns Starfish (COTS) robot (bot) using computer vision to identify COTS on Australia's Great Barrier Reef. The development of COTSbot was funded by QUTBlubox, QUT's commercialisation arm (see COTSbot story p. 54).

Table of Contents

About the Centre

We are an Australian research centre tackling the critical and complex challenge of applying computer vision to robotics. 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. Robotic vision is about integrating robotics with computer vision to enable robots to see, understand and interact with the world.

WHO WE ARE

We believe robotic vision is the key technology that will allow robotics to change the way we live and work. 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. We will create and nurture the next generation of research talent — training those who are just beginning a research career, and embedding research and innovative talent within industry. The Centre will help secure the future of our economy and build sustainable partnerships across the research sector and public and private enterprises. We are an interdisciplinary research team from four leading Australian research universities; Queensland University of Technology (QUT), The University of Adelaide (Adelaide), The Australian National University (ANU), and Monash University as well as DATA 61 (previously NICTA), and overseas universities and research organisations including INRIA Rennes, Georgia Institute of Technology, Imperial College London, the Swiss Federal Institute of Technology (ETH) Zurich, and the University of Oxford.

CENTRES OF EXCELLENCE

We receive $19m in public funding contributed 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) funding 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. 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.

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 as 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 modern computers and advanced 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 from both Australian and international research organisations (see p. 72) was assembled and an expression of interest (EOI) submitted to the ARC in May 2013. The team's EOI was announced as successful in June 2013, proceeding to a full proposal submitted in August 2013. After supplying a rejoinder in response to review comments, the Australian Centre for Robotic Vision was short-listed and then interviewed in Canberra in November 2013. On 17th December 2013, the Hon Christopher Pyne MP, then Minister for Education and Training, approved the allocation of $285m over seven years for 12 new ARC Centres of Excellence, including the ARC Centre of Excellence for Robotic Vision. Our Centre commenced on 1st January 2014 and became fully operational on 21st July 2014 following the signing of legal agreements between all collaborating organisations.

Oxford Univewrsity, Oxford, UK
Imperial, College, London, UK
ETH, Zurich, Switzerland
INRIA, Rennes, France

Geogia Tech, Atlanta, Georgia, USA

QUT, Brisbane
DAT61, Sydney, Canberra
University of Adelaide, Adelaide
Monash, Melbourne
ANU, Canberra

Map showing location of Centre partners

About this report

REPORT DESCRIPTION

Our report covers the activities of the ARC Centre of Excellence for Robotic Vision for the 2015 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 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 2015. 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

Dr Anjali Jaiprakash (QUT Medical Robotics) demonstrates robotics to high school students in her role as an ATSE Science Ambassador

Centre Performance

IT'S TIME FOR ROBOTS

Robots are autonomous machines that can move within their environment and perform physical tasks. 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.

We believe that vision is the key to unleashing the full potential of robots. The hype around robotics and computer vision has been building for more than 50 years but to date, our technologies have fallen short of our imagination — advanced machines in movies far surpass the abilities of modern robots. Thanks to accelerating advances in sensors, actuators, computation and machine-to-machine communication, we are reaching a tipping point in both robotics and computer vision where robots are gaining visual perception. This is the step required to make robots truly useful and safe to humans. 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 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.

By 2020 computers will be 30x faster than today, cameras will be cheap, and the need for robots will be pressing. The accelerating pace of technological change will provide both challenges and opportunities. In their 2015 book, "No Ordinary Disruption, The Four Forces Breaking All the Trends", McKinsey Directors Dobbs, Manyika, & Woetzel point to advanced robotics, autonomous vehicles and the automation of knowledge work as three of 12 disruptive technology trends, which will have significant impact on the natural forces of market competition. Gartner lists advanced machine learning, and "autonomous agents and things", as two of the top ten strategic technology trends for 2016. 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.

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
- reduce the cost of healthcare
- 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, not barcodes
- 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

QUT's Baxter robot by Rethink Robotics demonstrates its capabilities playing Connect 4 with QUT's Matt Kimball and Queensland Minister for Science and Innovation the Hon Leeanne Enoch MP

involved in using human divers (see p. 54)

- 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% (see p. 51)
- launched free global university level courses in robotic vision – over 30,000 students from 100 countries have taken the courses so far (see p. 68)
- created a vegetable picking robot that picks the right fruit the right way every time 24x7 (see p. 30)
- been active in the debate about robots, jobs and society (see p. 12).

Robotic vision is the key enabling technology that will allow robotics to transform labour-intensive industries, disrupt stagnant markets, and ensure robots become a ubiquitous feature of the modern world. Our Centre's goal is 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 be able to cooperate effortlessly with human co-workers. In this way we can create robots that see and understand their environment, working safely with us to help meet the challenges of the future.

Highlights for 2015

	22	JOURNAL ARTICLES
	15	ROBOTIC VISION WORKSHOPS (ACADEMIC & INDUSTRY)
	54	BRIEFINGS
	23	NEW PHD STUDENTS
	94	MEDIA PIECES
	98,000	WEBSITE HITS
	17	AWARDS

Director's Report

2015 was a year of growth and adaptation. In 2015 both Professor Peter Corke and Professor Ian Reid led the Centre. Director Peter Corke took a sabbatical visiting Professor Paul Newman's Mobile Robotics Group at Oxford University (one of our international partner organisations) from July-Dec 2015. During this time, ARC Laureate Fellow Professor Ian Reid served as Acting Director. The Director's Report is from both Peter and Ian.

Centre Director
PETER CORKE

Centre Deputy Director
IAN REID

CENTRE GROWTH

On 9th March 2015, our Centre was officially launched by (then) Federal Minister for Education, Christopher Pyne at a well-publicised event at QUT's The Cube (see Centre Launch p. 65). A significant early challenge for our Centre has been recruitment. We now have sixteen research fellows (the most recent appointees commenced late in the second quarter of 2015), more than thirty PhD students, a Chief Operating Officer, Centre Coordinator, and part-time support in the form of a Stakeholder Engagement Coordinator, Finance and Administration Officer, and three of the eventual four node

administrators (QUT, Monash and Adelaide) (see People p.72).

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 also accepted the nominations of three additional Associate Investigators (AIs), Dr Matthew Dunbabin (Environmental Robotics), Dr Laurent Kneip (Computer Vision) and Professor Ross Crawford (Medical Robotics) bringing our total number of AIs to 17. We created a new category of affiliate membership and have accepted four affiliate members and are introducing an industry affiliate program.

BOOTSTRAPPING THE CENTRE

We applied lean start-up principles to get the Centre going and have put in place internal management processes (policies and guidelines), a wiki-style intranet (Confluence by Australian company Atlassian) to facilitate communication between researchers at all nodes, and launched our external web site (roboticvision.org). By the end of 2015 the Centre had held 45 meetings of the Centre Executive, four meetings of all Chief Investigators as well as scheduling monthly research fellow, PhD researcher and research theme meetings via video-conference (Citrix Go-To-Meeting). We have been active with inter-node visits: Research Fellows and CIs making numerous short visits, as well as several longer visits to partner institutions by CIs and research fellows (CI Corke to Oxford, CI Hartley to ETHZ, CI Milford to Imperial, RF Sünderhauf to ETHZ). PI Andy Davison visited the Centre in March/April 2015 for 3 weeks, spending 2 weeks at the Adelaide node and a week at our Robotic Vision Summer School (see Robotic Vision Summer School p.69). François Chaumette (INRIA) also visited both our ANU and QUT nodes in November 2015. At present, Zetao Chen, one of our PhD researchers, is on an extended visit (to Adelaide from QUT for 6 months) and

another, Will Chamberlain, has made two shorter visits to Monash from QUT. A challenge for 2016 will be to encourage more and longer internode visits.

BRINGING ROBOTICS AND COMPUTER VISION TOGETHER

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) and large audiences. Examples include the ICRA 2015 Workshop on Robotic Vision: challenges and opportunities, (CIs Corke, Drummond and Greg Hager, Johns Hopkins University); CVPR 2015 Workshop on Semantics in Visual Reconstruction, Localisation and Mapping (CIs Reid, Gould and Silvio Savarese, Stanford);

RSS 2015 Workshop on Mobile Sensors: Setting future goals and indicators of progress for SLAM, (RFs Latif, Cadena and CI Reid). A special issue of the International Journal of Robotics Research on Robotic Vision was published in 2015 (eds. CI Corke, Jana Kosecka (George Mason University) and Eric Marchand (University de Rennes). A new special issue in this series will be edited by CIs Corke, Drummond and Reid and the call for papers is currently open.

BUILDING CENTRE CULTURE

We have held two symposia since starting in July 2014. Our first annual symposium "RoboVis" was held at ANU (Canberra) in November 2014 and was the first time all members of the Centre; CIs, AIs, research fellows, PhD researchers, and friends of the Centre came 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. The second RoboVis was held in the Barossa Valley in June 2015 and attended by 80 people. It incorporated a full training day for research fellows and PhD researchers covering a range of topics customised for each group. We also held our first Robotic Vision Summer School (RVSS) (see Robotic Vision Summer School p. 69).

ENGAGING OUTSIDE THE CENTRE

In 2015 the Centre ran a series of workshops for industry in both capital cities and regional areas. We partnered with AusIndustry, other local computer vision and robotics groups, with the aim to promote and educate people about robotic vision. The Centre garnered a lot of media coverage, provided globally accessible robotic vision education programs, formed networks to help

Refurbished space for the Centre of Excellence at the University of Adelaide

support people learning code, and have hosted a range of visits to Centre nodes from a wide range of stakeholders (Communication, Engagement, Outreach and Networks p. 59).

We aim to expand our outreach and industry engagement programs. It is encouraging that the number of industry contacts is picking up, and that industry is increasingly approaching us. The challenge for 2016 is to convert these into collaborative research projects.

DEMONSTRATIONS OF ROBOTIC VISION

Our Centre includes a $3m program in agricultural robotics, funded by the Queensland Department of Agriculture and Fisheries. One outcome of this program is a new field robot called the AgBot II, which has been developed for broadacre weeding and fertilising, by Centre-affiliated researchers at QUT. An agreement for manufacture of the robot is being negotiated with an Australian company. The AgBot has been field tested and ongoing development is concerned with

mechanical and microwave based weed destruction controlled by vision-based plant classification (see Agriculture p. 51).

At the end of 2014 Centre researchers at QUT received funding to develop a proof-of-concept visually-guided underwater robot to navigate, locate and cull crown-of-thorns starfish (COTS). These starfish are one of the most significant threats to the Great Barrier Reef, and ever more frequent outbreaks of COTS are destroying large areas. A successful prototype robot, "COTSbot" has been developed with negotiations commencing around commercial opportunities. As well as the benefits the robot can provide in controlling the number of COTS, the robot also collects a range of important information about the health of the Great Barrier Reef. We will work with the ARC Centre of Excellence in Mathematical and Statistical Frontiers (ACEMS) to interpret the data collected by the robot to further understanding about how to control the threat these starfish pose and how to manage the reef (see COTSbot p. 54).

CENTRE FACILITIES

Our Centre enjoys new purpose-built facilities at ANU. A major refurbishment of space at the University of Adelaide was completed in Q1 2015 and a similar refurbishment at QUT was completed in September 2015. A successful ARC Linkage Infrastructure, Equipment and Facilities (LIEF) grant titled "Computational infrastructure for developing deep machine learning models" was awarded to CIs Ian Reid, Stephen Gould, Anton van den Hengel, Gustavo Carneiro, Chunhua Shen and Peter Corke; RF Niko Sünderhauf; and other Australian researchers. The LIEF will provide AUD $400k to significantly increase the GPU computational capability available to Centre researchers.

CHALLENGES

The major challenges facing the Centre are around internal engagement and effectively harnessing our collective expertise in robotics and vision as robotic vision. Perhaps the biggest challenge is that robotics and computer vision

have very different approaches to research, the former encouraging real-world experimentation and the latter focussed on evaluation of standard datasets. As a Centre we need to better bridge this gap, and make it easier for robotics researchers to use cutting-edge computer vision algorithms, and for vision researchers to trial their results on robots. Increased interaction and visits between research nodes is the key to bridging the gap between the disciplines. With both robotics and computer vision booming, retention of key research staff will be increasingly challenging. Already our researchers are being head-hunted, particularly to the self-driving car industry. Our aim is to provide as stimulating and energetic environment as possible so that we can retain – and continue to attract – high quality researchers.

OUTLOOK FOR THE CENTRE

In August 2015, we initiated a review of our entire research program, matching the capabilities of the personnel we hired against our objectives, reflecting on our first year and mindful that the research landscape in the area has changed dramatically in the last two years. The inexorable rise of deep learning as a critical tool in both robotics and computer vision was barely envisaged when we first formed the Centre. We are well-positioned to exploit and lead this boom with our current personnel and expertise. In 2016 we will implement our new research strategy and continue to grow our team, welcoming more PhD researchers to the Centre and engaging them in our training and mentoring programs. With a growing team and renewed focus on what we need to achieve, the next twelve months will see us delivering in all the key areas where we promised our Centre would have impact; the development of breakthrough technologies, a true partnership between robotics and computer vision, and a growing collective awareness of the importance of robotic vision amongst industry, government and the general community.

We would like to thank all our colleagues - Centre researchers, operations staff and executive - for their contributions to a great collective outcome.

PETER CORKE
Centre Director

IAN REID
Centre Deputy Director

Market Trends

Over the lifetime of the Centre the computing power to cost ratio is likely to increase by a factor of 30, cameras will become cheaper, more highly integrated with computation and ever more ubiquitous – 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 (Figure 1). Computer Vision and Robotics are both booming technologies. Sales of machine vision gear that help computers understand their surroundings and identify images, much like a human can, are rapidly growing - by 22% in the first quarter of 2015 in the U.S alone. Camera and vision sensor sales reached $2.5 billion in 2015 (Fortune http://tinyurl.com/gthklbr) and large internet companies like Google and Facebook are key investors in vision technology and the software to make it work. The trend for leading companies like Google, Facebook, Uber, Qualcomm, and Amazon to acquire companies that specialise in robotics, vision and machine learning, slowed in 2015 although more than $1b worth of acquisitions were made. In 2015, Google acquired vision technology companies Odysee, Digisfera and Fly Labs; Facebook acquired computer vision companies Quickfire, Surreal Vision and Pebbles; and Amazon spent $690m on acquisitions in 2015 including IoT company 2lemetry, chip maker Annapurna Labs, and video processing company Elemental Technologies.

According to International Data Corporation (IDC http://tinyurl.com/glzu2nv), robotic capabilities are expanding while increasing investment in robot development is driving competition and helping to bring down the cost of robots, with two of the fastest adopters being process
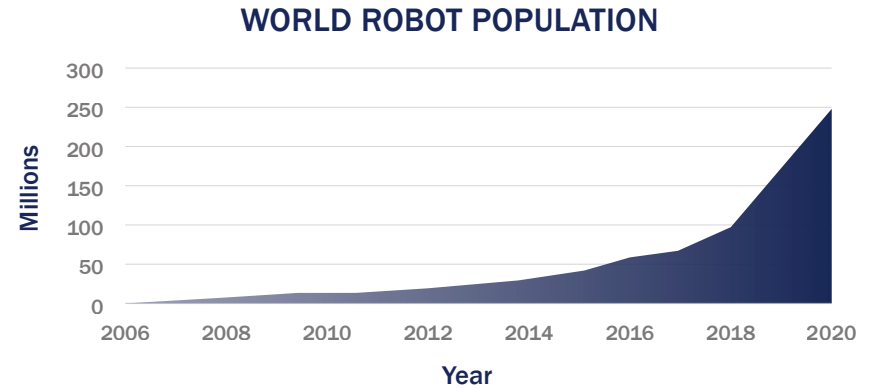

WORLD ROBOT POPULATION

Figure 1: World robot population interpolated from global robotics industry data sourced from the International Federation of Robotics (www.ifr.org)

manufacturing and healthcare. Spending $100,000 for a robot today will get you a machine that can do twice as much as a robot 10 years ago that used to cost the same price (Merrill Lynch http://tinyurl.com/gso39vu). A related boom in robotics-services industries is also expected, with the market for applications management, education & training, hardware deployment, computer integration, and consulting expected to grow to $32 billion in 2019. This will result in increased spending on computer hardware (servers & storage) and software (command & control, network infrastructure, and robotics-specific applications).

Worldwide spending on robotic systems, one of the main beneficiaries of advances in computer vision, reached $71 billion in 2015. Spending is expected to nearly double by 2019 to reach $135 billion (IDC http://tinyurl.com/glzu2nv), with most of that growth forecast for the Asia-Pacific region. Already the Asia-Pacific region accounts for 65% of current ($46.8 billion) robotic spending. China is the biggest buyer of robots, purchasing 25% of the world's industrial robots in 2014 (Merrill Lynch http://tinyurl.com/gso39vu). With a robotic density of 35 industrial robots per 10,000 workers, compared to South Korea's 478 per

10,000, Chinese spending on robotics still has plenty of room for growth, especially as China's population ages and the size of its labour force diminishes.

Advanced robotics can improve the productivity of many Australian industry sectors, the networked structure of the modern workplace allows robots to seamlessly integrate into operations. Ideally, in the current climate of demand, Australia could create a new export industry in robotic and vision technologies. Although the makers of such technologies in Australia are currently fragmented, it is anticipated that Australia will follow world trends in the increased demand for industrial and service robots. This will trigger the need for improved collaboration, more cohesive local supply chains, and skilled people with expertise to lead the new era of robotics and vision companies. The Centre anticipates playing a critical role in promoting Australian robotic and vision technologies by improving collaboration between researchers and industry.

Social Impact of Robotics

Our Centre is sensitive to community concerns about the impact of robots on the workforce, as well as privacy issues that arise from robots equipped with cameras. We will continue to proactively contribute to public education and debate in these areas. As part of this engagement, Centre Director Peter Corke appeared on SBS Insight as an invited guest. The SBS Current Affairs program aired on 21 April 2015 (http://tinyurl.com/haumjbx). The Episode was called "Trusting Robots" and the program's studio audience included researchers, users of robots, a cross-section of the community representing different views of the impact that robotics may have on people.

The show's presenter demonstrated some of the capabilities of robots by getting a NAO robot to perform a dance to Michael Jackson's "Thriller" and Rethink Robotics' Baxter industrial robot to stack dishes. The discussion revolved around the two juxtaposed points of how robots could be of benefit to people and the negative impact robots may have on people and society.

In one example a small robot was being used with a family who had two autistic children to assist in building their communication and social interaction skills. The children's mother said that interacting with the robot had provided a stimulus that had improved the children's ability to read and learn and their social interaction skills.

In a second example a robotic seal was being trialled by Professor Wendy Moyle of Griffith University, to improve dementia patients' quality of life. The robot was used as a substitute for a pet, such as a cat or dog. The argument for this approach was that while it has been shown that pets have been highly beneficial to the care of dementia patients there are challenges in maintaining pets in a care facility. The robotic seal provided an alternate solution and it could be deployed in numbers without increasing the workload of carers while providing similar benefits to the patient as a pet would.

While the initial results of this trial are positive the argument was put forward that robot seals are no substitute for human interaction and that reliance on this form of therapy could ultimately prove detrimental. When quizzed, our Centre Director, Peter, indicated that he hoped that the removal of human interaction would never occur.

Peter was also asked whether people would lose jobs as result of robots taking over the tasks people could do. In response he said "Australia has an aging population and as a result the country's productivity would seriously decrease and that it would become less competitive". He also said "Robots will be able to improve productivity by making up the shortfall in Australia's future workforce".

Associate Professor Rob Sparrow offered the opposing view of robots as displacing people from most areas of work, which would result in a society composed of extremes with the very poor and the very rich, with the former vastly outnumbering the latter.

The program did not seek to resolve the tension between the benefits or detrimental impact of robots, rather it sought to provoke a discussion around this disruptive technology.

Imperial Storm Troopers, courtesy of the Redback Garrison of the 501st Legion, consider purchasing an agricultural robot at QUT's Robotronica robotics festival *(photo courtesy of Owen Bawden)*

Final performance at Robotronica involving custom-made Perff robots and rock orchestra, Deep Blue.

ROBOTRONICA: A celebration of all things robotic

Thousands of people turned out to see a "glimpse into the possibilities of the future," at least when it comes to robots. Queensland University of Technology's robot expo, "Robotronica," was held on 23 August, and the Centre for Robotic Vision was truly one of the stars. The event was held at QUT's Gardens Point Campus, and the Centre had a presence in a tent, on the lawn, and even in the pool. If that's not enough, the Centre also had a starring role on stage as part of the closing performances, which featured robots performing with humans.

The Nao robots, as always, proved to be very popular with kids and adults. At the Centre's tent, kids of all ages had the opportunity to play with and hold the Naos. Jonathan Roberts is a Professor of Robotics at QUT, and also an Associate Investigator with the Centre. He told the ABC that robots would become more common as prices fell in the future. "It's all about the future and trying to show people what's just around the corner and what's happening in 10 to 15 years," said Jonathan.

Large numbers of people wanted to check out our vision-enabled agricultural robot AgBotII, as it was put through its paces. Even some Imperial Storm Troopers checked out the unit (with talk of a potential sale in the offing!).

Getting a few laps in the pool was Associate Investigator Matt Dunbabin's creation. His bright yellow Inference System robots were put on display for all to see. Normally deployed on waterways and dams, it's hoped these units will use their vision and sampling capabilities to provide 24 hour real-time data on environmental conditions and threats.

The day finished with a spectacular closing theatrical performance on stage. QUT's robotic group, the Perffbots, built a troupe of robots that performed on stage alongside the rock orchestra, Deep Blue.

The show was a feast of light, music, dance and robotic choreography at its best. Dancing cubes, and what looked like robotic tortoises, cleaning robots, a skipping robotic cube and even a robotic gymnast all had the crowd cheering as they performed on stage, all while being cheered on by our custom-built applausebots (clapping robotic cubes).

It was a truly engaging performance, and a day, that showed robots and robotic vision capabilities in a different light.

Research Performance

Our Centre will achieve breakthrough science and technology in robotic vision by addressing four key research objectives: Semantic Vision (SV), Vision and Action (VA), Robust Vision (RV), Algorithms and Architecture (AA).

RESEARCH STRUCTURE

Semantic Vision (SV) is key to robots learning visual contexts and the variable appearance of objects in the world. Vision and Action (VA) is key to robots using vision to move and moving to resolve ambiguity in the observed scene, basing their inference on evidence collected from different viewpoints. Robust Vision (RV) is concerned with developing and characterising sensing technologies and creating robust algorithms that allow robots to see in all viewing conditions. Finally, Algorithms and Architecture (AA) provides the computing and communications platform on which image capture and processing is built, allowing robotic systems to run in real-time and be deployed in large-scale real-world environments.

Together the four research objectives form the Centre's research themes, which, until now, have served as organisational groupings of the Centre's research projects. In the latter half of 2015, the Centre's Executive embarked on a wide-ranging audit and assessment of the Centre's research, particularly considering whether the research structure was best able to take advantage of the many rapid advances in the field and also to deliver on the Centre's research objectives. This has led to a comprehensive restructuring of our research program with a new level of accountability created below Theme level in the form of six new Research Programs, each to be led by a Centre Chief Investigator. The new Research Programs map directly to our original four research themes (see Figure 2). Themes will continue in 2016 but will not have discrete Theme Leaders, instead all programs will report to a new Centre Research Committee, which will meet monthly 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.

Centre projects are a quantum of research (PhD students, RFs, CIs, AIs) within each of our six new research programs, and have 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. In 2016 new projects will be created by our Research Committee, which will determine 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. An Honours project, which is an undergraduate project supervised by a CI, RF or AI, similarly must fall within the research areas of the Centre. All projects must develop technology that demonstrates robotic vision capability, 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.

Figure 2: Relationship between Centre Research Themes, Projects and new Centre Research Programs (commencing 2016)

Research Themes	Semantic Vision		Vision & Action		Robust Vision	Algorithms & Architecture
NEW Research Programs (commence 2016)	**Low level visual learning** *Develop & deploy machine learning techniques for application to robotic vision*	**Semantic reps for robotic vision** *Scene understanding, knowledge representation & Visual Question Answering (VQA) applied to robotics*	**Robots & Objects** *Vision control of robots for manipulating real objects, pushing the boundaries on speed, coordination & complexity*	**Humans & Robots** *Deliberative planned motion that allows understanding of human activities & intent, to create effective safe anticipatory co-workers*	**Robust Vision** *Developing sensing technologies & algorithms, which allow robots to see in all viewing conditions*	**Algorithms & Architecture** *Create algorithms & techniques to allow vision to be run in real-time on robotic systems deployed in large-scale real-world applications, using distributed sensing & computational resources*
2015 Projects	**SV1** Learning for robotic vision **SV2** Globally persistent visual recognition **SV3** Semantic representations for robotic vision		**VA1** Image-based visual servo control **VA3** Vision for human-robot Interaction		**RV1** Novel camera hardware **RV2** High- to low-level vision **RV3** Learning spatio-temporally robust visual representations of place	**AA1** Vision as a network service for robots in highly dynamic environments

Activity Plan for 2016

RESEARCH THEME: SEMANTIC VISION

Program: Visual learning (VL)

- Online acquisition of training data for deep nets
- Instance level/pixel level annotations for 10s of classes
- Small and fast deep nets for embedded systems

Program: Semantic representations (SR)

- Automatic object discovery system based on geometry; whole-scene labelling for RGBD SLAM; systems integration of semantic SLAM on a mobile robot; semantic segmentation of video, and video-rate segmentation

RESEARCH THEME: VISION & ACTION

Program: Robots & objects (RO)

- Demonstrate independent control of Baxter's two arms, one performing a task with a tool using visual information from a camera held by the other arm, which moves to maximise information gain
- Demonstrate visually guided robot motion using a light-field camera
- Develop a robotic system that can show learnt hand-eye behaviours (trained in a simulator and then applied to a robot)
- Successfully compete in the Amazon picking challenge

Program: Humans & robots (HR)

- Demonstrate a human-robot cooperative task

RESEARCH THEME: ROBUST VISION

Program: Robust Vision (RV)

- Release the world's first open-source light-field camera (3D printer and laser cutter CAD files, and software)
- Demonstrate low-cost camera system that creates a 3D hyperspectral image from a moving robotic platform

RESEARCH THEME: ALGORITHMS & ARCHITECTURE

Program: Algorithms & Architecture (AA)

- Demonstrate a distributed robot localisation system using on-board and CCTV cameras as first iteration of VOS
- Demonstrate a flexible front end for key point SLAM that can detect and represent semantic entities
- Integrate the above with a powerful sparse matrix SLAM back end
- Demonstrate high speed dense algorithms for curvature and ICP from RGBD sensors

INDUSTRY ENGAGEMENT

- 2 x End-User Advisory Board meetings
- 4 x Industry showcases
- Launch of Centre's Industry Affiliates program

COMMUNICATION

- 4 x media releases by partners related to robotic vision

- 4 x public lectures on robotic vision
- Host visits and tours featuring robotic vision to government, industry and community

OUTREACH

- Launch PD for teachers in robotics and coding in Qld, ACT, Vic, SA
- Present robotic vision displays during National Science Week
- Supply Nao robots for demonstrations
- Host visits and tours of newly refurbished CoE spaces at each node

NETWORKS

- Create a Robotic Vision resources hub
- Visit international partners and host visits
- Research Training
- Project Management training offered at all nodes
- Research Leadership training offered to ECRs

GOVERNANCE

- 1 x Host a visit from the Australian Research Council
- 2 x Centre Advisory Committee Meetings
- Centre Executive meetings at least fortnightly with quarterly F2F meetings
- Monthly Research Committee Meetings
- 2 x Chief Investigator Meetings
- 1 x Mock Centre Review

Research Performance (KPIs)

RESEARCH FINDINGS

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015
Number of research outputs	Annually				
Conference Publications		25	42	50	53
Journal Publications		12	15	25	22
Disclosures/Patents		0	0	0	0
Quality of research outputs % of papers published in peer-reviewed outlets		100	100	100	100
Quality of research outputs Number of paper prizes or awards each year		1	6	3	1
Number of invited talks/papers/ keynote lectures given at major international meetings (incl. int'l conferences held in Australia)		5	6	10	41
Number and nature of commentaries about the Centre's achievements (in professional magazines and the popular press - newspapers, TV etc)		5	49	10	94
Citation data for publications (list web hits for online articles separately) Measured using Google Scholar	At review	na	na	na	na

Recognising our People

The excellence of our people has been recognised this year with a number of our researchers receiving prestigious awards or gaining appointments or promotions during 2015.

A notable early success for the Centre was the workshop paper "Sequence Searching With Deep-Learnt Depth for Condition-and Viewpoint-Invariant Route-Based Place Recognition", which was a collaboration between QUT and Adelaide researchers, and won best paper at the CVPR Workshop on Intelligent Vehicle Technology. Authors included Michael Milford, Ian Reid, Stephanie Lowry, Guosheng Lin, Niko Sünderhauf, Fayao Liu, Sareh Shirazi, Chunhua Shen, Edward Pepperell, Cesar Lerma, and Ben Upcroft.

Centre DIrector Peter Corke won first place in the Engineering and IT category of the Wharton-QS Stars Reimagine Education Award 2015 for Reimagining Robotics Education for driving robotics education accessibility, his innovative teaching, curriculum, resources and global leadership in the field, and second place in the Teaching Delivery category, which recognised the best approach resulting in outstanding student satisfaction and/or enhanced learning outcome. Peter was also awarded a QUT Vice Chancellor's Performance Award in 2015. Peter also presented two invited scientific talks in the UK at the Royal Society (Robots and Autonomous Systems) and the first annual conference of the Edinburgh Centre for Robotics.

Michael Milford was named Queensland's Young Tall Poppy of the Year at an awards ceremony hosted by the Queensland Minister for Science, The Honourable Leeanne Enoch MP and awarded prize money for use in support of the awardee's research and/or directly related science promotion activities. Michael was also awarded a QUT Vice-Chancellor's Award for Excellence, which recognises exceptional performance of staff who demonstrate sustained and outstanding achievement in activities that are aligned to the University's vision and strategic goals. Michael gave a plenary presentation at the IEEE International Conference on Robotics and Biomimetics (ROBIO) in China.

Although it will not commence until 2016, Centre researchers were awarded an ARC LIEF Grant for ~$400k to support computational infrastructure for developing deep machine learning models (Ian Reid, Svetha Venkatesh, Peter Corke, Mohammed Bennamoun, Stephen Gould, Anton van den Hengel, Chunhua Shen, Anthony Dick, Gustavo Carneiro, Dinh Phung, Niko Sünderhauf, Ajmal Mian).

Robotic Vision researchers from the University of Adelaide rated No.1 in the Pascal visual object classes (VOC) Challenge. The VOC challenge compares methods that are trained using provided "trainval" (training + validation) data, with the objective to produce the most accurate results.The Adelaide team, comprising Guosheng Lin, Chunhua Shen, Ian Reid, Anton van den Hengel, has developed innovative Deep Learning techniques to outperform teams from Microsoft Research, Oxford and the University of California.

Centre AI Tat-Jun Chin received a Best Paper Honorable Mention Award from Chang Yuan from sponsor SenseTime at the conference Computer Vision and Pattern Recognition CVPR 2015 in Boston, USA. The award was for a paper, "Efficient Globally Optimal Consensus Maximisation with Tree Search". Tat-Jun accepted the award on behalf of fellow co-authors, Pulak Purkait and Centre AIs Anders Eriksson and David Suter. CVPR is the top computer science conference, and 55th of all publications in all fields according to Google's h5-index, and 8th in all computer science and engineering.

There have also been promotions amongst Centre researchers including to Associate Professor, Centre CIs Stephen Gould (ANU) , Gustavo Carneiro (Adelaide), and Michael Milford (QUT).

CVC Best paper award to Centre researchers at CVPR conference workshop

Michael Milford (3rd from left) and his Tall Poppy Award with (from left) Queensland Minister for Science and Innovation, The Hon Leeanne Enoch, Professor David Lovell (QUT) and Queensland Chief Scientist, Geoff Garrett.

Semantic Vision (SV) Theme

will produce novel learning algorithms that can both detect and recognise a large, and potentially ever increasing, number of object classes from robotically acquired images, with increasing reliability over time

Ian Reid is Deputy Director of the Australian Centre for Robotic Vision, leads the Semantic Vision Theme, and directs the Adelaide node of the Centre.

Ian is an Australian Laureate Fellow and Professor of Computer Science at the University of Adelaide, where he has been since September 2012. Prior to that he was a Professor of Engineering Science at the University of Oxford, a post held in association with Exeter College where he was the senior Engineering tutor.

He received a BSc in Computer Science and Mathematics with first class honours from University of Western Australia in 1987 and was awarded a Rhodes Scholarship in 1988 to study at the University of Oxford, where he obtained a D.Phil. in 1992. Between then and 2000 when Ian was appointed to a Lecturership he held various Research Fellowship posts including an EPSRC Advanced Research Fellowship.

Theme Leader
IAN REID

His research interests range across computer vision, and are currently focused on life-long visual learning, and developing high-level representations for image and video understanding, especially those that can be computed and queried sufficiently rapidly to enable real-time robotic decision making and control. He has previous published widely in areas such as active vision, visual SLAM, visual geometry, human motion capture and intelligent visual surveillance. He has published 160 papers on these topics in major journals and refereed conferences, has more than 11500 citations and an h-index of 52. His papers have won prizes in BMVC '05, '09, '10, CVPR '08 and 3DV '14.

Ian serves on the program committees of various national and international conferences, on the editorial board of Image and Vision Computing and IEEE T-PAMI, and has led a number of EU, UK and Australian Research Council sponsored research projects.

The Semantic Vision theme emphasizes the role that a semantic description of a scene (higher level understanding) will play in enabling robots to understand, respond to, and interact with a complex unstructured environment that includes humans. Current state-of-the-art robot systems rely on computing the geometry of the scene, and are rarely designed for operation when that geometry can change, so this will provide a step-change in the possibilities for robotics.

The SV theme places machine learning at the heart of the robot vision system to enable the system to acquire and develop the requisite understanding, to detect and respond to change and to adapt its knowledge over time. The research in this theme is divided into three projects in which we: develop foundational learning techniques for robotic vision (SV1); learn the most appropriate image features for recognition and control (SV1, SV2); learn to recognise objects and object classes reliably and in spite of change (SV2); and integrate learning with representations that enable continuous learning and semantic scene understanding, integrating geometry and semantic concepts (SV3).

A significant development since the inception of the Centre has been the rise of deep learning, in which deep Neural Networks have, for the first time, been shown to outperform more traditional forms of learning by a significant margin in a wide variety of tasks of interest to the Centre. In SV1 and SV2, we have invested considerable effort in developing techniques for scene segmentation from images (i.e. labelling each pixel of an image with an object type), image patch matching, matching disparate image modalities (e.g. sketch and photograph), and new techniques for generic feature learning in deep networks. A number of these methods are currently topping the international performance benchmarks, and publications have been forthcoming in T-PAMI, NIPS, CVPR and ICCV. This work has also informed much of the progress in the Robust Vision theme's project RV3 on place recognition. The impact of deep learning, and the expertise within the Centre in this area, is also impacting on our approach to our Vision and Action theme, as it now seems possible to link learning and action more directly through deep Recurrent Neural Networks.

SV3 is making progress to consolidate these ideas within the context of mapping systems: building systems for labelling of 3D scenes acquired using continuous observation of a scene by an RGB or RGB-D camera; creating the state-of-the-art system for depth prediction (i.e. scene geometry estimation) from a single view; exploring non-standard architectures for deep networks, in particular making use of multi-model auto-encoders for scene reconstruction and semantic labelling.

Straddling the boundary between SV2 and SV3 is work on action recognition and activity recognition of people, an essential component of any system that seeks to model a dynamic world including humans.

As with more traditional object and feature recognition work, our findings have been that a deep learning approach to this outperforms more traditional methods, but we have made interesting theoretical and practical advances that bridge the gap.

Computer vision scene understanding with labels to form a semantic map. The colours encode the semantic categories of different places encountered in an office environment (orange) with a kitchenette (dark green) and a long corridor (light green). Source: Sünderhauf, N. et al. (2016) Place Categorization and Semantic Mapping on a Mobile Robot, ICRA (to be published).

SV1 - Learning for robotic vision

developing foundational learning techniques for robotic vision

Project Leader
CHUNHUA SHEN

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 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 about six years. From 2006 to 2011, he held an adjunct position at College of Engineering & Computer Science, Australian National University. He moved back to Adelaide in 2011.

RESEARCH TEAM

Ian Reid (CI), Stephen Gould (CI), Richard Hartley (CI), Gustavo Carneiro (CI), Anton van den Hengel (CI), Vijay Kumar (RF), Guosheng Lin (RF), Anoop Cherian (RF), Basura Fernando (RF), Bohan Zhuang (PhD), Zhibin Liao (PhD), Saroj Weerasekera (PhD), Peter Anderson (PhD), Jian (Edison) Guo (PhD)

The goal of SV1 is to deliver the fundamental developments in machine learning required to support the advances in robotic vision required to create robots that see. Over the last decade, machine learning has underpinned much of the success in computer vision but the specific needs of, and opportunities within, robotic vision have been neglected. Towards this end, we are interested in developing: learning methods that support life-long learning, so that robotic systems can adapt and continually improve over their lifetime; efficient inference methods, since these are critical for the online operation of a robotic system; and light-weight learning models that are deployable to robotics, embedded systems.

As Deep Learning (DL) is changing the landscape of computer vision and many other AI related fields, the SV1 project team have been spending considerable effort on DL. One particular issue that we are working on is to learn compact models and at the same time, make the evaluation of DL significantly faster. Models learned by DL usually contain hundreds of millions of parameters and can consume a few gigabyte of RAM during inference. One has to compress the model and make the inference an order of magnitude faster before DL can be practically applied to many robotics applications, which typically have memory and computation power constraints.

The second main topic is Dr. Guosheng Lin's work on semantic segmentation. Dr. Lin and colleagues have developed DL based semantic segmentation systems which have topped the VOC image segmentation competition as well as on other public benchmark datasets (results published at NIPS 2015 and CVPR 2016). As the research focus of the computer vision community is shifting from image-level prediction to dense per-pixel prediction, semantic segmentation is attracting more and more research interest. From a practical point of view, semantic segmentation associates one of the pre-defined class labels to each pixel, thus provides much richer information for other high-level vision tasks. Beyond semantic segmentation, we will start working on instance-level segmentation shortly. Outcomes from this work feed into our Robust Vision themes RV2 and RV3 projects, supporting the robustness of those projects for contextual priming

and place recognition. The work also feeds into SV2 for general object recognition and SV3 for general scene classification, decomposition and understanding.

As computer vision is mainly about making sense of imagery data by finding the high-level information, which can be the label, depth, surface normal of each pixel; and object-level information. These tasks are typically achieved by building models like conditional random fields (CRFs) in order to exploit how the variables depend on each other. CRFs allow modelling of contextual relationships, correlations that exist between tasks such as coarse and fine-grained recognition, and correlations between detections and object pose. In SV1, we are designing innovative, efficient learning methods that enable us to train deep structured models so as to take advantage of the powerful feature learning capability of deep learning and, at the same time, the rich modelling capacity of structured learning for representing the complex relationships in images. We have successfully demonstrated the usefulness of our framework on the tasks of learning depth from monocular images; and semantic pixel labelling. We are exploring the possibility of solving image processing problems such as image denoising, deblur, and other image enhancement tasks with the developed deep structured models.

We are developing foundational machine learning techniques that are suited to the constraints of robotic vision; that are adaptive and flexible; that generalise to new environments with a relatively small amount of training data but that continue to learn over their lifetime. Our work underpins the Centre's other research projects.
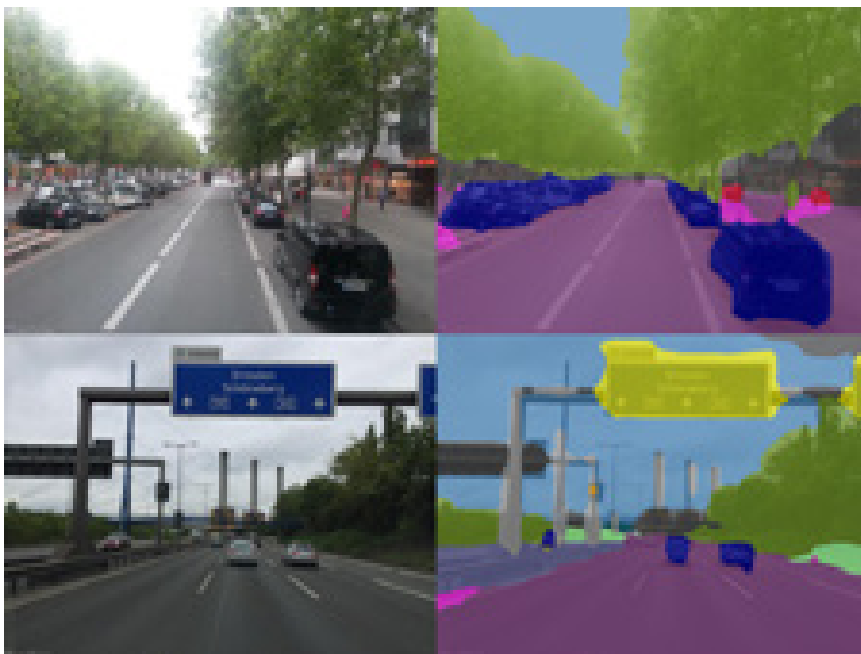
The images on the left are input and the images on the right are the output of our deep learning based semantic pixel labelling method. Different colours encode 20 categories. Our method represents the state-of-the-art for semantic understanding of urban street scenes. Source: Lin,G., Shen, C., van dan Hengel, A. & Reid, I. (2016) Efficient piecewise training of deep structured models for semantic segmentation. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'16).

SV2 - Globally persistent visual recognition

learning to recognise objects and object classes reliably and in spite of change

Project Leader
GUSTAVO CARNEIRO

Gustavo Carneiro is an Associate Professor of the School of Computer Science at the University of Adelaide. He joined the University of Adelaide as a senior lecturer in 2011, and has become an associate professor in 2015. His main research interests are in the fields of computer vision, medical image analysis and machine learning. In particular, in medical image analysis, Gustavo is developing multimodal methods to analyse medical images using deep learning techniques, and in computer vision his focus is in the development of new training procedures for deep learning methods and the development of feature learning approaches.

From 2008 to 2011 Gustavo was a Marie Curie IIF fellow and a visiting assistant professor at the Technical University of Lisbon (Instituto Superior Tecnico) within the Carnegie Mellon University-Portugal program (CMU-Portugal). From 2006 to 2008, Gustavo was a research scientist of the Integrated Data Systems Department at Siemens Corporate Research in Princeton, USA.

RESEARCH TEAM:

Chunhua Shen (CI), Gordon Wyeth (CI), Gustavo Carneiro (CI), Guosheng Lin (RF), Vijay Kumar (RF), Trung Tanh Pham (RF), Cesar Cadena (RF), Yuchao Jiang (PhD), Bohan Zhuang (PhD), Hui Li (PhD), Mehdi Hosseinzadeh (PhD), Zhibin Liao (PhD)

The main objective for SV2 is about the detection and recognition of large numbers of object classes. This outcome is fundamentally required for a robust, flexible robotic vision system yet remains an outstanding problem with no off-the-shelf solution. The problem is challenging since the appearance of an object is a function of both the object itself and extrinsic factors such as relative pose and lighting. Currently, object detection and recognition is typically considered a "retrieval" task over a database of images – there is an underlying assumption that the photographer has already framed salient aspects of the scene. However, for robotic vision the background and context of the images can vary dramatically, lighting is often uneven, object pose is highly variable and the object may be significantly occluded. An effective system must perform rapid (real-time) classification and inference, autonomously and with high reliability.

We are focusing on methods that are robust to context (places, appearances, pose, etc.) and using active training methods to discover new objects. The fundamental questions we are answering include: how to automatically select the best low-level features to use for detection and recognition; how to encode contextual knowledge and address the likelihood of objects of interest occluding one another; what are the most effective classifiers for multi-class classification and how can they be trained and inference be performed efficiently; and how to find and take advantage of natural hierarchies that exist amongst object classes?

To select low-level features, we are currently looking at: unsupervised metric learning of image patches to build pre-trained models that can be adapted to other domains and datasets; supervised metric learning that can be adapted to other datasets using active learning; and cross-modality distance metric learning. Our target is to develop methods of unsupervised training

The figure shows a novel visual SLAM system which is capable of detecting duplicate objects in the scene during run-time without requiring an offline training stage. The top view shows the scene, whereas the bottom view shows the reconstructed map. Background landmarks are shown as black points, while landmarks belonging to duplicate instances of an object are shown in the same colour. From a system, "Multi Object SLAM with Run-Time Object Discovery through Duplicates", developed by Thanuja Dharmasiri and Vincent Liu.
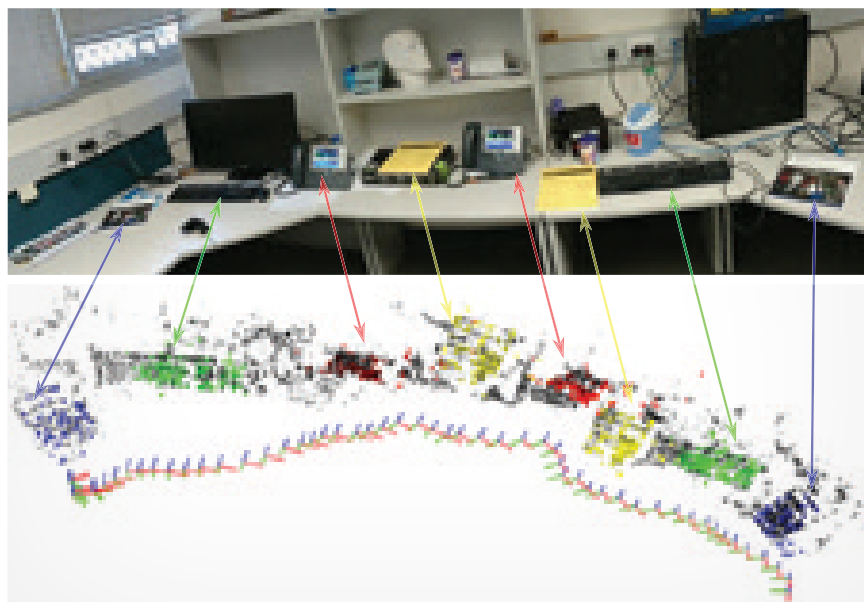
that can be effective at transfer learning and domain adaptation. In addition, we also aim at producing methodologies that are effective at discovering new objects in a dataset. To pinpoint the types of classification models and optimisation methods to use, we are currently exploring: new types of deep learning architectures and optimisation functions; and multimodal training approaches. Our goal is to produce training methodologies that are competitive with the state of the art and, at the same time, cheap to train in terms of computational resources. To encode contextual knowledge we are using: methodologies to estimate the likelihood of objects appearing together in a scene; affordances and

context-based navigation; and text/face recognition in the wild. We aim to produce a methodology that can be deployed in a robot that can figure out the context in order to be more effective at identifying the visual objects present in a scene. Finally, to encode object hierarchy, we are applying structured output learning using object hierarchy for semantic visual object segmentation.

We are currently planning the set up of a dataset for construction and agriculture that will serve as testbed for all methodologies developed for our project. A very large dataset of images and videos will be captured on construction and agricultural sites and we will then work on the annotation of that dataset, where the

images and videos will have visual objects of interest delineated and labeled. In parallel, we will define the experimental set up for classification, detection, and segmentation tasks and zero/one-shot learning. Our project will demonstrate a robot vision system that can learn to recognise objects for specific tasks and later to generalise from other task domains, which will reduce retraining time. We will develop a technology demonstrator for Construction and Agriculture, which integrates object detection and recognition with high-level semantic feedback and lifelong improvement.

SV3 - Semantic representations for robotic vision

integrating learning with representations that enable continuous learning and semantic scene understanding, integrating geometry and semantic concepts

Project Leader
IAN REID

RESEARCH TEAM

Stephen Gould (CI), Anton van den Hengel (CI), Gustavo Carneiro (CI), Ben Upcroft (CI), Gordon Wyeth (CI), Tat-jun Chin (AI), Trung Thanh Pham (RF), Yasir Latif (RF), Cesar Cadena (RF), Markus Eich (RF), Basura Fernando (RF), Anoop Cherian (RF), Cai Ming (PhD), Rodrigo Santa Cruz (PhD), Saroj Weerasekera (PhD), Mehdi Hosseinzadeh (PhD), Josh Boys (PhD)

In SV3 we are developing algorithms and the representations required to represent potentially dynamic environments in terms of semantically meaningful entities, properties and relationships. The key question we are addressing is how best to combine geometric information with semantic information. A significant capability that the project is aiming to deliver is a mapping system that is real-time and can be used as a building block for other work, especially in our Vision and Action projects.

Some of SV3's work has overlapped with SV2, with a concentration on systems for scene labelling either in outdoor scenes (for navigation) or for 3D indoor scenes. We have developed

methods for semantic labelling of 3D scenes such as those acquired by depth cameras or dense mono-SLAM systems. We have also begun work on incorporating deep learning methods into SLAM by labelling maps. Most notably we have developed a method for simultaneous estimation of semantic labels and scene depth (from a single view) using multi-modal auto-encoders. In the latter part of the year the project also embraced ideas in understanding human activities, a key component of semantic decomposition of dynamic scenes.

The team has acquired considerable experience with point-based (ORB-SLAM), semi-dense (LSD-SLAM) and RGB-D slam (Infinitam), including developing the first version of LSD-SLAM that works with rolling-shutter cameras. Work has also been

conducted on recognising tools from drawings (instruction manuals) and locating the tools in real imagery, with a view to being able to design plans for assemblies. We are also looking at how to semantically represent indoor scenes using geometric primitives such as planes, boxes, spheres, etc.

Our key milestone is to develop a system for semantic mapping in realtime from a handheld (or robotically controlled) monocular camera. This involves key questions of representation and the role of the outputs of SV1 and SV2, as well as the cross-over to our Algorithms and Architecture theme (distributed mapping).

The departure of Research Fellow, Dr Cesar Cadena to ETHZ at the end of the year has created a temporary staff gap.
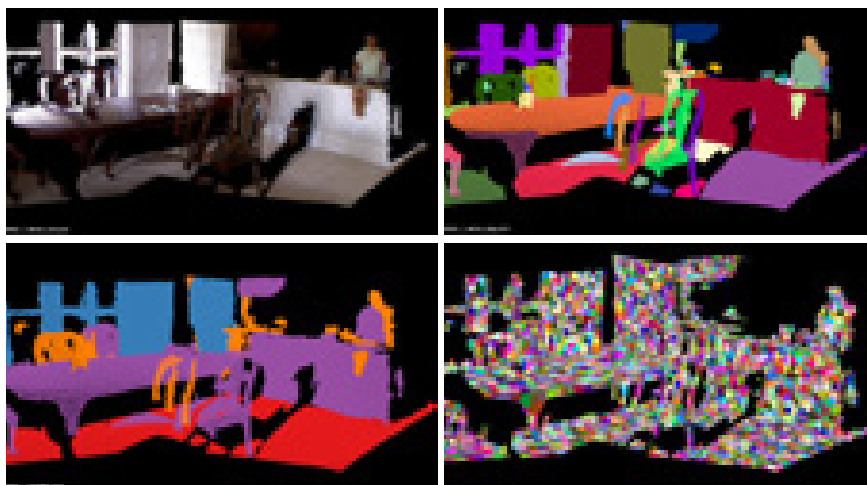
Two-stage approach to 3D semantic segmentation by researchers Trung Pham, Markus Eich, Ian Reid and Gordon Wyeth. Starting from a 3D point cloud (top let) reconstructed using InfiniTAM, we over-segment the point cloud into supervoxels (top right), then partition these supervoxels into connected components (bottom right), which represent individual 3D objects. Finally we used a simple Randomized Decision Forest classifier to predict semantic labels for these "components". We classify objects into 4 common semantic categories: Ground(Red), Structure (Blue), Furniture (Purple) and Props (Orange) (unpublished)

Vision & Action (VA) Theme

will create new theory and methods for using image data for control of robotic systems that navigate through space, grasp objects, interact with humans and use motion to assist in seeing

Rob Mahony obtained a science degree majoring in applied mathematics and geology from the Australian National University in 1989. After working for a year as a geophysicist processing marine seismic data, Rob returned to study at ANU and obtained a PhD in systems engineering in 1994.

Between 1994 and 1997 he worked as a Research Fellow in the Cooperative Research Centre for Robust and Adaptive Systems based in the Research School of Information Sciences and Engineering in ANU. From 1997 to 1999 he held a post as a post-doctoral fellow in the CNRS laboratory for Heuristics Diagnostics and complex systems (Heudiasyc), Compiegne University of Technology, France. Between 1999 and 2001 he held a Logan Fellowship in the Department of Engineering and Computer Science at Monash University, Melbourne, Australia.

Theme Leader
ROB MAHONY

Since July 2001 Rob has been working in the Department of Engineering, ANU. His research interests are in non-linear control theory with applications in robotics, mechanical systems and motion systems, mathematical systems theory and geometric optimisation techniques with applications in linear algebra, computer vision, digital signal processing and machine learning.

The relationship between motion in the world and changes in images are at the heart of robotic vision. Vision is tightly coupled to action and provides rapid and continuous feedback for control. Action also directs the eye and is required to inform the task. Motion in an image reveals much more about the world, and how the robot is moving within the world, than the content of the image alone. Knowledge of how objects move enables prediction of their future location and where visual attention should be devoted.

Vision-based control has been extensively developed in the robotics field over the last twenty years. Image-Based Visual Servo (IBVS) controls robot motion implicitly by describing how objects should move in one or more images of the world, rather than in the 3D world itself – this greatly simplifies motion planning tasks and has significant advantages in robustness to camera and target calibration errors, reduced algorithmic complexity, and is easily extended for multiple cameras and camera types.

Classical vision-based control depends on tracking ideal image features such as points, line segments or arcs of circles rather than realistic objects.

We are taking a different approach, investigating control using dense measures such as mutual information and image-structure such as homographies. Early visual servo control was deployed on industrial type robotic arms but this has proved difficult to generalise to real-world dynamic systems. Key challenges include dealing with dynamic systems

that are underactuated (e.g., flying vehicles, cars), and control, which is a function of object distance that is often unknown. For the case of multiple cameras we require a modular way to create the control system, for example, using special mathematical and navigation functions.

The Vision and Action theme is developing modular vision-based control algorithms designed to exploit networked vision services; robust vision sensors; and learned features and semantic visual information representations developed in our Centre's other themes. The image-based visual controllers we develop (VA1) will be applied for arbitrary robot dynamics and multiple networked cameras from our AA1 project; use improved image features from SV1 and scene understanding from SV2 for more intelligent control (VA2 - to commence in 2018); and extend these techniques to inferring and responding to human intent in cooperative work situations (VA3).

VA1 Image-based visual servo control

developing image-based visual controllers for application to arbitrary robot dynamics and multiple networked cameras
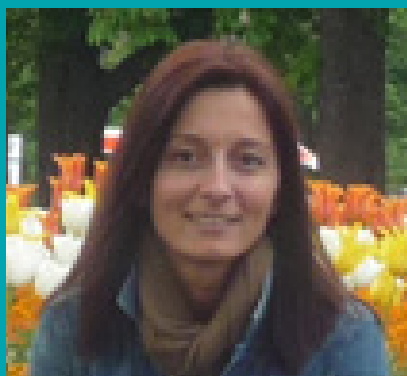
Project Leader
VIORELA ILA

Viorela Ila is a research fellow with the ARC Centre of Excellence for Robotic Vision at the Australian National University (ANU). Her research interests span from low-level image processing, parallel architectures, robot vision to advanced techniques for simultaneous localisation and mapping (SLAM) and 3D reconstruction based on cutting-edge computational tools such as graphical models, modern optimisation methods and information theory.

Viorela received an Engineering degree in Industrial Engineering and Automation from the Technical University of Cluj-Napoca, Romania, in 2000 and a PhD degree in Information Technologies from the University of Girona, Spain, in 2005. After her PhD studies, she joined the Robotics group at the Institut de Robòtica 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.

Our Baxter robot demonstrates Vision and Action by cutting the official ribbon at our Centre's launch, watched by then Minister for Education and Training, The Hon. Christopher Pyne MP.

RESEARCH TEAM

Peter Corke (CI), Tom Drummond (CI), Rob Mahony (CI), Ben Upcroft (CI), Juxi Leitner (RF), Niko Sünderhauf (RF), Sareh Shirazi (RF), Markus Eich (RF), Fangyi Zhang (PhD), Juan Adarve (PhD), Adam Tow (PhD)

Our project is creating the theoretical framework that allows active vision-based control for a variety of robotic platforms and applications. Perception is now directly integrated in the control loop. Our research lies on the intersection between computer vision, robotics and control theory. The challenge is how to adapt the perception and control strategies to the unstructured, dynamic environments that currently challenge modern robots.

Perception has to be performed in such a way as to produce useful

commands to direct the robot towards specific goals. The online nature of the applications restricts the time dedicated to image processing, therefore real-time algorithms are the key to obtaining on-time feedback for the control system. High-level understanding of the environment replaces low-level perception and the problem is formulated in the task space rather than image space. VA1 strongly benefits from image understanding techniques developed in our Robust Vision theme and real-time localisation and sensing developed in our Algorithms and Architecture theme.

Fast and accurate solutions exist for processing data from sensors in a static way. Nevertheless, when humans try to understand their environment they look around, get

closer to the objects of interest, change their position to see better and accumulate as much information as possible about the surroundings. Robots should be able to act in a similar way. Active perception is a sensing system that can change the viewpoint of the robot to enable it to explore the environment and get better information to perform given tasks.

A number of areas of research are being pursued, some of which involve the development of technology demonstrators. These include: active environment reconstruction with quadrotors; visual servoing using the bi-manual Baxter robot; learning of robotic control for manipulation; high speed vision algorithms for robotic applications; kineto-plenoptic models; and temporally scalable incremental SLAM.

Harvey the Capsicum-picking robot

As part of the Queensland Department of Agriculture and Fisheries (DAF) three-year Strategic Investment in Farm Robotics (SIFR), capsicum (bell or sweet peppers) was identified as an important Queensland crop that could benefit from robotic harvesting. The project team, led by Centre AI, Professor Tristan Perez, undertook a literature review on robotic capsicum harvesting and the challenges associated with this particular crop. The major challenges were associated with image processing (green fruit on green background and heavy occlusion), and manipulation due to the unstructured environment of the crop. The SIFR team developed an algorithm to detect in-field capsicum that improves on state-of-the-art vision systems. By using more effective visual texture features than previously considered in the literature, our algorithm is able to detect approximately 70% of in-field capsicum. We tested our algorithms using both day and night imagery, and as expected best performance is obtained at night due to better control of illumination. We also assessed the performance of our algorithm against capsicum detection of humans and found that the results are comparable. The horticulture industry in Australia has a gross value of more than $8 billion dollars per annum.

Harvey, a prototype robotic platform for harvesting capsicums, being tested at a Queensland Government protected cropping facility in November 2015

The SIFR team then developed a new agricultural robot prototype – nicknamed 'Harvey' – designed to harvest capsicums. In November 2015, the team conducted the first trials of the robot at a Queensland Government protected cropping facility in North Queensland (pictured). Tasked with identifying and picking red capsicums, Harvey performed significantly better than any capsicum-harvesting robot ever has.

Despite significant efforts by the worldwide horticultural research industry, progress in creating robots to harvest capsicums has been modest to date. In late 2014, worldwide literature indicated a success rate of only 6% in testing scenarios similar to those used for Harvey, and up to 30% when the crop is modified and leaves are removed.

Harvey achieved a harvesting success rate of 65% with unmodified crops (that is, with no leaves removed or fruit moved before harvesting). In 2016, the research team will fine-tune Harvey's performance and conduct further trials, and they believe only minor modifications will be required to achieve an overall success rate exceeding 90%.

How does it work? Harvey's robotic arm has a camera and a unique cutting tool attached to it. Using data from the camera, the robot creates a 3D model of each fruit and its surroundings and plans 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 crop-manipulation tools enables successful harvesting of the crop and promises significant benefits for horticulture growers, who export more than $2b in products every year.

We are now seeking partners both in Australia and overseas to commercialise this technology. In the future, our researchers also plan to investigate how automated harvesting technologies can be used for other crops, such as mangoes, strawberries and avocados.

VA3 Vision for human-robot interaction

inferring and responding to human intent in cooperative work situations

Project Leader
ROBERT MAHONY

RESEARCH TEAM

Tom Drummond (CI), Anton van den Hengel (CI), Gordon Wyeth (CI), Markus Eich (RF), Sareh Shirazi (RF), Niko Sunderhauf (RF), Anoop Cherian (RF), Fahimeh Rezazadegan (PhD), Ahmed Abbas (PhD), Ashesh Vasalya (PhD), Sean McMahon (PhD)

VA3 is developing computer vision tools to allow future robots to work intimately and collaboratively with a range of human users. For such robots to be accepted as useful members of a workforce they must be capable of recognising people, anticipating and reacting to the intent of humans, and recalling and using recent interaction history. Many of the robotic vision techniques required for interaction with humans will find

their origins in the other projects of the Centre, but will be specifically tailored to challenges in human robot interaction in this project. Humans are a special case for robotic vision, critical to the development of application areas where robots are to work safely and effectively alongside humans.

Our project focusses on natural interaction with human users based on visual interpretation of human pose and gesture. The core innovation is the development of mathematical models of human-robot interaction that build on the modular control systems construction techniques developed in VA1. The new mathematical techniques will allow models of human presence, motion, gesture and intent to be embedded into the robot's control system to guarantee human safety in all contexts. In a manufacturing context the robot could help to lift heavy loads guided by the natural motion and gestures of an operator.

A robot with semantic understanding (in SV3) can be directed to act on explicit categories or instances of an object by simple low bandwidth (possibly spoken) commands such as "move the pallet to the corner"; similarly the robot can report on the presence of objects at a location

based on its semantic understanding of a scene. Human robot interaction centred around symbolic description of scenes and tasks represents a new and exciting frontier in human-robot interaction research. The project will draw on SV3 to build tools to enable robots to perceive and understand relevant human behaviour and needs. The challenges are to develop techniques for learning subjective models of relevant human beliefs, desires and intentions based on vision and interaction, and to integrate the information provided by the semantic vision algorithms into a motion strategy that both deals with the uncertainty in the data and provides the human with clear indicators of the robot's intention. This problem will be addressed by adapting techniques from VA2 to incorporate the high level of uncertainty inherent in human intent.

The outcome of this project will be a suite of tools that enable natural interaction between humans and robots based on presence, gesture, commands or perceived intent. The tools will be presented in the technology demonstrators for Manufacturing and Construction, which will showcase robots and humans working safely side-by-side.

Robust Vision (RV) Theme

will develop new sensing technologies and robust algorithms that allow robots to use visual perception in all viewing conditions: night and day, rain or shine, summer or winter, fast moving or static

Project Leader
PETER CORKE

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.

He is a fellow of the IEEE, former editor-in-chief of the IEEE Robotics & Automation magazine, founding and associate editor of the Journal of Field Robotics, founding multi-media editor and editorial board member of the International Journal of Robotics Research, member of the editorial advisory board of the Springer Tracts on Advanced Robotics series, recipient of the Qantas/Rolls-Royce and Australian Engineering Excellence awards, and has held visiting positions at Oxford, University of Illinois, Carnegie-Mellon University and the University of Pennsylvania.

He received his undergraduate and Masters degrees in electrical engineering and PhD from the University of Melbourne.

For robots, the process of vision starts with images acquired by cameras, but images are a complex (non-linear and non-invertible) function of many variables including: the 3D structure of the scene and the materials within it, the lighting, and the viewpoint of the camera. Vision outdoors is confounded by changes in appearance from day to night and wet to dry, by shadows, and by changes due to season where objects (e.g., trees) dramatically change appearance. A robot's vision system must be able to understand its world in a robust way and we will tackle this problem in several ways.

Firstly, current cameras fall far short of human capability when it comes to sensitivity, dynamic range and the ability to focus. We are rigorously investigating new sensing technologies including light-field and night-vision cameras. We are also investigating how to intelligently combine multiple low-cost sensors, to improve dynamic range, colour acuity, frame rate, and field of view.

Secondly, we are developing robust visual processing algorithms that work despite distractors such as changes in camera viewpoint, motion blur and shadows. For example, our recent results show how colour information can eliminate shadows and that it may be possible to use hyper-low resolution images (tens of pixels) to significantly reduce the effect of image change due to blur, defocus, rain or viewpoint change. We are developing robust algorithms that can recognise places despite poor viewing conditions: night and day, rain or shine, summer or winter, fast moving or static.

Finally, robots must infer the state of the world from images despite the limitations of the image formation process. Current vision systems can

not match human vision when it comes to using information from multiple eyes on a body moving deliberately through the world. Human vision is augmented by visual memory and experience, which allows us to maximise our understanding of our surroundings. We humans also bring a number of tricks to bear on the problem of vision while moving, including fusing information from two eyes, deliberate head and body motion, visual memory and a sense of visual context. To replicate human vision, robots can use multiple views to allow the reconstruction of 3D structure. This helps to resolve the ambiguities present in monocular (single eye) views, for example, to distinguish a picture of a door from a

real door, and to eliminate changes due to view point.

We are developing mechanisms to exploit visual context, to constrain the set of objects we might expect to encounter in a particular environment. We will also use contexts to emulate the human "high-to-low" level pathways — to choose or condition the low-level image processing stages (for example, feature detectors) based on knowledge of the scene, e.g. when tracking cars we might look for edges or shadows of cars in the daytime and for clusters of lights at night. Different cues can be combined, or actively selected based on feedback from image processing quality.

The research in the Robust Vision theme builds upon, and contributes to work in other themes. Semantic vision is key to learning visual contexts and the variable appearance of objects in the world. Vision and Action is key to exploiting motion to resolve ambiguity in the observed scene, basing our inference on more evidence collected from different viewpoints. Finally, Algorithms and Architectures provides the computing and communications platform on which this early stage image capture and processing is built.

A Lytro Illum light field camera mounted on an ABB industrial robot arm allows for precisely repeatable camera trajectories

RV1 - Novel camera hardware

improving the visual information available to robots

**Project Leader
HONGDONG LI**

Hongdong Li is a Chief Investigator for the Centre and a tenured Associate Professor at ANU where he teaches computer vision and robotics courses. He joined ANU in 2003 and was then seconded to NICTA until 2009. In 2010, Hongdong moved back to ANU.

Hongdong's main research interests include 3D computer vision, vision navigation, and mathematics optimisation. He made important contributions to 3D vision, incl. camera pose computation, robust geometry estimation, 3D registration, and non-rigid deformable shape structure-from-motion. He is an Area Chair in the area of "3D computer vision" in recent major international conferences in computer vision including ICCV, CVPR and ECCV. He was a winner of the IEEE CVPR 2012 Best Paper Award, jointly with his student and colleague, for their work on non-rigid 3D vision reconstruction. He was also the recipient for the IEEE ICIP 2014 Best Student Paper Award for work on multibody motion segmentation, a winner for IEEE ICPR 2010 Best Student Paper Award for multi-view human gait identification, a winner of DICTA'13 Canon Best Paper Award, and a winner for DSTO Best Fundamental Paper Prize in 2014. He was an organiser and program chair for ACRA in 2015 - the Australasian annual robotics conference.

His past research projects include a major autonomous car driving project, the Australia Bionic Eyes Project with BVA. His research has been funded by the Australia Research Council (ARC) Discovery Program, Linkage Program, NSFC China International collaboration program, Toshiba Corp, General Motors Global Research, and Microsoft Research (via a Linkage grant).

RESEARCH TEAM

Peter Corke (CI), Robert Mahony (CI), Richard Hartley (CI), Ben Upcroft (CI), Andy Davison (PI), Donald Dansereau (RF), Viorela Ila (RF), Chuong Nguyen (RF), Laurent Kneip (AI)

Modern image sensors have abundant resolution but other factors are just as important for robotic vision. For example, navigation requires wide fields of view, outdoor operation requires high dynamic range and shadow elimination, robust object recognition require greater colour acuity, and high-speed robot motion requires high frame rate. New camera technologies are constantly being developed, and commoditised. Lightfield and random pixel cameras are becoming generally available and have potential as depth sensors for robots. New mixed-pixel sensor arrays will improve dynamic range and colour accuracy, while single-pixel cameras promise low-cost hyper-spectral sensing. Operation at night is essential for vision-guided robots but our recent experience indicate that standard cameras on a moving robot at night require light sources that are on the edge of human eye safety limits. Very sensitive cameras have been developed for security, animal studies and military operations based on special sensors or image

intensifiers. Classical computer vision techniques that were once impractical from a speed or cost perspective are now feasible, for example, depth from focus using high-performance variable focus optics developed for mobile phones.

We are also keeping abreast of emerging camera technology and evaluating its utility for robotic vision. We have established a small laboratory to characterise the temporal and radiometric characteristics of current and emerging imaging devices. Deep knowledge of camera performance will be key know how for the Centre and will be incorporated into our technology demonstrators.
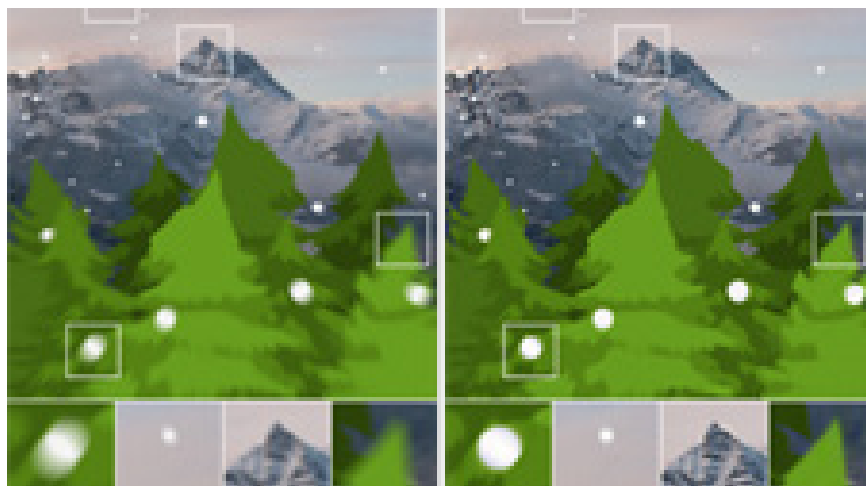
In this work Research Fellows Donald G. Dansereau, Jürgen Leitner and Anders Eriksson explored the simplification of deblurring using light field cameras. Motion blur depends on the shape of the scene, typically requiring complex algorithms that jointly estimate the 3D shape and remove blur. Their work has shown that in the case of a light field camera no shape estimate is required, rather blur can be removed by directly exploiting the depth information implicitly captured in every light field image.
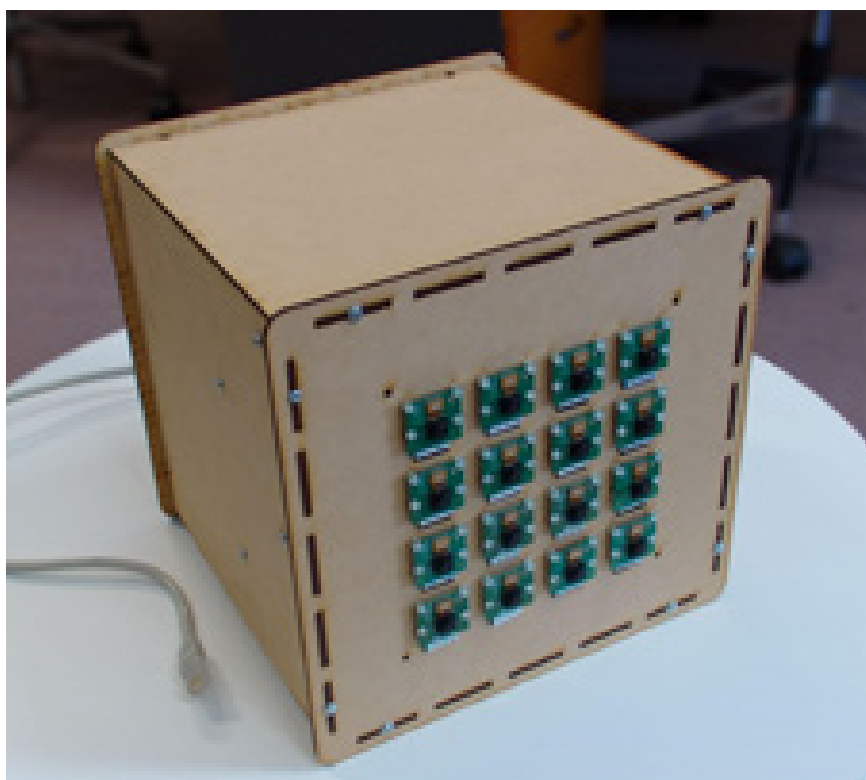
The 16-Pi Light Field Camera Array integrates 16 Raspberry Pi single-board computers, each equipped with a 5 MPix camera, to yield 80 Megaray light fields at full video frame rates. Originally constructed by Rafe Denham as part of a final-year undergraduate project.

RV2 - High- to low-level vision

giving robots the ability to choose the vision data they require to effectively operate

Project Leader
BEN UPCROFT

Ben Upcroft is an Associate Professor and the leader of QUT's Robotics and Autonomous Systems Discipline. He applies his expertise in long-term visual navigation in the Robust Vision theme, with connections to Algorithms and Architecture, and contributes to the development of real robotic systems.

Ben's interests lie in the development of vision systems for long term robotic applications ranging from underwater and ground to airborne autonomous platforms. The overarching theme of his work is in robust visual perception for robotics. Ben joined QUT in 2011 after previously working at the Australian Centre for Field Robotics (University of Sydney) and with CRCMining at The University of Queensland (UQ).

RESEARCH TEAM

Michael Milford (CI), Chunhua Shen (CI), Peter Corke (CI), Niko Sünderhauf (RF), Sareh Shirazi (RF), Juxi Leitner (RF), Cesar Cadena (RF), Guosheng Lin (RF), Adam Tow (PhD), Fahimeh Rezazadegan (PhD), Sean McMahon (PhD), Fangyi Zhang (PhD), Zetao Chen (PhD), James Sergeant (RA), Peter Kujala (RA)

If robotics is to be ubiquitous in our everyday lives, scene understanding and an understanding of the types of features that are required to facilitate the interpretation of vision data depend entirely on the task required of the robot. This task is not necessarily object recognition. For example, a self-driving car does not need to recognise a sail boat but does need to know the difference between a footpath and the road. A house-cleaning robot does not need to recognise a car but needs to

know that it can pick up a knife from the handle but not the blade. Thus affordances or the functionality of the scene is important to inform the low-level visual features required for a task as well as the abstract high-level semantics of a scene. Our project will enable robots to observe the world in a truly robust way. We are using semantic knowledge of the world to inform how the image is sensed and processed by a robot while also guiding how a robot understands the scene using low-level visual information. This bidirectional information flow between raw feature extraction and abstract representations of the scene enables feedback at all levels of the robot's visual processing pipeline. We achieve this by employing deep recurrent neural nets (RNNs) in which training at all visual levels of the RNN, from gradient-based features at the first layer to abstract classification

at the top layers, are informed by all other layers in the network. In this way we can develop a robotic system that can learn and adapt to real-world changes over a large set of challenging scenarios such as extreme lighting changes, low-light conditions and highly dynamic motion.

RV2 acts as an interface between the semantic vision (SV) and vision & action (VA) themes. It integrates and extends concepts from both themes and ties them together in two mobile robot demonstrators that highlight a fully integrated learning process. The first is learning to drive from scratch, building on our existing strengths of navigation in real world environments while also providing an application domain in which we can develop deep learning methods for a concrete task in the real world. The second demonstrator is a house-helper robot, that is, a robot you can take home

with you to perform chores without manual training. This demonstrator encompasses much of what we aim to achieve in this project, in which new tasks should be learnt online without manual intervention. Intelligent behaviour should be exhibited in new domains. The tasks carried out should inform the lower level behaviours and perceptual information.

Learning target reaching from exploration in simulation and skill transfer between simulation and the real world (from Centre PhD researcher, Fangyi Zhang)
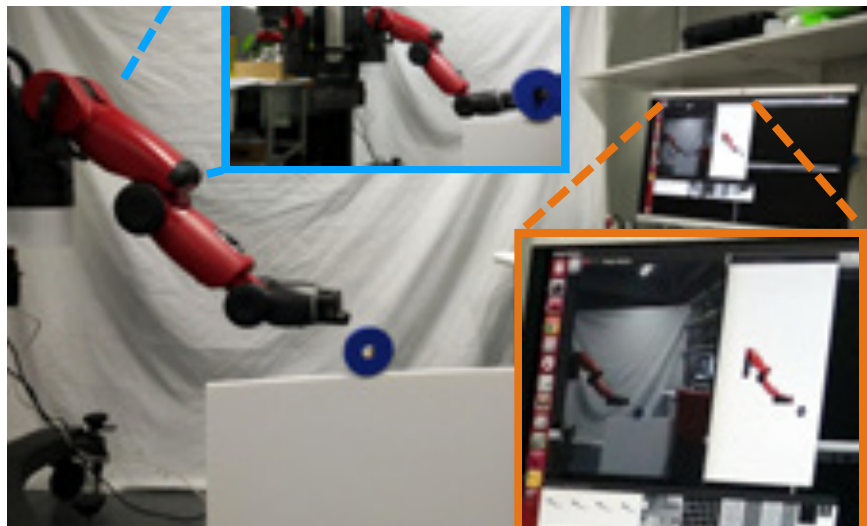
RV3 - Learning spatio-temporally robust visual representations of place

determining robot location using visual sensing alone

Project Leader
MICHAEL MILFORD

Michael Milford is an Associate Professor and Australian Research Council Future Fellow at QUT with a research focus, although he continues to teach Introduction to Robotics every year. After a brief postdoc in robotics at UQ, Michael worked for three years at the Queensland Brain Institute as a Research Fellow on the Thinking Systems Project. In 2010 he moved to the Queensland University of Technology (QUT) to finish off his Thinking Systems postdoc, and then was appointed as a Lecturer in 2011. In 2012 he was awarded an inaugural Australian Research Council Discovery Early Career Researcher Award, which provides him with a research-intensive fellowship salary and extra funding support for 3 years. In 2013 Michael became a Microsoft Faculty Fellow and lived in Boston on sabbatical working with Harvard and Boston University. Michael holds a PhD in Electrical Engineering and a Bachelor of Mechanical and Space Engineering from the University of Queensland (UQ), awarded in 2006 and 2002 respectively.

RESEARCH TEAM

Ben Upcroft (CI), Chunhua Shen (CI), Ian Reid (CI), Andy Davison (PI), Matt Dunbabin (AI), Niko Sünderhauf (RF), Sareh Shirazi (RF), Guosheng Lin (RF), Adam Jacobson (PhD), Zetao Chen (PhD), Sourav Garg (PhD), Sean McMahon (PhD), Fayao Liu (PhD)

Our project will develop robust means for determining a robot's location in the world using visual sensing alone. This is a crucial robotic competence yet presents significant challenges given a wide range of environments as well as variation due to weather, daily or seasonal cycles and structural changes. The most effective solutions to date have cast the problem as one of image retrieval, but a major weakness is that this relies on the images being acquired under similar perceptual conditions.

We are approaching this as a multi-faceted machine learning problem, in which local motion and place recognition will leverage the power of deep and structured learning methods. First, we will develop coarse localisation techniques that treat place recognition in a manner analogous to object recognition (SV-2) based on the most suitable learnt features (SV-1). We will learn mappings from low-resolution imagery to descriptors that capture the overall gist of the scene, which will be used for contextual priming for feature selection, as in RV-2, to improve long-term robustness of place recognition. Second, we will learn the mapping from local image change to differential motion, also casting visual odometry as a learning problem that will assist in learning the appearance of locations online. Third, we will address the question of place recognition under varying environmental and imaging conditions as instances of domain adaptation and transfer learning, learning the functions that map the appearance of scenes across viewing conditions, for example mapping daytime appearance to night-time. This mapping can be used to predict the scene features that would have arisen at a given locale for a known set of imaging conditions.

We are leading research into the problem of visual place recognition on robots on an international scale, and translating these research outcomes into other research fields including computer vision and machine learning and industry application areas. Our two primary research streams involve developing:

1. innovative methodologies for performing visual place recognition including deep learning approaches, hand crafted approaches, sequence-based approaches, learning-based techniques including both offline and online approaches, training-free methods and closed-loop place-semantic techniques.

2. a modular "black box" place recognition system which is widely used by researchers across robotics, computer vision and machine vision, and applied fields including construction and infrastructure monitoring, on robots, vehicles and personal devices such as personal navigation systems, and a sequence-based version leveraging the success and popularity of the SeqSLAM algorithm.
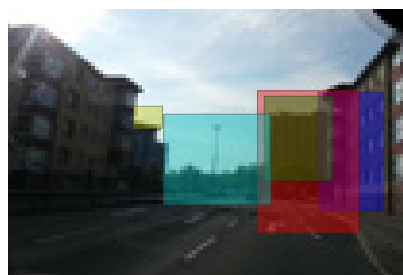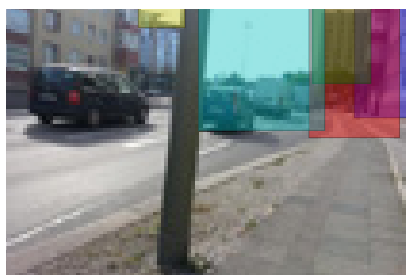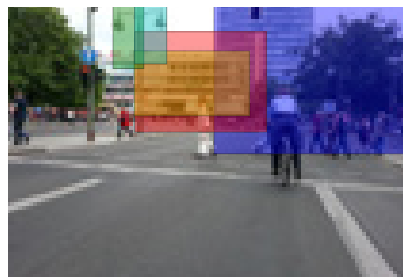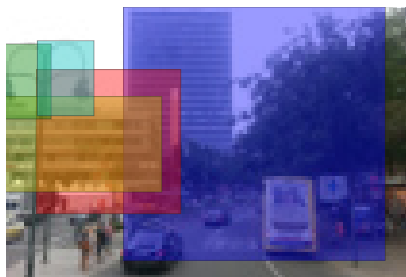
We present a novel place recognition system that adapts state-of-the-art object proposal techniques to identify potential landmarks within an image. The proposed system utilises convolutional network features as robust landmark descriptors to recognise places despite severe viewpoint and condition changes, without requiring any environment-specific training. The coloured boxes in the images above show ConvNet landmarks that have been correctly matched between two significantly different viewpoints of a scene, thus enabling place recognition under these challenging conditions. Source: Sünderhauf, N. et al (2015) Place Recognition with ConvNet Landmarks: Viewpoint-Robust, Condition-Robust, Training-Free, Proceedings of Robotics Science and Systems (RSS).

Algorithms & Architecture AA Theme

will create novel technologies and techniques to ensure that the algorithms developed across the themes can be run in real-time on robotic systems deployed in large-scale real-world applications

Theme Leader
TOM DRUMMOND

Tom Drummond grew up in the UK and studied mathematics for his BA 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 and was subsequently promoted to Senior University Lecturer. In 2010 he returned to Melbourne and took up a Professorship at Monash University.

Tom's research is principally in the field of real-time computer vision (i.e. processing of information from a video camera in a computer in real-time typically at frame rate). This has applications in augmented reality, robotics, assistive technologies for visually impaired users as well as medical imaging.

Spatio-temporally robust visual representations are essential for long term reliable robot operation in real world environments and we are building robust systems for place recognition that will work under a wide range of environment types and perceptual conditions.

Our Algorithms and Architecture theme will revolutionise robotic vision by developing platforms to enable robots to use networked services. Currently robots carry both the vision sensors and computational resources they require relying on old-fashioned system architectures to integrate information on board the robot. AA will reduce this task to a single framework that will enable the distribution of vision sensing and computational resources across a network of robots and cameras.

Cameras produce a prodigious amount of data, from which the essential information must be extracted. Typically images with order 10^7 of pixels are converted into order 10^3 features that are more manageable by computers. However, robotic devices have limited computational capability and this has led to important work on computationally efficient low-level vision algorithms for automatic detection as well as describing and matching of features on computationally-limited devices like mobile phones. Some of these descriptors have compact representations that allow ready communication.

Features are tracked between cameras and over time to create 3D geometric models of the robot's environment. The ability to do this using video data from a single camera, and to localise a camera within this model (visual SLAM), all in real-time, have been major breakthroughs in robotic vision in the last decade. The reason for this is that the computational complexity of dealing with real-time video data pushes the boundaries of current computational resources. There has been considerable research on improving computational efficiency including particle-based methods, Kalman filters, or hybrids such as FastSLAM. More recently, sparse matrix methods have been applied to this problem to reduce complexity while avoiding inconsistencies arising from incorrect assumptions of linearity. Features have also been used

for statistical descriptions of places for recognition to aid tasks such as loop detection. Even with these advances, online dense 3D scene understanding is only feasible at small spatial scales and requires energy intensive computing not available on embedded systems.

We are working towards developing a Vision Operating System (VOS) that provides a framework for distributed implementation of robotic vision tasks across a dynamic network of sensors and computational resources. We will exploit this architecture to develop algorithms that integrate semantic information into persistent visual representations of the environment created by a network of robots and to model highly dynamic visual environments in real time. Our research will develop an architecture that provides vision as a networked service to make available pre-optimised features incorporating knowledge of scene dynamics as required by the client (AA1). We will also use the results of SV2 and SV3 to apply distributed computation to scene understanding (AA1). In 2018 we will start working on an architecture that allocates computing resources as required by the task (AA2). We are creating novel technologies and techniques to ensure that the algorithms developed across the themes can be run in real-time on robotic systems deployed in large-scale real-world applications.

AA1 Vision as a network service for robots in highly dynamic environments

*delivering realtime solutions to demanding robotic vision problems via networking*

Project Leader
TOM DRUMMOND

RESEARCH TEAM

Hongdong Li (CI), Peter Corke (CI), Frank Dellaert (PI), Ahmet Sekercioglu (AI), Laurent Kneip (AI), Khurrum Aftab (RF), Dinesh Gamage (RF), Viorela Ila (RF), Andrew Spek (PhD), Ben Harwood (PhD), Vincent Liu (PhD), Thanuja Dharmasiri (PhD), Josh Weberruss (PhD), Yan Zuo (PhD), Will Chamberlain (PhD), Pan Ji (PhD), Juan Adarve (PhD)

Our project is transforming a robot's sense of vision into a resource accessed over the network. A key aspect of our work is treating all vision sensors throughout the environment as potential resources for a robot moving in a complex environment. We will develop a framework for smart, resource-aware network camera access. It will create a communications protocol and embedded processing architecture that provides low-level vision services to a network of distributed vision processes. Due to the extremely high data rates associated with raw vision sensor outputs, network bandwidth limitations will always prevent the sharing of full image sequences between distributed vision sensors and computing processes. This project will resolve this technological road-block by introducing a bi-directional communications protocol and tailored processing architecture that adaptively specifies the vision information that is exchanged between the embedded sensor and the computing process.
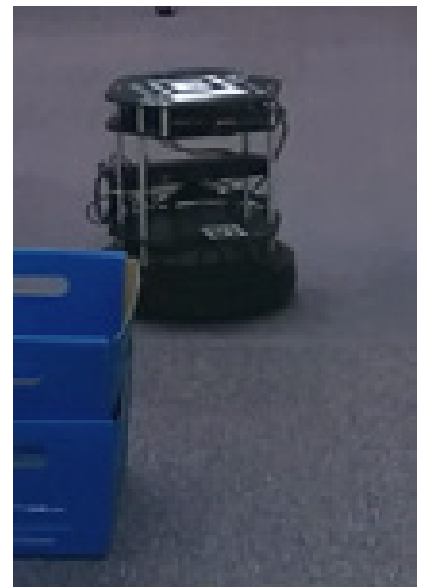
Mobile robot navigating and avoiding obstacles using offboard networked cameras. Source: Chamberlain W., Leitner J., Drummond T., and Corke P. (in press) A Distributed Robotic Vision Service, 2016 IEEE International Conference on Robotics and Automation (ICRA).

Our project addresses the problem of combining multiple sensing and computational resources distributed on robots and embedded in the environment to solve large collaborative robotic vision problems. To achieve our objective requires developing and integrating a suite of technologies. We plan to develop a VOS, Vision Operating System, that allows sensing, and computational resources to be combined to solve robotic vision problems. To build our VOS we are developing computationally efficient algorithms that can be provided as services. These services are the key to constructing rich robotic systems within a network architecture. Our emphasis is on delivering realtime (framerate) solutions to demanding robotic vision problems. There are three kinds of activities needed to do this:

1. Creating a low level toolkit of computationally efficient algorithms for robotic vision
2. Curating additional tools (e.g. existing open source software)
3. APIs for providing these algorithms as services in a network environment.

There are many components that need to be developed within the toolkit to address a range of computationally demanding problem types:

- dense methods using every pixel in RGB (and possibly) range images
- problems with large numbers of parameters
- regularisation and preconditioning of large sparse matrix problems
- problems with large datasets
- representations of large scale environments
- dynamic environments robustness to and modelling of moving objects
- support for computationally efficient filtering and semantics

Because of the recently available computational capabilities of GPUs, these play an extremely important role in the development of high performance algorithms. We are fortunate to have access to excellent GPU facilities at all our nodes, facilities which form the basis of high performance convolutional neural networks and machine learning (for example, the CuDNN plugin for Caffe).

National Benefit

The technologies developed by the Centre will deliver important economic, environmental and social benefits for Australia.

NATIONAL RESEARCH AND INNOVATION PRIORITIES

Robots that see will be a transformative technology that offer solutions to important economic and social problems facing Australia in this century such as labour shortages and low productivity growth in key industries (by decoupling labour growth from population growth); diminishing competitiveness due to high-wage rates; rising OH&S compliance costs; ageing infrastructure; rising healthcare costs; and growing demands for minerals, energy and food that must be met in the context of the community's growing environmental expectations.

In mid-2013 the Australian government defined 15 Priorities for Strategic Research to gradually replace the established National Research Priorities (NRPs) by mid 2014. These Strategic Research Priorities (SRPs) address five societal challenges facing Australia and the world (see SRPs p. 44).

All of these challenges 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 that can respond to change. 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 all of Australia.

In April 2015 the government defined a set of national Science and Research Priorities that we have also mapped our research program against (see SRPs p. 45), identifying overlap in the areas of Food; Soil & Water; Transport; Advanced Manufacturing, Environmental Change; and Health.

The 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 DATA61 and DSTG 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

National Benefit

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015
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%
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
▪ Number of fellowships		2	3	3	3
▪ International grants		0	1	1	0

Strategic Research Priorities (SRPs)

Five major societal challenges have been defined, with three strategic research priorities (SRPs) identifed to help address each challenge

1. **Living in a changing environment**
   - Identify vulnerabilities and boundaries to the adaptability of changing natural and human systems
   - Enable societal transformation to enhance sustainability and wellbeing
   - Manage risk and capture opportunities for sustainable natural and human systems

2. **Promoting population health and wellbeing**
   - Optimise effective delivery of health care and related systems and services
   - Maximise social and economic participation in society
   - Improve the health and wellbeing of Aboriginal and Torres Strait Islander people

3. **Managing our food and water assets**
   - Optimise food and fibre production using our land and marine resources
   - Develop knowledge of the changing distribution, connectivity, transformation and sustainable use of water in the Australian landscape
   - Maximise the effectiveness of the production value chain from primary to processed food

4. **Securing Australia's place in a changing world**
   - Improve cybersecurity for all Australians
   - Manage the flow of goods, information, money and people across our national and international boundaries
   - Understand political, cultural, economic and technological change, particularly in our region

5. **Lifting productivity and economic growth**
   - Identify the means by which Australia can lift productivity and economic growth
   - Maximise Australia's competitive advantage in critical sectors
   - Deliver skills for the new economy

overlap of our Centre's research program with the priorities is highlighted in colour
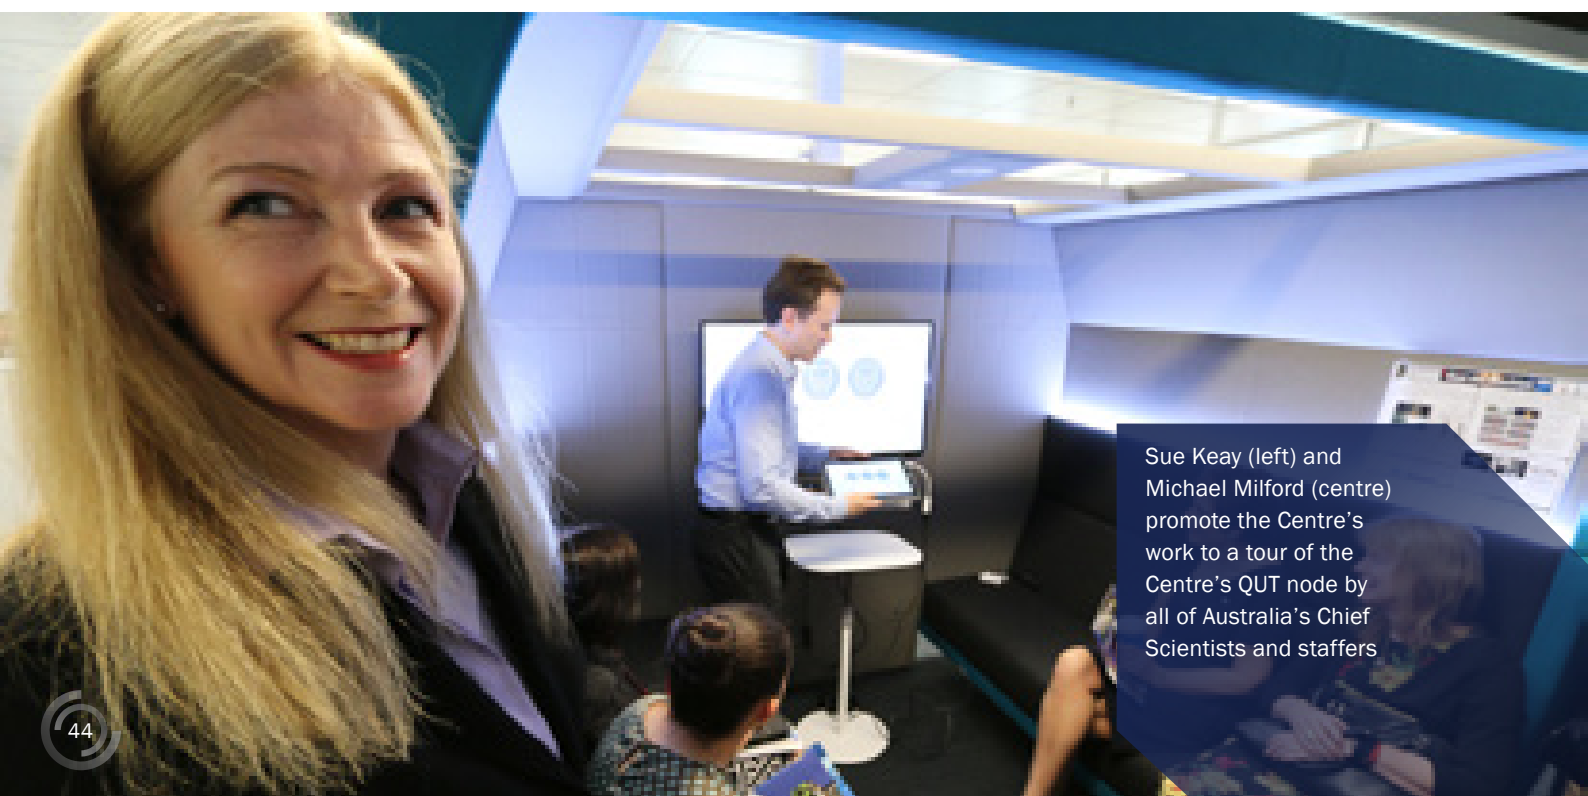
Sue Keay (left) and Michael Milford (centre) promote the Centre's work to a tour of the Centre's QUT node by all of Australia's Chief Scientists and staffers

SRPs (April 2015)

In April 2015, the Australian government identifies nine science and research priorities along with associated practical challenges:

overlap of our Centre's research program with the priorities is highlighted in colour

1. Food

- knowledge of global and domestic demand, supply chains and the identification of country specific preferences for food Australian can produce
- knowledge of the 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 conditions

2. Soil and Water

- new and integrated national observing systems, technologies and modelling frameworks across the soil-atmosphere-water-marine systems
- better understanding of sustainable limits for productive use of soil, freshwater, river flows and water rights, terrestrial and marine ecosystems
- minimising damage to, and developing solutions for restoration and remediation of soil, fresh and potable water, urban catchments and marine systems

3. Transport

- low emission fuels and technologies for domestic and global markets
- improved logistics, modelling and regulation: urban design, autonomous vehicles, electrified transport, sensor technologies, real time data and spatial analysis
- effective pricing, operation, and resource allocation

4. Cybersecurity

- highly-secure and resilient communications and data acquisition, storage, retention and analysis for government, defence, business, transport systems, emergency and health services
- secure, trustworthy and fault-tolerant technologies for software applications, mobile devices, cloud computing and critical infrastructure
- new technologies and approaches to support the nation's cybersecurity: discovery and understanding of vulnerabilities, threats and their impacts, enabling improved risk-based decision making, resilience and effective responses to cyber intrusions and attacks
- understanding the scale of the cyber security challenge for Australia, including the social factors informing individual organisational and national attitudes towards cyber security

5. Energy

- low emission energy production from fossil fuels and other sources
- new clean energy sources and storage technologies that are efficient cost-effective and reliable
- Australian electricity grids that can readily integrate and more efficiently transmit energy from all sources including low- and zero-carbon sources

6. 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

From left to right, QUT's Vice-Chancellor Prof Peter Coaldrake, Dr Anjali Jaiprakash and orthopaedic surgeon Dr Ross Crawford show Minister for Education Senator Simon Birmingham work on medical robotics that could lead to improved healthcare.

7. Advanced manufacturing

- knowledge of Australia's comparative advantages, constraints and capacity to meet current and emerging global and domestic 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 polymers

8. Environmental Change

- improved accuracy and precision in 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

9. 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, for example, using mobile apps, remote monitoring and online access to therapies.

Application Areas

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?

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. Our aim is to create vision-enabled robotic systems that can understand and respond to their environment, that can operate reliably over long periods in complex unstructured surroundings, and that can interact safely and effectively with humans as well as objects. Our success will see robots become a ubiquitous feature in most industries.

Just as the minds of animals developed around the need to support a sense of vision, so too will the capabilities of robots.

ADDRESSING KEY CHALLENGES

Australia will face many challenges in the near future, 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 creating wealth and wellbeing. We expect robots to help provide solutions to meet Australia's future challenges, in line with the nation's new Science & Research Priorities (see SRPs p. 45) including:

- Labour shortages and low productivity growth in key industries;
- Diminishing international competitiveness due to high-wage rates;
- Rising OH&S compliance costs;
- Need for increased productivity due to an ageing population;
- Rising healthcare costs;
- Ageing infrastructure and asset base;
- Growing demands for minerals, energy and food; and
- The need to preserve the environment.

FIVE APPLICATION DOMAINS

We have identified five areas of economic importance to Australia that will benefit from robots with advanced vision capabilities. Robots have the potential to transform industry, yet currently they are underutilised. Labour-intensive industries characterised by unstructured working environments will particularly benefit from vision-enabled robots. The five areas we have identified include:

- **Small and medium scale smart manufacturing**, in which robots will use vision to rapidly learn to produce small runs of highly customised products in flexible collaboration with humans

- **Infrastructure monitoring**, in which robots will visually inspect and monitor infrastructure such as roads, rail, pipelines, tunnels and bridges

- **Agriculture, aquaculture and bio-monitoring**, which require robots to visually and reliably perceive native flora and fauna, crops, weeds, and pests in varying lighting, weather and seasonal conditions

- **Building and construction**, which requires both indoor and outdoor operation as well as interaction with human operators in a dynamic, unstructured environment

- **Medical and healthcare**, where robots can be used to assist surgery, to transport people and materials, and to provide care and companionship.

TECHNOLOGY DEMONSTRATORS

Our research projects will contribute to technology demonstrators relevant to our five application areas. These proof-of-concept demonstrations show how vision-enabled robotics can operate in realistic scenarios motivated by each of our application domains. We will work with end-users to generate investment in the engineering resources required to build fully functional prototypes, which can then be tested by end-users in practical applications. We aim to attract an additional $7 million in investment for prototype development and other end-user deliverables over the lifetime of the Centre, we have already attracted $3m. Industry interaction will be encouraged via many mechanisms including industry forums and invitations to see technology demonstrators (see Industry Engagement p. 62). Where significant industry demand is identified, our Centre will run continuing professional education courses in robotics, vision and/or robotic vision and has already conducted a national robotic vision workshop series (see Construction Industry workshop p. 57).

Robotic boats can be used for environmental monitoring and search and rescue, pictured are QUT's Roboboat (left) and the yellow Inference system of environmental monitoring boats developed by Matt Dunbabin

Smart Manufacturing

Developing advanced robots that can use their sense of vision as smart co-workers, watching what human workers do and anticipating what comes next

Manufacturing Portfolio Lead
TOM DRUMMOND

In 1961, the first manufacturing robot went to work in a General Motors (GM) automotive plant, unloading a die casting machine. Today we see more sophisticated versions of these robots working in the Tesla motor factory where 10 of the world's largest robots, given heroic names from X-men like Wolverine, Vulcan and Collosus, lift cars and install battery packs. In many respects, manufacturing robots are a mature technology, however we see the potential for tremendous growth.

The use of advanced industrial robots is at hand, with factories set to begin using automation at an increasing pace.

Reports of the demise of the Australian manufacturing sector have been greatly exaggerated. The Australian manufacturing sector continues to represent a core source of economic prosperity for Australia – and all indications suggest it will continue to do so. In terms of its contribution to Australia's sustained economic performance and its capacity to generate quality jobs, manufacturing has outperformed many other sectors and produces around $100 billion of Australia's gross domestic product, employs around 900,000 Australians (7.4%), and contributes a disproportionate 25.7% ($4.8b) to Australia's total business R&D expenditure (ABS data).

However Australia's manufacturing workforce is less skilled than in other sectors (45% without post-school qualifications versus 39% average across all industries). To become competitive, Australia must shift from heavy industrial manufacturing towards higher value-added, technologically advanced production. Our economy is undergoing a significant transition in which traditional industries are evolving and new opportunities are emerging. Together industry and research are poised to forge an exciting new path and accelerate Australia's transition into high-value, knowledge based manufacturing, creating sustainable jobs in high-tech industries.

Australia is a relatively high labour cost country and consequently the health of the manufacturing sector is dependent on both automation and also maximising the value that humans bring. Our Centre is therefore particularly interested in human-robot cooperative solutions that combine the best capabilities of both. To that end, we are developing a demonstrator to illustrate how humans and robots can cooperate in a single workspace, for example, something complex like assembling an IKEA flatpack.

Infrastructure monitoring and asset inspection

Robotic vision systems that can see what needs to be seen, analyse those images and provide condition reports on the health of these assets

Infrastructure Portfolio Lead
ROBERT MAHONY

Australia has a vast and aging network of infrastructure assets that require inspection and monitoring. The quality of Australia's infrastructure is both a reflection of our economic prosperity and an indicator of our potential for future growth. According to Engineers Australia (http://tinyurl.com/ln4t6ha), to capitalise on our productive capacity, we must make the best use of the infrastructure we have and that means regular inspection and monitoring. Critical assets typically have asset protection schedules, while many are monitored online as well to ensure their integrity. Second tier assets are often monitored manually with spot-checks from the field or preventive maintenance checks. Automated monitoring and inspection of essential assets using robotic vision is more likely to catch sudden failures or abnormal operation than maintenance based on historical reliability. Operation, safety, and environmental incidents can be avoided, thus increasing the reliability of important infrastructure assets and potentially extending their life.

Infrastructure monitoring is a prime target for deployment of the Centre's expertise, which can enable:

- Precise localisation of robots or cameras performing monitoring

- Precise repeated visual servoing of robots or people through an environment in order to facilitate change detection (for example, identifying crack propagation in a road)

- Precise sensory alignment – alignment of sensory data at both a "video" level and within sensory snapshots – for example registering and aligning two images of the same place under significantly different environmental conditions.

We are currently investigating applying our place recognition technology to facilitate aircraft maintenance, and road crack monitoring.

Agriculture, aquaculture and bio-monitoring

requiring robots to visually and reliably perceive their environments and conduct operations in agriculture as well as environmental monitoring

**Agriculture Portfolio Lead
TRISTAN PEREZ**

Our vision is to develop and fast track farm robotic technology that will reinvigorate productivity through increased production and reduced costs

With a world population projected to reach almost 9 billion by 2050, sustainability and food security are significant challenges for all countries. Australia, in particular, faces a real challenge to ensure its participation in food production is both competitive and sustainable. Robots can enhance the agricultural productivity required to achieve sustainable production of food. Robots can maximise the efficiency of inputs and, at the same time, collect data that can help in the management of agricultural produce (crops and livestock). Robotic technology will soon have a significant impact on agricultural practices.

Our vision is to develop and fast track farm robotic technology that will reinvigorate productivity through increased production and reduced costs. The Strategic Investment in Farm Robotics (SIFR) program is funded by Queensland's Department of Agriculture and Fisheries (QDAF), and is part of the Australian Centre for Robotic Vision at QUT. Such robot technology can be used to conduct autonomous multi-vehicle operations in applications of weed management, fertilising, and seeding. Weeds cost Australian farmers around $1.5billion pa in weed control activities, and a further ~$2.5billion pa in lost agricultural production. The use of multiple, relatively low cost, vision-enabled field robots enables novel alternative weed destruction methods based, for example, on mechanical and thermal principles rather than herbicides. Such advances could reduce the input cost associated with weed management operations in terms of energy, labour, and improved chemical delivery by up to 90%.

Agricultural robots also have the capability to communicate, not only with each other, but also with

AgBot II with the QUT project team, which includes (not all pictured) Prof Tristan Perez (Project Leader and Lead – Dynamics and Control); Mr Ray Russell (Lead – Mechanical Design); Mr Owen Bawden (Lead –Industrial Design); Dr Jason Kulk (Lead – Software Design); Dr Christopher McCool (Lead – Robotic Vision); Dr Christopher Lehnert (Lead – Manipulation); Dr Inkyu Sa (Lead – Visual Servoing); Dr Feras Dayoub (Robotic Vision), Dr Andrew English (Engineering Support), A/Prof Ben Upcroft (Robotic Vision), Dr Tanya Scharaschkin (Plant Physiology), Prof Margot Brereton (Computer-Human Interaction); Prof Paul Hyland (Business and Economics), A/Prof Jennifer Firn (Ecology), Ms Victoria Sullivan and Ms Tyaka Ringrose (Administrative Support)

unmanned aircraft and operations managers in order to combine different kinds of environmental and field information. Our robots will be equipped with sensors for navigation and crop data collection. We will use sophisticated algorithms for data-fusion to extract information from the sensor data, allowing us to use low-cost sensors, such as cameras, and yet obtain high accuracy in robot localisation and navigation. Cameras are used to navigate, detect and avoid obstacles, and also for weed detection and classification as well as to control variable rate technology in herbicide application.

As well as being used for tasks related to field and crop management, robots enable new management practices and data collection, which will lead to advances in the field of precision agriculture. When data collection is combined with an appropriate digital infrastructure, the result can be more refined site-specific crop management leading to increased performance and robustness of crop agricultural enterprise systems. This will optimise productivity (increase of yield and quality) and profitability (optimising the return on investment in energy, water, and labour), while maintaining performance (reduce volatility) in the face of climate variability, incomplete information, market movements and other threats that may be biophysical or socio-economic.

Apart from applications in farms, field robotics have the potential to play a key role in the understanding and management of our natural environments. As in the case of farm robotics, this technology can be used to both collect key data and also apply actions by interacting with the natural environment. These robots could soon become a key tool for ecologists, and environmental scientists who are trying to understand and manage complex systems with spatio-temporal variability.

In our 2014 Annual Report we profiled AgBot II, an innovative agricultural robot prototype fully designed and fabricated by QUT researchers and engineers, with support from the Queensland Government. Part of a new generation of crop and weed management machinery, AgBot II is designed to work in autonomous groups in both broadacre and horticultural crop management applications.

The robot's cameras, sensors, software and other electronics enable it to navigate through a field, apply fertiliser, detect and classify weeds, and kill weeds either mechanically or chemically. AgBot II takes information on the geometry of a paddock, sets its path and follows it, covering

the area for weed control and crop nutrient management. This technology promises to reduce the cost of weeding operations by approximately 90%, which could save the Australian agricultural sector $1.3 billion pa.

In June 2015, trials of the AgBot II prototype were carried out at the Queensland Government's Redlands Research Station, with outstanding results. In these field trials, AgBot II achieved an overall success rate in weed detection and classification above 90%. Its highest performance was with cotton (97.8% recognition) and wild oats (97.3%), its lowest performance with sowthistle (82.0%). The trials also demonstrated the success of spot-spraying selected weed species and using a robotic hoe to mechanically remove weeds from the soil.

QUT is now seeking commercial partners both in Australia and overseas to move the AgBot II prototype into real world applications. The SIFR team has also developed a fruit-harvesting robot called, "Harvey" profiled in Vision and Action p.30.

Aquaculture and Bio-Monitoring

Centre CI Michael Milford and team have been working with ecologists over the past 12 months on developing techniques for automated monitoring and analysis of ecological systems including grassland and forests.

Large-scale ecological studies are labor intensive and not necessarily repeatable – because the workload is shared between many ecologists, often performed over multiple disparate locations. Analysis performed by hand (as most analysis is performed currently) can benefit from automated approaches. We are currently

developing an automated approach to monitoring grassland and its response to pollutants and other chemicals with the feasibility of it being deployed across 40 monitoring sites around the globe. This work will benefit from collaboration with novel sensing modalities developed in the Centre's Robust Vision theme (see RV1 p.37).

Matt Dunbabin, a Centre AI, has been contributing to solve an ongoing environmental threat to Australia's Great Barrier Reef (see p. 54). COTSbot (Crown-Of-Thorns Starfish robot) aims to be a revolutionary advancement in robotic environmental monitoring and management, specifically to increase the efficiency of Crown-Of-Thorns Starfish (COTS) eradication.

Crown-Of-Thorns Starfish (Acanthaster planci) are described as one of the most significant threats to the Great Barrier Reef. Since the 1960's, land-based nutrient runoff has accelerated outbreaks of COTS, which are destroying large areas of reef. With few natural predators, traditional control of COTS required manually injecting the starfish in excess of 10 times with a biological agent. In 2014, a new agent developed by James Cook University (JCU) requires only one injection per starfish. This advancement provided the stimulus for us to revisit automated (robotic) COTS population control and monitoring.

The intention of COTSbot is to provide a proof-of-concept robotic system that consolidates prior and ongoing research into image-based COTS detection, robotic vision, manipulator control, and shallow water Autonomous Underwater Vehicle (AUV) design, navigation and control, to directly facilitate COTS reduction. This multi-stage project will validate and demonstrate AUV performance to

stakeholders and ensure the system components are a useful and flexible enabling foundation technology for environmental monitoring beyond the problem of COTS control.

A fundamental component of COTSbot is its ability to automatically and robustly detect Crown-Of-Thorns Starfish in complex coral reef environments. A new state-of-the-art vision and classification system has been developed to run on-board the AUV.

Previous research on image-based COTS detection has increasingly improved the detection accuracy of the system and its ability to run on-board modern lower power computers that can be installed within the space and power constraints of the underwater robot. The current COTSbot system is a significant further advancement that exploits state-of-the-art techniques in machine learning to achieve a detection performance of well over 99% whilst being able to operate real-time on-board the AUV.

Queensland University of Technology (QUT) commenced the first stage of the proof-of-concept COTSbot AUV testing in January 2015, funded by QUTBluebox (QUT's commercialisation arm). A working prototype was tested on the reef in October 2015. The COTSbot team are actively working with potential industry commercialisation partners to bring this technology into the real world.

The Crown-Of-Thorns Starfish (COTS) robot (COTSbot)

It's one of the biggest threats to an Australian national treasure – the Great Barrier Reef. But now, a revolutionary new advancement in robotic environmental monitoring and management could help in the fight against Crown of Thorn Starfish, or COTS.

The Crown of Thorns Starfish Robot, the COTSbot, is the world's first robotic submersible that will seek out and inject COTS with a toxin. The goal is to control their numbers, which are responsible for an estimated 40 per cent of the reef's total decline in coral cover.

Looking like a mini-yellow submarine, COTSbot can move up and down through the Great Barrier Reef with the help of five thrusters, GPS and pitch-and-roll sensors.

The unique aspect of the COTSbot is that it can think for itself. It uses a visual recognition system made possible by Centre of Robotic Vision algorithms to identify crown of thorns starfish in the visually challenging environment of the Great Barrier Reef.

The robot is designed to cruise about a metre above the coral surface, looking for COTS. When it sees one, a robotic arm extends to inject the starfish with vinegar, proven by James Cook University to be effective in controlling COTS numbers within 24 hours of application.

The COTSbot was designed by Dr Matt Dunbabin, a QUT Research Fellow and Associate Investigator with our Centre. He said the robot was intended as a first responder system to beef up the existing program that uses divers to hunt for COTS.

"Human divers are doing an incredible job of eradicating this starfish from targeted sites but there just aren't enough divers to cover all the COTS hotspots across the Great Barrier Reef," said Matt.

"Imagine how much ground the programs could cover with a fleet of 10 or 100 COTSbots at their disposal, robots that can work day and night and in any weather condition."

QUT roboticists spent months developing and training the robots to recognise COTS among coral, using still images of the reef and videos taken by COTS-eradicating divers.

Dr Feras Dayoub, a QUT Research Fellow specialising in Robotic Vision, trained the COTSbot to pick out the pest from other sea life.

"The system has seen thousands of images of COTS and not COTS and now it's able to detect and decide which one is COTS and which one is not," Feras told the ABC.

The story about the COTSbot first aired on ABC Television on 31 August, with the BBC World News and NBC picking up the story.

Check out the ABC story here **http://tinyurl.com/zt5hnhx**.

COTSbot researchers Matt Dunbabin (left) and Feras Dayoub (right)"

Building and construction

Robots with vision have strong potential to transform the building and construction industry, which has significant economic importance to Australia

Construction Portfolio Lead
IAN REID

Visual sensing for robotics has strong potential to transform numerous aspects of the building and construction sector. This industry, which has significant economic importance to Australia, is characterised by a large, expensive labour force, and encounters issues such as safety, productivity, and quality control.

Building sites are largely unstructured, which is problematic for current generation robots. Our objectives in the Centre are to develop visual sensing technologies for scene understanding in such dynamic and unstructured environments. This will lead to increased automation, enabling robots to operate autonomously or semi-autonomously in ways that improve safety, assist humans or replace unskilled labour. Within the construction environment examples will include: smart cranes, pallet trucks or scissor lifts that can work with humans on site; mapping of dynamic scenes to assess the geometry of buildings under construction; situation awareness for safety monitoring and assessment such as looking for people entering no-go areas or finding trip hazards.

This latter problem – finding trip hazards – has been selected as a test case for work by one of the Centre's

PhD students, Sean McMahon, who is investigating learning affordances using deep neural networks.

The Centre has been actively developing relationships within the construction industry and negotiating site access to acquire data; gaining an understanding of existing working practices; discussing potential applications of visual sensing to the automation and monitoring in the construction industry; and using these to inform the research program (see Construction Industry Workshop p. 57).

In particular the Centre is delighted to have Prof. Andrew Harris on its End User Advisory Board. Prof. Harris is Australian Director of Laing O'Rourke's engineering and innovation consultancy, the Engineering Excellence Group. Laing O'Rourke is Australia's largest private engineering and construction business.
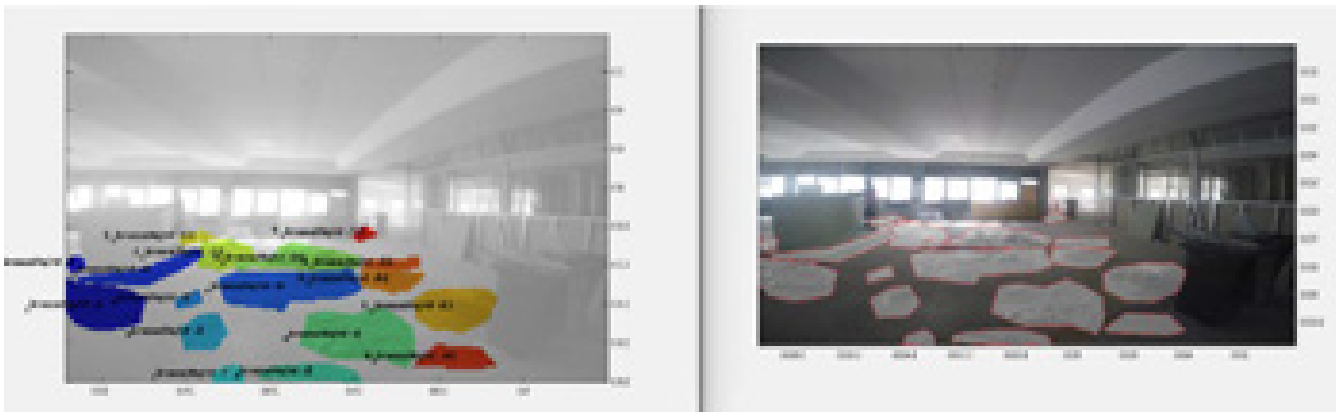
Example image from a training set with hand-labeled potential trip hazards on a construction site (images courtesy of Sean McMahon)

Construction Industry Workshop

It's a sight we've seen for decades - robotic technology in the controlled environment of an assembly line, but what about a construction site, where the site is cluttered, uneven, and outdoors? Questions like that were the focus of the Australian Centre for Robotic Vision's "Industry Engagement Collaboration Workshop."

Construction industry association representatives, the Gold Coast City Council and the Institute of Engineers attended the workshop on 20 April 2015.

Centre Director Peter Corke provided an overview of the evolution of robots in industry from Devol and Engelberger's UNIMATE, robots that performed spot welding and extraction of die-castings for General Motors in 1961, to more recent high tech robots that walk and run. As robots evolve they are beginning to perform diverse tasks ranging from the extremely hazardous to the everyday task of keeping your floors clean.

Peter had two key messages for the construction industry. The first is that robots are different to how they are portrayed in the media, and second, just like the rise of the computer in the 1960's, robots that see will be the next disruptive technology. Vision will give robots the ability to work alongside people safely, efficiently and effectively in any environment. This will ultimately change the way we work and live.

The main outcome of the workshop was to develop an understanding how research into robotic vision should progress given the requirements, demands and constraints of the construction industry.

The construction industry members put forth some interesting challenges that robots could face:

- Acceptance of new technologies by the construction industry is slow

- There is a difference between domestic construction and commercial construction sites

- Construction sites never have the same surface, and that surface can be cluttered

- How to handle materials for a Just In Time (JIT) inventory system

- Constructions sites are outdoors, and face varying weather conditions

- Regulated construction working hours

- Workplace, health and safety requirements

However, many opportunities for robotic vision in the construction industry were identified. They include cleaning construction sites, materials handling, lifting equipment and inspection of sites and structures.

The group also identified one interesting opportunity involving the solar power industry. Robots with vision could be deployed to identify solar panels that require cleaning and repair. This is a critical issue, because damaged or dirty solar panels can really hurt efficiency when it comes to capturing solar radiation.

A special thanks goes to Samantha Perkins from the Australian Institute of Commercialisation for organising and facilitating the workshop.

Medical and healthcare

We are seeking to develop a robotic tool to take care of manipulation tasks while the surgeon takes care of the decision-making
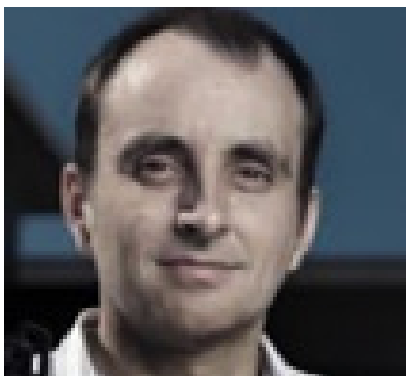
Medical and Healthcare Portfolio Lead
JONATHAN ROBERTS

As our population ages and as new medical techniques and procedures are developed, there will be a need to reduce the cost and improve productivity in the health sector. Hospitals are very large businesses where logistics are just as important to the costs of running as the medical treatment. There are already examples of mobile robots to transport hospital trolleys where they are needed, rather than relying on human porters (e.g., the Aetheon Tug robot). According to the International Data Corporation IDC http://tinyurl.com/glzu2nv, robotic capabilities are expanding while increasing investment in robot development is driving competition and helping to bring down the cost of robots, with two of the fastest adopters being process manufacturing and healthcare. The healthcare sector is expected to see a bump in robotics purchases, with spending expected to double by 2019.

We are working in the area of medical and healthcare robotics that focus on ways to help improve the efficiency of surgeons, and in particular in the area of Minimally Invasive Surgery such as arthroscopy. The use of arthroscopy,

where a camera and cutting tool is inserted into a joint such as the knee, is now common place in the developed world. However, it is not yet widely available in the developing world due to a lack of specialised surgeons. Currently, surgeons use their own vision to understand the operation before them and their own hands, and feet, to manipulate the camera and their tools. The procedure is difficult to master and takes a long time to learn. Our research is developing the techniques and technology that will enable future robotic assistants that will work in tandem with a human surgeon in performing joint arthroscopy. In particular, the research is developing robotic vision systems that are capable of mapping joints in real-time via arthroscopically sourced video streams. Our research will also explore control schemes that allow robots to hold and manipulate both the arthroscope and the surgical tools using robotic vision in the control feedback loop (visual servoing).

Our medical robotics researchers also have an interest in the use of vision-based machine learning techniques to lower the cost of diagnostic testing of blood and retinal images. The

team is exploring the use of machine learning techniques while at the same time developing new cost-effective mobile phone-based image acquisition systems.

2015 was the first full year of operation of the Centre's Medical and Healthcare robotics program and the focus was on building a team and obtaining the equipment required to carry out the research. As of the end of 2015 there were three full-time research fellows (Dr Anjali Jaiprakash, Dr Anay Pandey, Dr Liao Wu) dedicated to working on medical and healthcare robotics projects. We have commenced three PhD students (Dr Chris Jeffrey, Andres Marmol Velez, Mario Strydom) and we have won a number of equipment grants to enable us to start a viable surgical robotics lab. One of the new labs at QUT is currently being used as our medical robotics lab and visitors have been impressed with our first (basic) demos. We have spent the year networking with industry - both end users of medical robotics (hospitals) and developers of medical robotics. Our first publication in the medical robotics area was a paper presented at the Australasian Conference on Robotics & Automation.
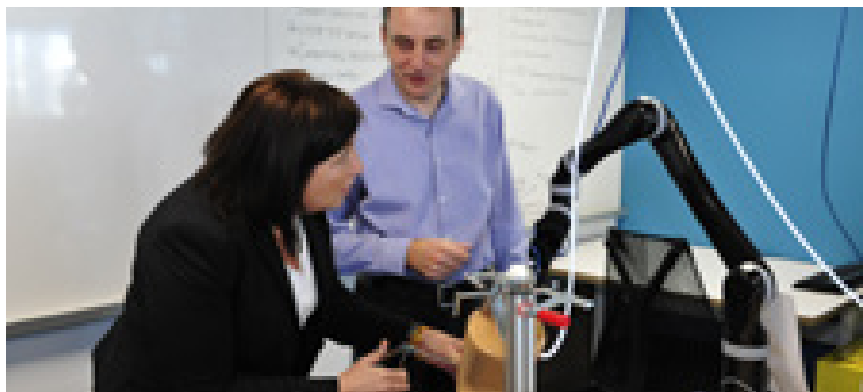
Jonathan Roberts (right) demonstrates medical robotics research to Queensland's Minister for Science and Innovation, the Hon. Leeanne Enoch MP.

Communication, Engagement, and Networks

Our Centre engages with industry, business, government and the community through its Communication, Engagement and Outreach Programs. Our robotic vision research is creating technology to enable robots to deliver on their potential and transform the world.

For our research to be effective our Centre needs to engage with a wide range of stakeholders to ensure that we successfully meet our objectives and generate support for our research in the future.

Critical components of our Communication, Engagement and Outreach strategy are illustrated in our Centre Stakeholder Map (Figure 3), which shows the focus of our Centre and how it relates to our stakeholders in industry, business, government and the community. As the first research centre in the world to bring the fields of computer vision and robotics together, our researchers are directly working towards developing technologies that will help solve the challenges of the future. Successfully communicating this to our stakeholders will ensure that we are well-positioned to meet the many objectives of the Centre.

We aim to build sustainable long-term partnerships across the research sector and public and private enterprises. The Centre's engagement targets can be divided into five key groups including End-Users (see Industry Engagement p. 62). Our End-Users cover the range of Application areas (see p. 47-58) that we have identified as economically important to Australia with high potential for application of robotic vision technologies. We aim to attract an additional $7 million in investment in the Centre to develop prototypes in these application areas.

More generally the Centre also seeks to develop links and potential partnerships with those companies already in the robotic vision space in Australia and we are undertaking an audit of these companies to get a clear picture of the industry in Australia. We aim to take a leadership role in developing a network of these companies (see Industry Engagement Highlights p. 63). It is not the Centre's role to compete with commercial enterprises, instead our aim is to partner with them in delivering technological solutions across a range of industries.

During 2015 we delivered 54 government, industry and business community briefings and 12 public awareness and outreach programs (see End-User links KPIs p. 61). As well as hosting visits from all Australia's Chief Scientists and many industry partners, the Centre also hosted visits from the then Minister for Education and Training The Hon Christopher PyneMP , Queensland Minister for Science and Innovation The Hon Leeanne Enoch MP, Minister for Education Senator Simon Birmingham, the Assistant Minister for Innovation, Wyatt Roy MP, and Queensland Opposition MPs Tim Nicholls and John McVeigh.
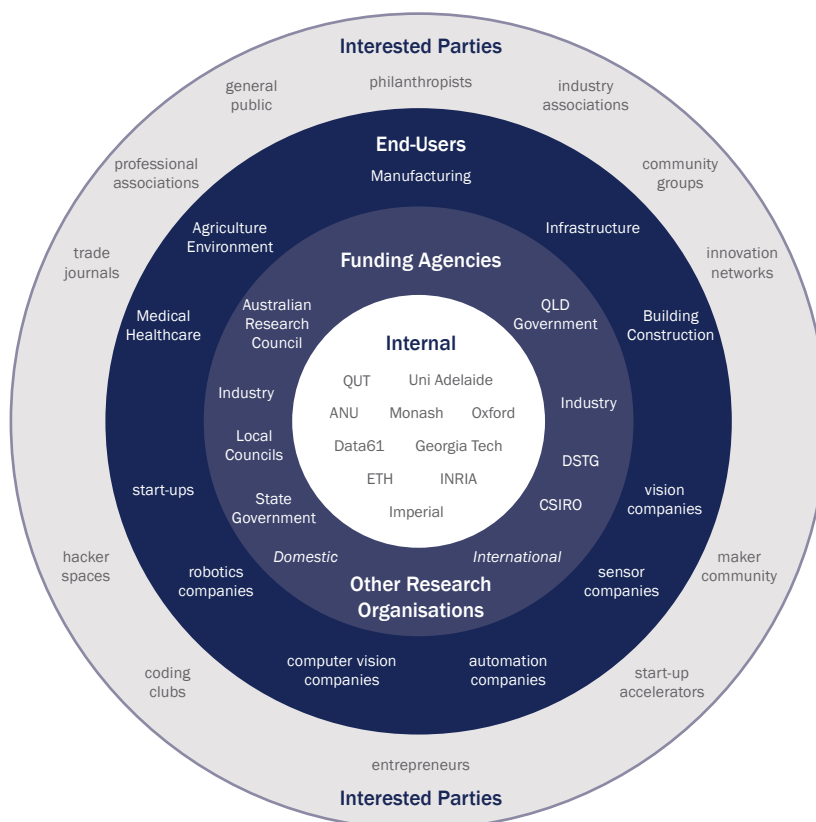

Figure 3: Centre Stakeholder Map

Queensland Minister for Science and Innovation The Hon Leeanne Enoch MP drives our agricultural robot AgBotII

Assistant Federal Minister for Innovation, Wyatt Roy MP (left) visits the Centre's QUT node

End-User Links KPIs

End-user Links

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

Centre Director Peter Corke (right) explains robotic vision to Minister for Education Senator Simon Birmingham

Industry Engagement

Our industry engagement strategy supports the Centre in taking a leadership role in developing the fledgling robotic vision industry in Australia. Outreach to industry has many facets.

End-users and potential end-users have already participated in workshops and briefings held by the Centre and we will continue this engagement by inviting end-users to demonstrations of the Centre's technologies at breakfasts held at our nodes. Centre personnel attend and participate in end-user events. Our Engagement Coordinator, Greg Lee, is actively working with peers at our partner universities to identify and engage with relevant end-users. Further engagement will be through the general media, articles in trade journals, as well as the general community engagement channels (media coverage, website, social media, printed material). 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. We started with an industry engagement day (Ideas Exchange) as part of our official launch (attended by ~80 industry representatives) (see Industry Engagement highlights p. 63). We have also worked with the Australian Institute for Commercialisation to meet with and present to representatives from Queensland's Building and Construction industry (~15 attendees) (see Construction Industry workshop p. 57). In partnership with AusIndustry our Centre also ran a series of successful national robotic vision workshops in 2015 to build links with industry, to engage with other research groups in the robotic vision area and to form links with our regional development associations. These industry workshops were held in Geelong (with the Centre for Intelligent Systems Research CISR, Deakin University), Wollongong (with University of Wollongong's ARC CoE for Electromaterials Science (ACES)), Melbourne (at Monash University), Mackay (with Central Queensland University CQU), Sydney (with University of Sydney's Australian Centre for Field Robotics ACFR) and in Adelaide (at Adelaide University). Total attendance of our industry workshops exceeded 150 people (see Industry Engagement highlights p.63). All of the CIs, the COO and our engagement coordinator have had one-on-one meetings with members of various industry segments related to our Application areas (see Application areas p. 47). Our Engagement Coordinator, Greg Lee, has instituted a register of contacts to keep track of Centre visits and meetings. Our customer-relationship management (CRM) system forms the backbone to our Engagement strategy.

In 2015 we convened our Centre's End-User advisory board (EUAB), chaired by Russel Rankin from Food Innovation Partners. Our EUAB met for the first time in Adelaide on 12th November 2015. Our board has requested more regular (and longer) meetings, at least twice per year. The meetings involve our board members engaging with our Centre Executive, participating in one of our industry breakfasts and assessing how our industry engagement strategy is working, and offering suggestions for improvement.

We aim to expand our industry engagement programs. It is encouraging that the number of industry contacts is picking up, and that industry is approaching us. The challenge for 2016 is to convert these into collaborative research projects. In 2016 we will hold two EUAB meetings, at least four industry showcases and launch an Industry Affiliates program.

Industry Engagement Highlights

The best way to focus research is to know who's going to use your research, and how they want to use it in the 'real world.' Since the launch of the Australian Centre for Robotic Vision, we've been very active in engaging with industry.

We've held an "Ideas Exchange" industry forum, industry workshops in six different cities, hosted numerous visitors, and toured robotics companies in America's Silicon Valley.

"It is very important that we listen to industry to get an understanding of some of the challenges they face," said Centre Director Peter Corke.

The "Ideas Exchange" industry forum was held as part of the launch celebrations. More than 70 people attended the event. They were invited to record their thoughts on the challenges and opportunities of robotic vision.

Professors Jon Roberts (QUT), Tom Drummond (Monash), Peter Corke (QUT), Tristan Perez (QUT) and Ian Reid (Adelaide) presented on the Centre's five focus areas of Medical and Healthcare, Smart Manufacturing, Infrastructure Monitoring, Agriculture and Automation in Construction. Vodcasts of the presentations are available on the Centre's website at roboticvision.org.

The Ideas Exchange was followed later by a series of industry workshops around Australia. In partnership with AusIndustry, we delivered six workshops in four different states:

- Queensland: Mackay
- Victoria: Melbourne & Geelong
- New South Wales: Sydney & Wollongong
- South Australia: Adelaide

Over 150 people attended the six workshops, with attendees responding enthusiastically to the opportunity to learn more about the Centre and its research.

We also had the opportunity to show off the Centre to several key members of industry. They included Stryker Medical Technologies, BMW and Softbank Robotics. Industry also didn't just come to us. We went to them.

Together with our colleagues CSIRO, Blue Ocean Robotics (Denmark), and Callaghan Innovation (New Zealand), we took an international delegation to explore robotics companies in Silicon Valley in September 2015 and to attend the RoboBusiness conference in San José. The Centre's COO Sue Keay worked with CSIRO to develop the program for participants in the Industry track of the Knowledge-Innovation tour, which attracted delegates from eight countries.

The companies we visited in Silicon Valley included: GE (the multinational conglomerate with an interest in future tech), Fetch Robotics (logistics robots), Precise Automation (human safe robotic manipulation), Blue River Technology (agricultural robotics), 3D Robotics (UAVs), and Singularity University, with a detour to see the OSHbot (Orchard Supply Hardware robot) developed by Fellow robots, a company with a strong Australian connection.

Tours like these spark ideas for the fledgling robotics industry in Australia and, as our local industry grows, may lead to the formation of our own Australian-version of Silicon Valley.

Communication

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 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.

The general community has an enduring fascination with robots (see Social Impact of Robotics p. 12). Technology demonstrations, a key outcome from our Centre's research projects, will be disseminated through YouTube or more traditional media such as print, radio and TV. School students will be reached through TV programs such as Totally Wild and Scope, in 2016 we will partner with experts to provide leading professional development workshops for high school teachers, as well as supporting talks in schools by our researchers. Our annual report highlights the achievements of our researchers and is available in electronic form and distributed in hardcopy to selected stakeholders.

Major communication activities in 2015 included, our Centre launch on 9th March attended by Minister Christopher Pyne (see Launch p.65), an associated robot press conference (see Robot Press Conference p.66), multiple TV appearances (see Robots on TV p.67), the publication of the Centre's first annual report, and involvement in QUT's Robotronica, a festival of all things robotic (see Robotronica feature p.14). In 2016 we aim to distribute four media releases related to robotic vision and give at least four public lectures, and continue to host visits and give tours featuring robotic vision to interested stakeholders from government, industry and the community.

Chris Lehnert (right) appears on TV's Scope and Totally Wild with QUT's Baxter Robot and Channel 11 reporter Kellyn Morris (left)

Centre Launch

A Centre with big plans and big goals deserves a big launch. Certainly, the launch of the Australian Centre for Robotic Vision was just that.

More than 200 people gathered at QUT's "The Cube" on 9 March 2015 to watch the Federal Education Minister, The Hon. Christopher Pyne, officially open our Centre. Baxter the robot assisted him by cutting the ribbon!

"Robotic vision is the key enabling technology that will allow robotics to transform labour-intensive industries and improve productivity," Mr Pyne told the crowd at the launch.

"What this new ARC Centre of Excellence will do," he said, "is improve robot technology to create a new generation of robots that can visually sense and understand complex and unstructured real-world environments."

Other notable attendees included:

- **Professor Aidan Byrne**, CEO of the Australian Research Council (ARC)

- **The Hon. Leeanne Enoch**, Queensland Minister for Science and Innovation

- **Dr Geoff Garrett**, Queensland Chief Scientist

- **Professor Peter Coaldrake AO**, QUT Vice Chancellor

- **Councillor Ryan Murphy**, representing the Brisbane Lord Mayor

QUT Deputy Vice-Chancellor (Research and Commercialisation) Professor Arun Sharma said the university was proud to host the Robotic Vision Centre.

"This global initiative is developing the next generation of innovating robotic vision experts who will forge whole new industries by converting their research into new products, services and enterprises," Professor Sharma said.

The Centre, of course, brings together experts from QUT, The University of Adelaide, Australian National University, Monash University, as well as DATA61 and partnering international universities.

"Once robots can see and understand the environment they operate in they can make decisions that allow them to work safely beside us," said Centre Director Peter Corke.

"Current industrial robots are dangerous for humans to work around because they're simply not equipped to recognise and avoid obstacles in their way."

To capture the atmosphere and vibe of the event the Centre engaged Matuese Pingol, a QUT digital ambassador, and Juxi Leitner, one of the Centre's own researchers to create a vox pop of the evening.

Matuese and Juxi conducted impromptu interviews with the guests to get their views and opinions on the Centre and the future development of "creating robots that see".

You can see their end result here: https://www.youtube.com/watch?v=tBfBA1Nohp0 .

Robot Press Conference

Talk about a coming out party! A crush of media outlets turned out at the Cube at QUT to see the stars of a new venture. That venture – The Australian Centre for Robotic Vision.

They all jockeyed for the best position to grab shots of Baxter, Casper, the Freeari (the car owned by Frosty the Nao robot – our own Instagram superstar), and, oh, also our own researchers.

The news conference on the morning of 9 March, 2105, preceded the official launch of the Centre that evening. The news conference (and launch) received over 118 media hits, one of the highest rated stories to come out of QUT.

Centre Director Peter Corke and Professor Gordon Wyeth, Executive Dean of QUT's Science and Engineering Faculty, spoke about our new Centre, and about the challenges of giving robots the ability to see.

Peter talked about the environments where robots currently are found, in places like car or computer factories where everything they do is very precise and organised.

"The challenge really for robots," said Peter, "is how do we get them to move out of the those very highly structured and organised work environments that they have in factories? The world we live in is not that at all."

Baxter showed off what it had learned so far, by playing a game of ConnectFour with Peter. Baxter also practiced picking artificial capsicums.

The Nao robots showed off by driving a pint-sized car around without human help, and by taking their pet humans for a walk.

The media got a chance to see Agbot II, a robot designed to seed, weed and fertilise farming land automatically. Agbot II was designed in-house at the Centre.

Also on display, was the Guiabot, which can automatically move around a room, map its environment, and avoid obstacles in its path.

For anyone feeling a bit uneasy about the future featuring the type of robots that have been portrayed in a lot of Hollywood movies, Peter said that robots with a temperament or personality are most likely a decade away.

But Peter did say that roboticists dream of creating a machine that can communicate and learn.

"That's the light on the hill," he said. "That's what gets me out of bed in the morning."

Robots on TV

They've become media darlings. The robots from our Centre certainly have been getting a lot of TV exposure over the past year.

The robots and our researchers have been the focus of two different SCOPE TV episodes on Channel 11, as well as on the station's "Totally Wild" show. On Saturday 6 June, 2015, three of the Centre's robots and researchers were featured in a SCOPE TV episode titled, "Hi Tech Science." Centre Director Peter Corke showed off Baxter, and talked about why robotic vision is important if robots are going to continue to develop.

"If we want robots to be smart and to be able to do all sorts of things that human beings do," said Peter, "then clearly, they need to be able to see as well."

QUT PhD student Gavin Suddrey showed off some of the things that Nao robots can do, including some of their signature dance moves as well as driving the Centre's "Freeari" around.

The show also featured Centre Associate Investigator Tristan Perez and his project, AgBot II.

The purpose of AgBot II is to help farmers with big, painstaking jobs like weeding and fertilising. Something it can't do without vision.

The Agbot gets its vision from a combination of cameras and software that has been developed by the Centre.

"We will not only be able to identify different weeds," Tristan said, "but we will be also be able to determine the best way to get rid of them."

AgBot II is also equipped with its own weather sensors.

"This is to make sure that weather conditions won't move the spray from the weeds to the crops and we can calculate exactly where the chemicals will land," said Tristan. "We're also hoping to program the Agbot to release fertiliser when it's needed and to gather information in the early detection of pests and diseases"

A few weeks later, Channel 11's "Totally Wild" featured Baxter and Research Fellow Chris Lehnert.

Baxter played – and won – a game of Connect 4 against reporter, Kellyn Morris.

Chris then explained to her how Baxter uses a Kinect sensor on its head, and cameras on each wrist to see. Chris also talked about how Baxter is human-safe and that it can detect if it's bumped into a person, or if a person bumps into it.

Gavin Suddrey also showed off the Nao robots to Kellyn, with one of the robots taking her for a walk.

Finally, just a couple weeks after that, SCOPE TV featured Baxter and Chris Lehnert in its episode titled, "Science in Motion."

Chris talked about how the Kinect sensor on Baxter's head helps give Baxter the vision it needs to do a task, like picking capsicums (see Harvey p. 30).

"We're using the Kinect sensor to use the colour information to detect where the capsicum is and use the depth information to find the location of the capsicum."

"So my research will be able to help farmers in the future to autonomously pick fruit which is great news for everyone," said Chris.

No doubt, as the robots continue to improve their vision and the tasks they can perform, they will continue to attract media attention.

http://tenplay.com.au/channel-eleven/scope/episode-details/science-in-motion

http://tenplay.com.au/channel-eleven/scope/episode-details/hi-tech-science

https://www.youtube.com/watch?v=0l5rktnzld8

http://www.abc.net.au/news/2015-08-18/robotronica-natural-language-programming-next-step-home-robots/6686974

Outreach

During 2015 we delivered a range of ad hoc public presentations, demonstrations and school visits promoting robotic vision at different locations. Robotronica, a festival celebrating all things robotic, was held on 23rd August at QUT in Brisbane attracting more than 15,000 people. Our work in robotics and vision was showcased at Robotronica, a public event held as part of National Science Week. This garnered a large volume of media around our robotic and vision activities (see Robotronica p. 14). We also enjoyed a range of visits from schools and the broader community and have enjoyed media coverage in print, on radio and TV including the TV shows Scope, Totally Wild and Insight (see Robots on TV p. 67).

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 Australia at the graduate level to connect the largely disjoint fields of robotics and computer vision and undertake hands-on project work. The first of these Robotic Vision Summer Schools (RVSS 2015) was held at ANU's Kialoa campus in March 2015 (see Summer School p. 69).

In 2015 we partnered with CoderDojo in Queensland to supply a robot, "Franko" the Nao robot, to encourage children to code. We delivered world first Massively Open Online Courses (MOOCs) in "Introduction to Robotics" and "Robotic Vision". MOOCs, free, online courses are open to everyone and offer access to world-class tertiary education and provide a platform for interactive user forums and communities. CI Peter Corke developed the two 6-week MOOCs, which were run globally in 2015. Each course was run twice, with a total of over 40k enrolments with students from over 150 countries. "Introduction to Robotics" covers the fundamental principles of robotics with a focus on robotic manipulator arms, while "Robotic Vision" is a more advanced course. The courses were developed at QUT with the support of Mathworks and Springer. The courses will be run again in 2016. More details, and instructions to enrol, can be found at https://www.qut.edu.au/study/open-online-learning

In 2016 we will launch our professional development program for teachers in robotics and coding in Queensland, the ACT, Victoria and South Australia. We will present robotic vision displays during National Science Week, supply Nao robots for demonstrations, as well as hosting visits and tours of our refurbished Centre facilities at each node to interested parties.

Robotic Vision Summer School (RVSS)

How do you attract the next generation of roboticists? It's a challenge we decided to tackle by holding our first ever Robotic Vision Summer School (RVSS).

We held it in March 2015 in New South Wales, in Kioloa, at the Australian National University's Coastal Campus. The Centre invited our early stage PhDs to attend, and we invited 4th year university students who are considering our PhD program.

RVSS was structured around a series of high profile talks from:

- Peter Corke, Centre Director, QUT

- Andy Davison, Partner Investigator from Imperial College (UK)

- Racquel Urtasan, the University of Toronto

- Lourdes Agapito, University College London (UCL)

- Tom Drummond, Chief Investigator, Monash University

- Rob Mahony, Chief Investigator, Australian National University (ANU)

- Richard Hartley, Chief Investigator, Australian National University (ANU)

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 and OpenCV 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.

We want to give a special thanks to Centre Chief Investigators Stephen Gould and Hongdong Li for pulling the Summer School together.

We plan to hold the school annually at the same venue.

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 2015 Centre Partner Investigators Andrew Davison (Imperial College London) and François Chaumette (INRIA, France) visited our nodes in Australia and QUT received a visit from University of California Berkeley's Pieter Abbeel. 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. 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.

CI Michael Milford did a short 2.5 month sabbatical in the UK, which included a period of time working with partner investigator Andy Davison,

resulting in joint RSS and ICRA workshop papers. This research was on performing place recognition using the relatively new sensing modality of event cameras. Michael gave invited presentations at a number of leading corporations and universities including Google Deepmind, Microsoft Research Seattle, Edinburgh University, St Andrews University, Heriot Watt University, Imperial College London, University of Leeds, University of Reading, Cambridge University, University of Lincoln and University College London. Centre Director, Peter Corke, spent six months at Centre partner Oxford University and also visited the Edinburgh Centre for Robotics. Centre CI, Richard Hartley spent several months at our partner organisation ETH Zurich in 2015. Centre Research Fellow Niko Sünderhauf spent time overseas collaborating with our partner investigators at ETH (Marc Pollefeys) and with colleagues at Chemnitz University of Technology (Peter Protzel). Centre PhD student Zetao Chen has recently spent several months in Adelaide collaborating with researchers there with the intent to submit a joint publication as well.

As part of creating the field of robotic vision and building our profile,

Centre researchers had a number of robotic vision workshops accepted for presentation at the top robotics and computer vision conferences in 2015. In May, Michael Milford, Niko Sünderhauf and Peter Corke ran Place Recognition workshops at CVPR2015 (Computer Vision and Pattern Recognition), ICRA2015 (International Conference on Robotics & Automation), and ICVS2015 (International Conference on Computer Vision Systems), the premiere large robotics and computer vision conferences. CIs Tom Drummond and Peter Corke ran a workshop at ICRA on Challenges to Robotic Vision. CIs Ian Reid and Stephen Gould conducted a workshop on Semantics for Structure from Motion and Visual SLAM at CVPR. CI Ian Reid worked with Centre Research Fellows Cesar Cadena and Yasir Latif organised a workshop on mobile sensors and SLAM at RSS 2015 (Robotics Science and Systems) and Centre Research Fellows Sareh Shirazi, Juxi Leitner and Anoop Cherian presented a workshop on Deep Learning in Vision and Robotics at ACRA2015 (Australasian Conference on Robotics and Automation).

Guest speakers at some of the robotic vision workshops included: David Cox from Harvard University; Centre PI, Paul Newman from Oxford University; Silvio Savarese from Stanford University; John Leonard from MIT; Luca Carlone from our partner organisation Georgia Tech; Jose Neira from University of Zaragoza; and Henry Carillo from Colombia's University Javeriana.

Centre research fellow, Don Dansereau, was invited to present his recently published ACM Transactions on Graphics paper at SIGGRAPH
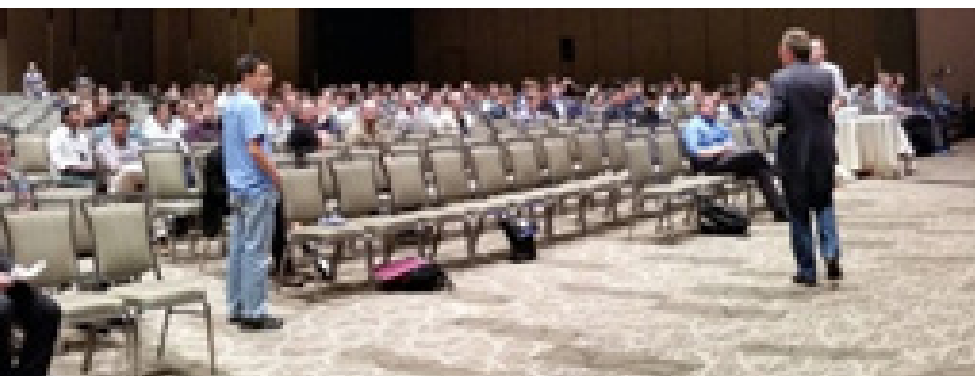
in Los Angeles (August, 2015). SIGGRAPH is the world's premier conference on computer graphics and interactive techniques, attracting tens of thousands of film makers, game developers, students and research scientists from over 70 countries.

We will continue to promote robotic vision in major international conferences and conventions in 2016 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 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 is a premier international forum for robotics researchers to present their work. The conference joins experts in the field of robotics and automation 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.

Centre CIs Stephen Gould (left) and Anton van den Hengel (right)

International and National Links & Networks KPIs

International, National, and Regional Links and Networks

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015
Number of international visitors and visiting fellows	Annually	2	8	8	11
Number of national and international workshops held/organised by the Centre	Annually	1	1	2	15
Number of visits to overseas laboratories and facilities	Annually	2	11	8	25
Examples of relevant interdisciplinary research supported by the Centre *Number of interdisciplinary projects within the Centre	Annually	0	0	1	3
Number of talks given by Centre staff open to the public	Annually	5	10	10	12

People

As we build our Centre we are also all developing a Centre culture, with a set of pervasive values, beliefs and attitudes that characterise who we are.

CENTRE CULTURE/ ONBOARDING STRATEGY

As people join the Centre, we introduce the Centre's Vision and Values, our history, and our plans for the future. We aspire to be a nimble enterprise, with the flexibility and capacity to handle diversity and a range of new ideas.

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. 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 23 PhD researchers bringing our total PhD cohort numbers to 29. Three new Associate Investigators joined us in 2015 including environmental roboticist Dr Matt Dunbabin (see COTSbot p. 54), orthopaedic surgeon and medical robotic enthusiast Prof Ross Crawford (see Medical robotics p.58) and mobile computer vision specialist Dr Laurent Kneip.

Mix of Staff

Position	#	FTE	Gender Ratio (female:male)
Chief Investigator	13	3.9	0 : 13
Partner Investigator	6	0.3	0 : 6
Associate Investigator	17	-	0 : 17
Professional Staff	8	5.6	6 : 2
Research Fellows	17	16	3 : 14
PhD researchers	28	28	3 : 25

Internal Engagement

To truly operate as a Centre, rather than as a collection of individuals, requires us to focus on internal engagement of all of the Centre's researchers and operations staff. We need Centre researchers and staff who understand our strategy and how their work contributes to the performance of the Centre. Research suggests that two-thirds of all employees are 33% as productive as they could be because they don't understand what they are being asked to do. People who are engaged have significantly higher productivity, profitability, and customer ratings, less turnover and absenteeism, and fewer safety incidents (Gallup Report, State of the Global Workplace, 2014).

We promote internal engagement by a number of means. We have adopted a wiki-style intranet tool called Confluence, developed by Australian company, Atlassian. Our Centre has an active internal Google+ community (robotvision) with 69 members, as well as a public-facing identity on Google+, Facebook, Twitter, LinkedIn, and Flickr. Our Centre website, internal Confluence wiki internet and social media channels are updated with new content at least once a week. The Centre's Executive meets fortnightly (moving to weekly in 2016) via Citrix Go-To-Meetings. Research Fellows and PhD researchers each have monthly virtual meetings. Centre Director, Peter Corke keeps in regular with CIs via global emails and to promote Centre achievements we produce an online newsletter for internal stakeholders. Regular communication from the Centre Director also kept everyone up-to-date on the recruitment and arrival of new people to the Centre and their specialist skills, as well as announcing good news.

To be effective we need to be cohesive and this requires conversation and collaboration, and travel is an important enabler of these given that we are spread over four cities. We encourage and support travel for chief investigators, research fellows, and PhD researchers to other Centre nodes, including overseas partner universities, and to host visits by researchers from our overseas partners (see Node Visits below). Within Australia we hold an annual symposium, a 3 day event, that includes presentations from Centre researchers, invited guests and showcases demonstrators of current research work (see RoboVis p.74).

Node Visits

In early February 2015 a group of Centre researchers from the QUT node travelled to Adelaide and spent 6 days at the University of Adelaide. The purpose of this trip was to discuss future research collaboration avenues between our nodes. It was a successful trip that established the foundations for short-term and long-term collaboration goals. In a series of meetings, various research avenues at the intersection of computer vision and robotics were explored. Special interest groups brainstormed potential demonstrators for RoboVis 2015. This first visit led to a better understanding of the specific expertise available within our Centre and paves the way for defining future common research directions.

RoboVis2015

One of the main goals of any Australian Research Council (ARC) Centre of Excellence is collaboration between researchers at different universities. It's a goal our Centre takes very seriously.

To help meet that goal, the Centre holds what we call "RoboVis," an annual meeting for 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.

In June 2015, 80 people turned out for "RoboVis" in the spectacular Barossa Valley in South Australia. The two and a half day program consisted of a training stream and a scientific stream.

The training stream was run as parallel sessions for research fellows and PhD researchers. It included sessions on developing careers, managing people, engaging industry, writing papers and giving talks. The group heard from Robert Chalmers, the Managing Director for Adelaide Research & Innovation, on "Intellectual Property and Commercialisation", and also heard from CI Stephen Gould on his experiences in "Start-ups and Academia". All presentations were very popular.

The scientific programs delivered were an opportunity for our researchers to talk about and share the exciting work being conducted in the Centre. This part of the program also included great demonstrations and three-minute spotlight talks.

One of the highlights of the program was the keynote address by our Centre Advisory Committee member Mandyam Srinivasan who talked about his work with the birds and the bees. No, it's not what you think. Instead, his talk provided fascinating insight into how bees and birds navigate their way through the environment while in flight.

RoboVis will continue to be an annual showcase and celebration of the Centre's research with RoboVis2016 scheduled to be hosted by our Centre's Monash University node somewhere in Victoria in September 2016.

Research Training

Twenty-three new postgraduate researchers started with our Centre in 2015 as well as nine additional research fellows. To support the development of our early career researchers, a range of research training opportunities are being provided to both our research fellows and PhD researchers. These courses have been a mixture of face-to-face and online opportunities, supplemented by a Research Training Toolbox hosted on our wiki intranet, Confluence.

In 2015, Peter Corke's Massively Open Online Courses commenced with an Introduction to Robotics and Robotic Vision, and was used by commencing PhD researchers at QUT. In March we ran our first Robotic Vision Summer School (RVSS), organised by Stephen Gould and held at ANU's coastal campus at Kioloa on the South Coast of NSW. Although the primary focus of the summer school is to recruit PhD researchers to the Centre it is also an opportunity for new PhD researchers and research fellows to meet and learn from the range of invited experts talking about all things robotic vision. The Centre was lucky to host Partner Investigator Andy Davison from Imperial College London as well as three other international guest speakers, Lourdes Apapito (University College London), Fredrik Kahl (Chalmers University) and Raquel Urtasun (University of Toronto) (see Summer School p. 69)

At our annual symposium, RoboVis, a day was set aside exclusively for research training for PhD researchers and Research Fellows. Both cohorts were asked their preferences regarding topics and there were many overlapping interests. Courses were presented in 30 minute time slots by the Centre's CIs and AIs as well as the COO and a guest speaker, Robert Chalmers, Managing Director of Adelaide Research and Innovation. Topics covered included:

- Writing papers and giving talks
- Managing Code
- Work/Life Balance
- Managing People
- Developing Your Career
- Intellectual Property and Commercialisation
- Start-ups and Academia
- Topics specifically for PhD researchers:
- Writing up a Thesis
- Managing Your PhD Supervisor
- Topics for Research Fellows:
- Engaging with Industry
- Supervising Undergraduate and PhD projects

In the second half of 2015 the Centre also delivered two courses, internally via Citrix Go-To-Meeting, on the subjects:

- Creating a start-up
- Careers

These additional courses were delivered by COO Sue Keay, the first involved 10 meetings tackling the Lean Launchpad method using Steve Blank's Udacity MOOC with Eric Ries book on The Lean Start-Up used as reference material. The second course involved eight meetings and workshops around different stages of career planning, underpinned by Peter Fiske's Career Guide for Scientists (and Engineers) and involving guest speakers, Howard Leemon (UniQuest) and Surya Singh (UQ).

Courses to be offered in 2016 include Project Management and Research Leadership. A combination of internal delivery and external consultants will be used to deliver content, potentially including sessions at the next RoboVis to be held in Victoria in September 2016. The Start-up and Careers courses will also be run as half day workshops at each node as opportunity permits.

In addition to the training activities mentioned, we have been hosting monthly PhD researcher and Research Fellow meetings via Citrix Go-To-Meeting (recorded on Confluence) to encourage internode collaboration.

Research Training and Professional Development KPIs

Research Training and Professional Education

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015
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**
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
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
Number of new postdoctoral researchers recruited to the Centre working on core Centre research	Annually	7	6	9	11
Number of new Honours students working on core Centre research and supervised by Centre staff	Annually	4	10	8	7
Number of postgraduate completions and completion times, by students working on core Centre research and supervised by Centre staff	Annually	0	0	0	0
Number of Early Career Researchers (within five years of completing PhD) working on core Centre research	Annually	7	6	16	14
Number of students mentored	Annually	45	6	90	29
Number of mentoring programs offered by the Centre (include programs for students, new staff, external stakeholders and clients)	Annually	0	1	2	2

Gender Diversity

Our Centre 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 such as robotics and computer vision to attract more women into IT and engineering at the undergraduate level. As more women join the Centre, a Centre-wide female support network will be established based on the group that meets every 6

weeks at our Centre's QUT node, which also hosted a meeting with MIT's Professor Daniela Rus to talk about challenges for female researchers in science and engineering. In 2015 Centre COO Sue Keay was invited to present at Geoscience Australia as part of their Women in Leadership series and also chaired a panel discussion on Females in STEM held by QUT's Science and Engineering Faculty and featuring panellists Abigail Allwood (NASA scientist) and Anne-Marie Birkill (OneVentures Innovation fund). In 2016 we will continue to explore opportunities to encourage diversity.

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.

Each Theme Leader, and the Director, has a nominated deputy who fulfils their role during short-term absences (see Governance p. 88).

Research Fellow Donald Dansereau (centre) demonstrates the operation of a light field camera to Centre PhD researchers

Commercialisation and Entrepreneurship

Protecting and developing Centre Intellectual Property (IP) is a fundamental objective of the Centre. All people 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 will consult with the Centre Executive over protection of Project Intellectual Property.

The Centre Executive is responsible for deciding whether any item of Project Intellectual Property merits protection or commercialisation. Project parties then work with the Centre Director to determine a Commercialisation Lead and follow the conditions stipulated in the Centre Agreement to commercialise the Project Intellectual Property.

Centre researchers consult with the Executive before open sourcing any Centre IP. If the Executive supports open sourcing then they will recommend a license suitable for wide-ranging application. If the Executive decides the IP might have commercial potential then they will recommend another type of license and/or IP protection strategy.

Centre researchers are encouraged to learn about opportunities for commercialisation and to develop their entrepreneurial skills. We firmly believe that our early career researchers, as well as making substantial contributions to research and to any companies they may enter, will be well-placed to start their own companies and help develop the fledgling robotic vision industry in Australia.

For example, two of the Centre's researchers have already been developing their entrepreneurial skills. CI Michael Milford has launched an educational start-up called, "Math Thrills" and is seeking investment, while Centre PhD researcher Zongyuan (Tony) Ge spent some time interning with Chinese AI start-up company, "Deepglint". Here we detail their experiences in 2015.

Math Thrills

CI Michael Milford has just founded the company Math Thrills, which produces innovative, fun and exciting fictional mathematical material for students. Math Thrills received Kickstarter funding in 2014, and seed funding from QUTBluebox (QUT's commercialisation arm) in 2015. It has led to several awards including 2015 Queensland Young Tall Poppy of the Year Award and a 2015 TedXQUT talk (https://www.youtube.com/watch?v=m_U7qIjvMJw), as well as media coverage. Its first generation product was a young adult novel stealthily embedded with mathematical concepts ("A Question of Will"), accompanied by extensive explicit mathematical worksheets and animated tutorials. In late 2015 the second generation product, which included live action workshops, went into trials at schools around south east Queensland. In 2016 we are working on further product developments including automated workshops and a YouTube series of maths in fiction videos.

Michael Milford's Maths Thrills educational program is being trialled across Queensland Schools, image courtesy of Michael Milford

Deepglint's offices in Beijing where Centre PhD researcher Zongyuan Ge interned, image courtesy of Zongyuan

Deepglint internship

In 2015, Zongyuan (Tony) Ge spent two months at a really cool start-up artificial intelligence company called Deepglint, Beijing (www.deepglint.com). Tony's main responsibility was to use his deep learning skills to develop a product called DeepV vehicle feature analysis system with Deepglint's Apollo development team. The development team has an all-star colleague list including people from Yoshua Bengio's group, Feifei Li's group, Google X etc. During Tony's internship, the knowledge and theory he developed while a PhD candidate with the Centre served as a strong foundation of creativity on product design. The experience he gained, the system design, coding, product testing skills, helped him to complete his research arsenal.

Tony feels it is important to mention the side benefits of his experience; getting to know young and senior researchers, access to high quality industry seminars every week (meeting people from Baidu, Amazon, Linkedin), nice food (Deepglint built a high standard kitchen and all staff are from 5-star hotels), a beautiful garden to walk around and a shining line on his resume. Tony strongly recommends all PhD students from our Centre spend some time in these rising star companies who are very keen to recruit our graduates.

Research Fellow Vijay Kumar (left) explains his research to Centre researchers at our annual symposium RoboVis

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

Peter Corke

Ian Reid

Tom Drummond

Rob Mahony
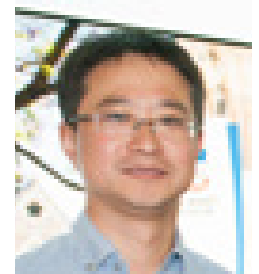
Hongdong Li

Michael Milford

Chunhua Shen

Gustavo Carneiro
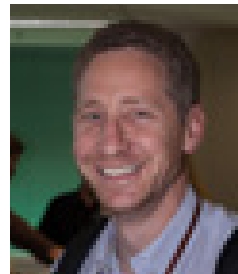
Stephen Gould
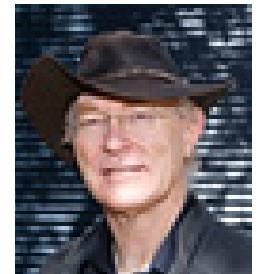
Richard Hartley

Ben Upcroft
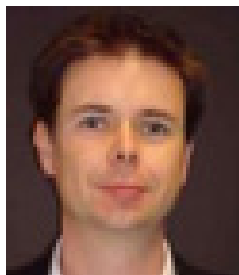
Anton van den Hengel

Gordon Wyeth

Associate Investigators

Tat-Jun Chin

Ross Crawford

Anthony Dick

Matt Dunbabin

Anders Eriksson

Clinton Fookes

Jason Ford

Jonghyuk Kim

Luis Mejias Alvarez

Laurent Kneip

Tristan Perez
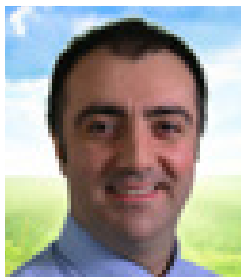
Fatih Porikli

Jonathan Roberts

Ahmet Sekercioglu

Qinfeng (Javen) Shi

David Suter

Jochen Trumpf

Partner Investigators

François Chaumette

Andrew Davison

Frank Dellaert

Paul Newman

Phillip Torr

Marc Pollefeys

Research Fellows

Khurrum Aftab

Cesar Cadena

Anoop Cherian
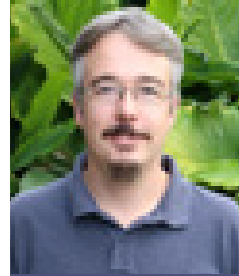
Donald Dansereau

Markus Eich

Basura Fernando
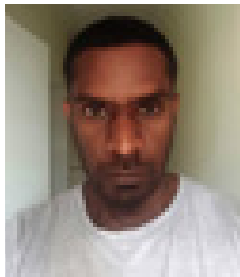
Dinesh Gamage

Viorela Ila

Vijay Kumar

Yasir Latif
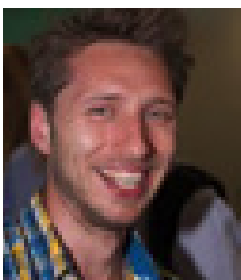
Jürgen "Juxi" Leitner

Guosheng Lin
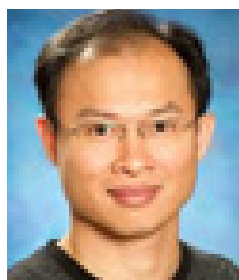
Chuong Nguyen

Trung Than Pham

Sareh Shirazi

Niko Sünderhauf
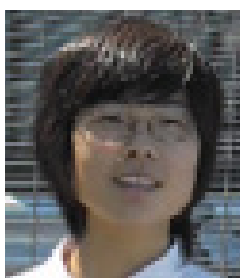
Lin Wu

PhD Researchers

Peter Anderson

Juan Adarve

William Chamberlain

Zetao "Jason" Chen
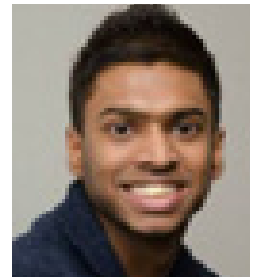
Thanuja Dharmasiri
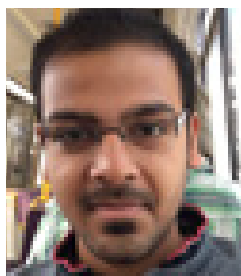
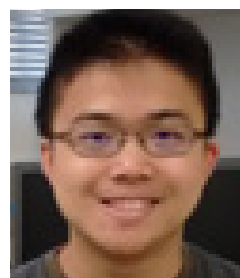Sourav Garg

Zongyuan Ge

Jian "Edison" Guo
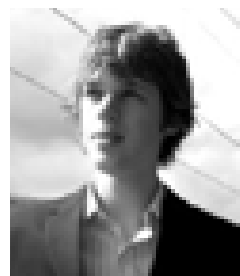
Ben Harwood
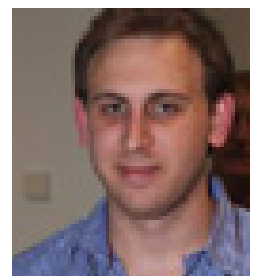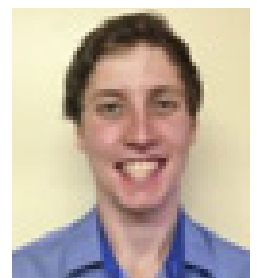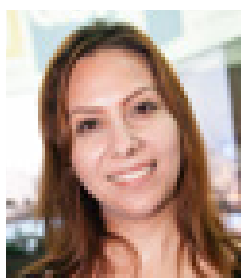
Adam Jacobson

Yuchao Jiang
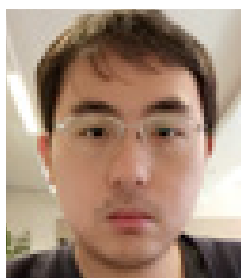
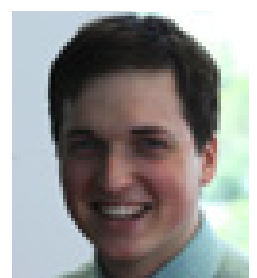Zhibin Liao

Hui Li

Vincent Lui

Sean McMahon

Ben Meyer

Fahimeh Rezazedegan

Andrew Spek

Rodrigo Santa Cruz

Adam Tow

PhD researchers (cont.)

Andres Felipe Marmol Velez
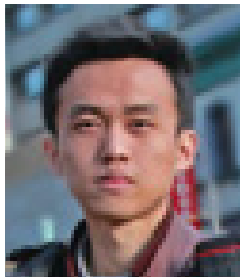
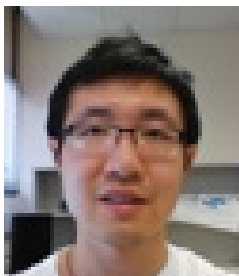**Xiaoqin Wang**

Bohan Zhang

Fangyi Zhang

Yi "Joey" Zhou

Yan Zuo

PhD researcher Josh Weberruss is not represented in these images

PhDs starting in 2016

Dan Richards

James Sergeant
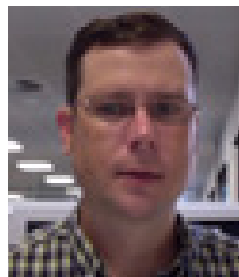
Peter Kujala
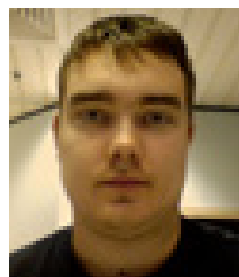
Jean-Luc Stevens
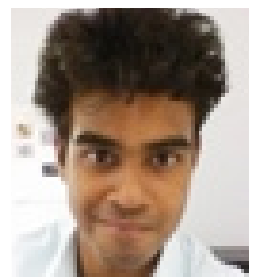
Jeffrey Devaraj

John Skinner

Mina Henein
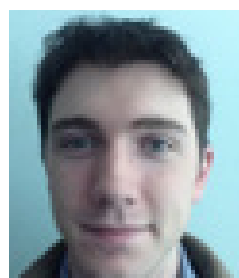
Sean O'Brien

Tong Shen

Research Affiliates

Felipe Gonzalez

Mark McDonnell
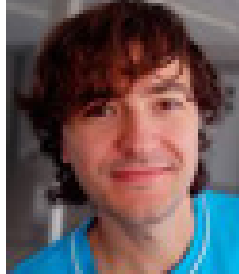
Anton Milan

Thierry Peynot

Professional Staff

Kate Aldridge

Sarah Allen

Sue Keay

Tracy Kelly

Greg Lee

Thuy Mai

Sandra Pederson and Alex Martin are professional staff members with the Centre but are not represented in these images

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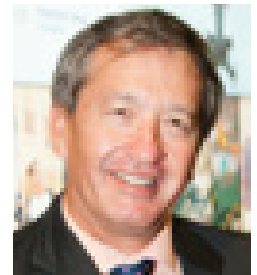 Australian Research Council, a statutory agency responsible for Australia's National Competitive Grants Program, which contributes $19m in public funding to our Centre. The Executive also represents our four domestic research partners; Queensland University of Technology (QUT), The University of Adelaide (Adelaide), The Australian National University (ANU), and Monash University.

We have held 40 Centre Executive meetings and four meetings of all Chief Investigators to the end of 2015. In mid-2015 we instituted quarterly
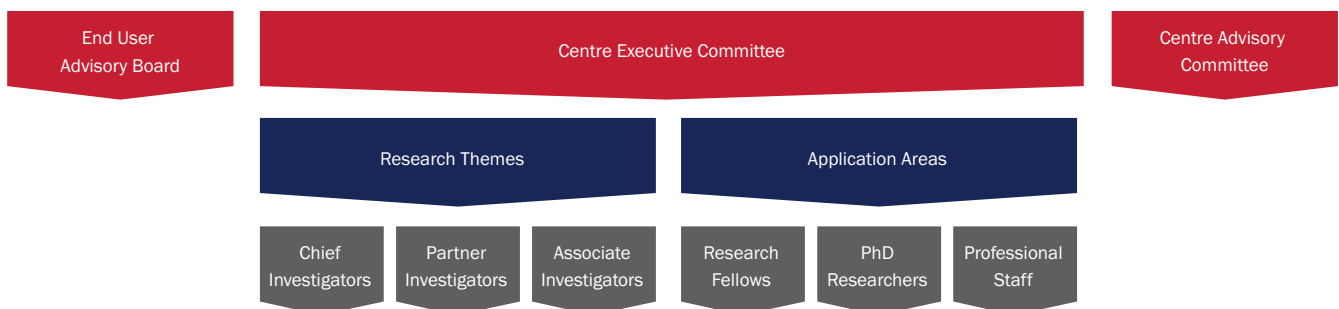
face-to-face (F2F) meetings of the Executive, realising that there is still no substitute for direct contact. The first of these was held in Melbourne on 5-6th August, and the second in Adelaide on 12th November. Representatives of the Centre (Centre Acting Director, Ian Reid and COO Sue Keay) visited the ARC in Canberra in July to report on the Centre's progress in person.

The Centre holds monthly research fellow and PhD researcher meetings and publishes a research training toolbox on our intranet, and has developed research training and mentoring programs for all early career researchers (research fellows and PhD researchers). Monthly "Theme meetings" in Robust Vision and Semantic Vision bring together

all researchers in a theme via video-conference. Vision and Action held a F2F workshop meeting after the Australasian Conference on Robotics and Automation in Canberra, a local robotics conference, in December 2015.

In 2016 we aim to host a visit from the Australian Research Council, to hold two Centre Advisory Committee meetings, to hold regular Centre Executive meetings with quarterly F2F meetings, two all CI meetings, to initiate monthly meetings of our newly formed research committee, and to conduct one mock Centre mid-term review.

OUR STRUCTURE

End User Advisory Board	Centre Executive Committee			Centre Advisory Committee

Research Themes			Application Areas		
Chief Investigators	Partner Investigators	Associate Investigators	Research Fellows	PhD Researchers	Professional Staff

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.

The first meeting of our Centre Advisory Committee, chaired by Alex Zelinsky, took place via teleconference on 17th August 2015.

CENTRE EXECUTIVE COMMITTEE (CEC)

The purpose of our Centre Executive Committee is to ensure that the Centre meets its agreed goals. The committee comprises our Centre Director, three CIs representing Monash, University of Adelaide and ANU, and the Chief Operating Officer. The CEC governs the Centre. The Committee meets fortnightly via Citrix Go-To-Meeting: 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; aquaculture & bio-monitoring, 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.

We convened our Centre's end-user advisory board (EUAB), chaired by Russel Rankin, in Adelaide on 12th November 2015. Our EUAB has requested more regular (and longer) meetings, at least twice per year. The 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.
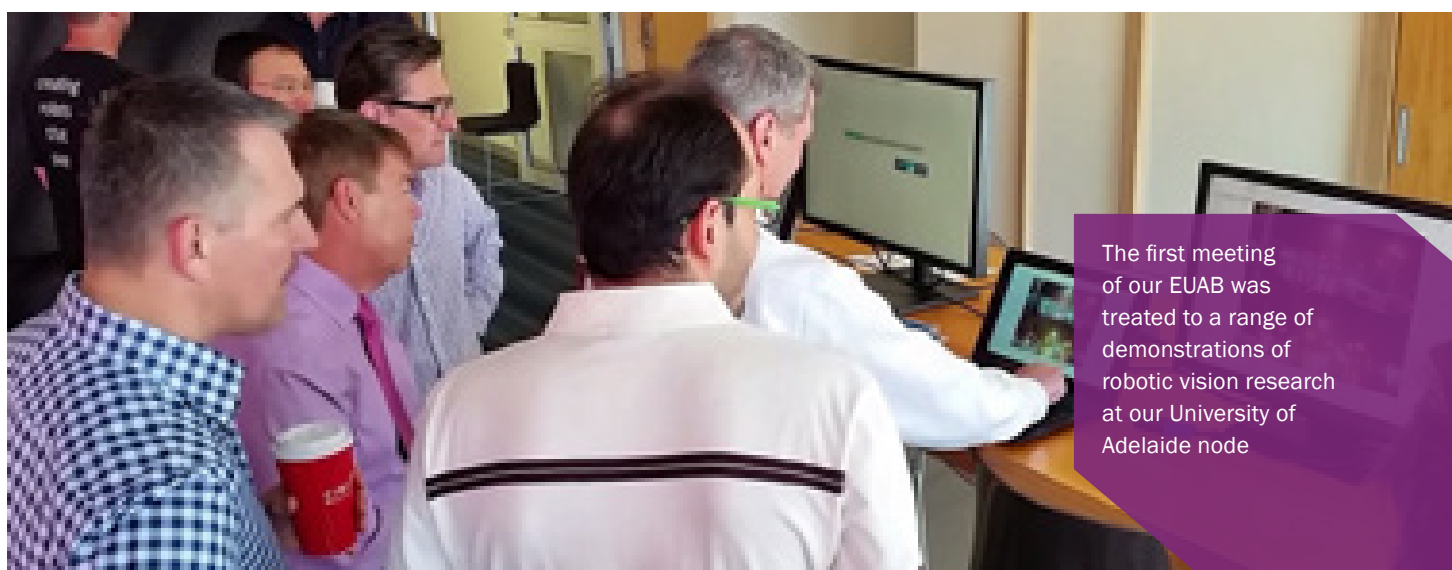
The first meeting of our EUAB was treated to a range of demonstrations of robotic vision research at our University of Adelaide node

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.
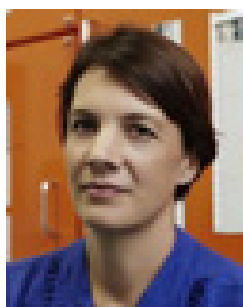
Scientia Professor Michelle Simmons

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. She has recently become Editor-in-Chief of Nature Quantum Information and in 2015 was awarded the CSIRO Eureka Prize for Leadership in Science.

Professor Mandyam Srinivasan

Srini is presently Professor of Visual Neuroscience at the Queensland Brain Institute and the School of Information Technology and Electrical Engineering at The University of Queensland. He 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, 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, and the Membership of the Order of Australia (AM) in 2012. With a research focus on bees, Srini has explored how simple animal systems display complex behaviours. This broad field has applications in robotics, especially unmanned aerial vehicles because of the competing needs for autonomy and a lightweight control system.

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 Sir Mike 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.

Brisbane City Council Lord Mayor Graham Quirk at a coding lesson from Coder Dojo with the Centre's donated NAO robot, Franko

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 connect companies, research organisations, financial institutions and Government and understands how to manage the innovation process and the steps required to commercialise innovative ideas.

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 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) 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 Harris is a Professorial Fellow of Chemical and Biomolecular Engineering at The University of Sydney, and the Australian director of Laing O'Rourke's 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 billion per annum.

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.

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 glob - including in the Asia-Pacific, United Kingdom and the Middle East. He has worked in business consultancy and new ventures always where innovation is leveraged to identify new sources of value.

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.

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.

Governance KPIs

Governance

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015
Breadth, balance and experience of the members of the Advisory committee	At review				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
Vision and usefulness of the Centre strategic plan	At review	Annual Review by Board			at review
The adequacy of the Centre's performance measure targets	At review	As above			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
▪ Number of new joint research projects		0	0	1	1
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 fellowship applications PhD students will be developed through supervisor mentoring, a Centre-run mentoring program, and opportunities for supervision of undergraduate students.			

Finance

Our Financial Statement (p. 95) provides a summary of our financial performance for the 2015 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 as in-kind contributions, mainly consisting of researcher time. As the Centre progresses, additional funding will also be sourced via industry engagement and relevant industry projects.

Statement of Operating Income and Expenditure for year end 31 December 2015

Income	2014	2015
ARC Centre Grant distributed as follows	2,714,290	2,714,284
Monash University	116,929	116,929
University of Adelaide	244,400	244,400
Australian National University	230,000	230,000
Queensland University of Technology	360,000	350,000
ARC Indexation	82,027	132,110
Total Income	3,747,646	3,787,723

Expenditure	2014	2015
Purchased Equipment	42,256	51,829
Shared Equipment/Facilities	-	815
Travel and Professional Development	24,276	248,352
Maintenance (IT and lab)	8,350	19,639
Salaries/Personnel exp	373,061	2,337,787
Other	38,002	185,704
Total Expenditure	485,946	2,844,126

	2014	2015
Surplus/Deficit	3,261,700	943,597
Previous year carry forward	-	3,261,700
Total carry forward surplus*	3,261,700	4,205,297

*Surplus from 2014

** Based on combined data from all nodes. Each node's data is compliant with it's own policy and procedures.

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 2014 as ARC funds were released to the Centre more than six months before legal agreements governing the operation of the Centre could be put in place. No recruitment could be undertaken until these agreements were executed and so expenditure was below target. This carry-forward has continued through 2015 and has increased by just under $1m. The carry-forward reflects the ongoing recruitment of personnel until the Centre reached full capacity in mid-2015. 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 2016 and beyond (where "new" income is the ~$3.7m cash contracted to be received by the Centre each year until 2021).

FINANCIAL OUTLOOK

The forecasted cash budget for the Centre for 2016, totalling $3,655,086 (Figure 5).

2015 Expenditure vs Income

$943,597

$2,337,787

$205,343

$248,352

$52,644

- Personnel
- Equipment
- Travel
- Other
- Carry-forward
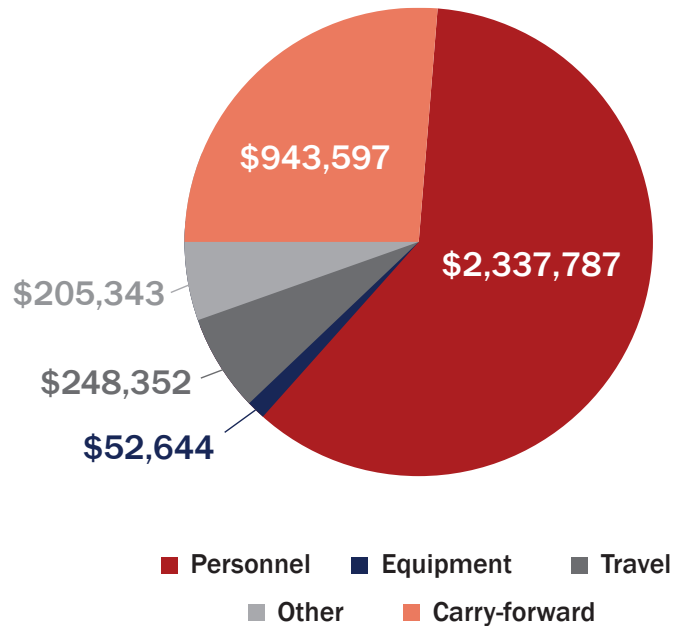
Figure 4: Allocation of Centre Resources

2016 Forecast Cash Income

$350,000

$230,000

$244,400

$116,400

$2,714,286

- ARC Centre of Excellence Grant
- Monash University
- University of Adelaide
- Australian National University
- Queensland University of Technology

Figure 5: Anticipated income by source for 2016.

Summary of Contributions from All Parties

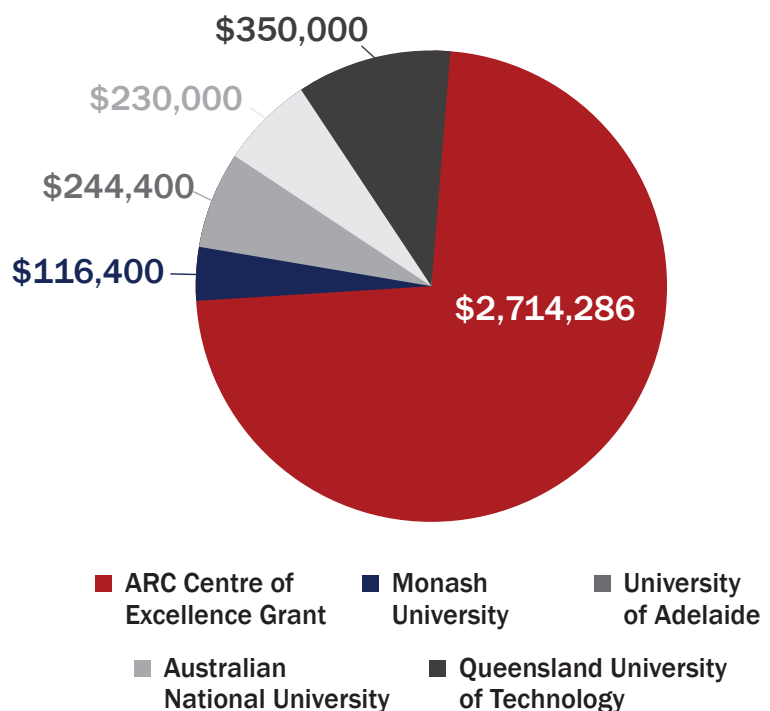The administering and collaborating organisations are contributing $980k per annum in cash, which amounts to $6.86m in cash over the life of the Centre, and nearly $997k per annum in-kind totalling $6.98m over the life of the Centre. Our

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	$2.26m	$1.23m	$1.14m	$224k	$127k	$126k	$101k	$185k	$210k

Finance KPIs

Organisational support

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015
Annual cash contributions from Administering and Collaborating Organisations	Annually				
▪ QUT		$350,000	$350,000	$350,000	$350,000
▪ Monash		$116,400	$116,400	$116,400	$116,400
▪ ANU		$230,000	$230,000	$230,000	$230,000
▪ University of Adelaide		$244,400	$244,400	$244,400	$244,400
Annual in-kind contributions from Administering and Collaborating Organisations	Annually				
▪ QUT		$309K	$309K	$309K	$482K
▪ Monash		$151K	$151K	$151K	$150K
▪ ANU		$332K	$114K	$332K	$342K
▪ University of Adelaide		$153K	$200K	$153K	$205K

Finance KPIs cont.

Organisational support

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015
Annual cash contributions from Partner Organisations*list each Organisation separately	Annually	Nil	Nil	Nil	Nil
Annual in-kind contributions from partner Organisations	Annually				
▪ Georgia Tech		$18K	$3.6K	$18K	$18K
▪ INRIA		$14.5K	$0K	$14.5K	$14.5K
▪ Imperial College		$18.1K	$3.6K	$18.1K	$18.1K
▪ NICTA		$30K	$5.8K	$30K	$30K
▪ Swiss Federal Institute		$26.4K	$0K	$26.4K	$26.4K
▪ Oxford		$32K	$3.2K	$32K	$32K
Other research income sourced by Centre *End User (industry, public sector, ARC Linkage and Discovery in non-core areas, CRC	Annually	$0	$1M	$1M	$4M
Number of new organisations collaborating with, or involved in, the Centre	Annually	2	2	2	2
Level and quality of infrastructure provided to the Centre	At review				

Publications

BOOK CHAPTER

Leitner, J. (2015). A Bottom-Up Integration of Vision and Actions To Create Cognitive Humanoids Cognitive Robotics (pp. 191-214): CRC Press.

JOURNAL ARTICLES

Aftab, K., Hartley, R., & Trumpf, J. (2015). Generalized Weiszfeld Algorithms for Lq Optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(4), 728-745.

Andersh, J., Cherian, A., Mettler, B., & Papanikolopoulos, N. (2015). A vision based ensemble approach to velocity estimation for miniature rotorcraft. *Autonomous Robots*, 39(2), 123-138.

Chen, Z., Lowry, S., Jacobson, A., Hasselmo, M. E., & Milford, M. (2015). Bio-inspired homogeneous multi-scale place recognition. *Neural Networks*, 72, 48-61.

Cherian, A., Morellas, V., & Papanikolopoulos, N. (2015). Bayesian Nonparametric Clustering for Positive Definite Matrices. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Dansereau, D. G., Pizarro, O., & Williams, S. B. (2015). Linear Volumetric Focus for Light Field Cameras. *ACM Transactions on Graphics* (TOG), 34(2), 15.

Edussooriya, C. U. S., Dansereau, D. G., Bruton, L. T., & Agathoklis, P. (2015). Five-Dimensional Depth-Velocity Filtering for Enhancing Moving Objects in Light Field Videos. *IEEE Transactions on Signal Processing*, 63(8), 2151-2163.

Fernando, B., Tommasi, T., & Tuytelaars, T. (2015). Joint cross-domain classification and subspace learning for unsupervised adaptation. *Pattern Recognition Letters*, 65, 60-66.

Hui, L., Chunhua, S., van den Hengel, A., & Qinfeng, S. (2015). Worst case linear discriminant analysis as scalable semidefinite feasibility problems. *IEEE Transactions on Image Processing*, 24(8), 2382-2392.

Jacobson, A., Zetao, C., & Milford, M. (2015). Autonomous Multisensor Calibration and Closed-loop Fusion for SLAM. *Journal of Field Robotics*, 32(1), 85-122.

Jingxin, X., Denman, S., Sridharan, S., & Fookes, C. (2015). An Efficient and Robust System for Multiperson Event Detection in Real-World Indoor Surveillance Scenes. *IEEE Transactions on Circuits and Systems for Video Technology*, 25(6), 1063-1076.

Li, X., Shen, C., Dick, A., Zhang, Z., & Zhuang, Y. (2015). Online Metric-Weighted Linear Representations for Robust Visual Tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(5), 931-950.

Liu, F., Lin, G., & Shen, C. (2015). CRF learning with CNN features for image segmentation. *Pattern Recognition*, 48(10), 2983-2992.

Lowry, S., Sunderhauf, N., Newman, P., Leonard, J. J., Cox, D., Corke, P., & Milford, M. J. (2015). Visual Place Recognition: A Survey. *IEEE Transactions on Robotics*, 32(1), 1-19.

Miaomiao, L., Hartley, R., & Salzmann, M. (2015). Mirror Surface Reconstruction from a Single Image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(4), 76-73.

Neubert, P., Sünderhauf, N., & Protzel, P. (2015). Superpixel-based appearance change prediction for long-term navigation across seasons. *Robotics and Autonomous Systems*, 69, 15-27.

Ryan, D., Denman, S., Sridharan, S., & Fookes, C. (2015). An evaluation of crowd counting methods, features and regression models. *Computer Vision and Image Understanding*, 130, 1-17.

Sa, I., Hrabar, S., & Corke, P. (2015). Outdoor flight testing of a pole inspection UAV incorporating high-speed vision. *Field and Service Robotics*, 105, 107-121.

Shi, Q., Reid, M., Caetano, T., Van den Hengel, A., & Wang, Z. (2015). A hybrid loss for multiclass and structured prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(1), 2-12.

Wang, X., Şekercioğlu, Y. A., & Drummond, T. (2015). Vision-Based Cooperative Pose Estimation for Localization in Multi-Robot Systems Equipped with RGB-D Cameras. *Robotics*, 4(1), 1-22.

Wang, X., Sekercioglu, Y. A., Drummond, T., Natalizio, E., Fantoni, I., & Fremont, V. (2015). Fast Depth Video Compression for Mobile RGB-D Sensors. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(4), 673-686.

Warren, M., Mejias, L., Kok, J., Yang, X., Gonzalez, L. F., & Upcroft, B. (2015). An automated emergency landing system for fixed-wing aircraft: Planning and control. *Journal of Field Robotics*.

Zetao, C., Yuen, J., Crawford, R., Jiang, C., Chengtie, W., & Yin, X. (2015). The effect of osteoimmunomodulation on the osteogenic effects of cobalt incorporated -tricalcium phosphate. *Biomaterials*, 61, 126-138.

CONFERENCE PAPERS

Aftab, K., & Hartley, R. (2015). *Convergence of Iteratively Re-weighted Least Squares to Robust M-Estimators*. Paper presented at the IEEE Winter Conference on Applications of Computer Vision (WACV), 2015, Los Alamitos, CA, USA.

Aftab, K., & Hartley, R. (2015). *LQ-bundle adjustment*. Paper presented at the IEEE International Conference on Image Processing (ICIP) 2015.

Ajanthan, T., Hartley, R., Salzmann, M., & Li, H. (2015). *Iteratively Reweighted Graph Cut for Multi-label MRFs with Non-convex Priors*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

Bewley, A., & Upcroft, B. (2015). *From ImageNet to mining: Adapting visual object detection with minimal supervision*. Paper presented at the Field and Service Robotics (FSR) 2015, University of Toronto, Canada.

Buyu, L., Xuming, H., & Gould, S. (2015). *Multi-class Semantic Video Segmentation with Exemplar-Based Object Reasoning*. Paper presented at the IEEE Winter Conference on Applications of Computer Vision (WACV), 2015.

Chamberlain, W., Drummond, T., & Corke, P. (2015). *Distributed Robotic Vision as a Service*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015, Canberra, Australia.

Chen, Z., Lowry, S., Jacobson, A., Ge, Z., & Milford, M. (2015). *Distance Metric Learning for Feature-Agnostic Place Recognition*. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015, Hamburg, Germany.

Chen, Z., Lowry, S., Jacobson, A., Ge, Z., & Milford, M. (2015). *Distance metric learning for feature-agnostic place recognition*. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015.

Cunningham-Nelson, S., Moghadam, P., Roberts, J., & Elfes, A. (2015). *Coverage-Based Next Best View Selection*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015.

Dansereau, D. G., Williams, S. B., & Corke, P. I. (2015). *Closed-Form Change Detection from Moving Light Field Cameras*. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015.

Dayoub, F., Dunbabin, M., & Corke, P. (2015). *Robotic Detection and Tracking of Crown-of-Thorns Starfish*. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015, Hamburg, Germany.

Eich, M. (2015). *Towards Fuzzy Knowledge Based Scene Understanding from RGB-D Images*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

Faraki, M., Harandi, M. T., & Porikli, F. (2015). *Material Classification on Symmetric Positive Definite Manifolds*. Paper presented at the IEEE Winter Conference on Applications of Computer Vision (WACV), 2015, Los Alamitos, CA, USA. http://dx.doi.org/10.1109/WACV.2015.105

Fernando, B., Gavves, E., Muselet, D., & Tuytelaars, T. (2015). *Learning to rank based on subsequences*. Paper presented at the International Conference on Computer Vision (ICCV) 2015.

Fernando, B., Gavves, E., Oramas, J., Ghodrati, A., & Tuytelaars, T. (2015). *Modeling video evolution for action recognition*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015

Gamage, D., & Drummond, T. (2015). *Reduced Dimensionality Extended Kalman Filter for SLAM*. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015.

Ge, Z., McCool, C., Sanderson, C., & Corke, P. (2015). *Subset feature learning for fine-grained category classification*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, Piscataway, NJ, USA.

Hall, D., McCool, C., Dayoub, F., Sunderhauf, N., & Upcroft, B. (2015). *Evaluation of Features for Leaf Classification in Challenging Conditions*. Paper presented at the IEEE Winter Conference on Applications of Computer Vision (WACV), 2015, Los Alamitos, CA, USA.

Hanxi, L., Yi, L., & Porikli, F. (2015). *Robust online visual tracking with a single convolutional neural network*. Paper presented at the Asian Conference on Computer Vision (ACCV) 1-5 Nov. 2014, Cham, Switzerland.

Ila, V., Polok, L., Solony, M., Smrz, P., & Zemcik, P. (2015). *Fast Covariance Recovery in Incremental Nonlinear Least Square Solvers*. Paper presented at the IEEE International Conference on Robotics and Automation (ICRA), 2015, Seattle.

Jacobson, A., Chen, Z., Rallabandi, V. R., & Milford, M. (2015). *Multi-Scale Place Recognition with Multi-Scale*

Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015, Hamburg, Germany.

Sensing. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015.

Jacobson, A., Zetao, C., & Milford, M. (2015). *Online place recognition calibration for out-of-the-box SLAM*. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015, Piscataway, NJ, USA.

Jia, X., Gavves, S., Fernando, B., & Tuytelaars, T. (2015). *Guided long-short term memory for image caption generation*. Paper presented at the International Conference on Computer Vision (ICCV) 2015.

Leitner, J. (2015). *Extending Visual Perception With Haptic Exploration for Improved Scene Understanding*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

Leitner, J., Dansereau, D. G., Shirazi, S., & Corke, P. (2015). *The Need for Dynamic & Active Datasets*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, Massachusetts, USA.

Liao, Z., & Carneiro, G. (2015). *The use of deep learning features in a hierarchical classifier learned with the minimization of a non-greedy loss function that delays gratification*. Paper presented at the IEEE International Conference on Image Processing (ICIP) 2015.

Lin, G., Shen, C., Reid, I., & Hengel, A. v. d. (2015). *Deeply Learning the Messages in Message Passing Inference*. Paper presented at the Neural Information Processing Systems (NIPS), 2015, Montreal, Canada.

Lui, V., & Drummond, T. (2015). *Image based optimisation without global consistency for constant time monocular visual SLAM*. Paper presented at the IEEE International Conference on Robotics and Automation (ICRA), 2015, Piscataway, NJ, USA.

Lui, V., Gamage, D., & Drummond, T. (2015). *Fast Inverse Compositional Image Alignment with Missing Data and Re-weighting*. Paper presented at the British Machine Vision Conference (BMVC) 2015, Swansea, UK.

Maire, F., Alvarez, L. M., & Hodgson, A. (2015). *Automating marine mammal detection in aerial images captured during wildlife surveys: A deep learning approach*. Paper presented at the AI 2015: Advances in Artificial Intelligence.

Malti, A., Bartoli, A., & Hartley, R. (2015). *A Linear Least-Squares Solution to Elastic Shape-from-Template*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

Marouchos, A., Muir, B., Babcock, R., & Dunbabin, M. (2015). *A shallow water AUV for benthic and water column observations*. Paper presented at the MTS/IEEE OCEANS, 2015, Piscataway, NJ, USA.

McMahon, S., Sünderhauf, N., Milford, M., & Upcroft, B. (2015). *TripNet: Detecting Trip Hazards on Construction Sites*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015.

McMahon, S., Sünderhauf, N., Upcroft, B., & Milford, M. (2015). *How Good Are EdgeBoxes, Really*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

Milford, M., Shen, C., Lowry, S., Suenderhauf, N., Shirazi, S., Lin, G., Upcroft, B. (2015). *Sequence Searching with Deep-learnt Depth for Condition-and Viewpoint-invariant Route-based Place Recognition*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

Ok, K., Gamage, D., Drummond, T., Dellaert, F., & Roy, N. (2015). *Monocular image space tracking on a computationally limited MAV*. Paper presented at the IEEE International Conference on Robotics and Automation (ICRA), 2015.

Paisitkriangkrai, S., Sherrah, J., Janney, P., & Van-Den Hengel, A. (2015). *Effective semantic pixel labelling with convolutional networks and conditional random fields*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, Piscataway, NJ, USA.

Pepperell, E., Corke, P. I., & Milford, M. J. (2015). *Automatic image scaling for place recognition in changing environments*. Paper presented at the IEEE International Conference on Robotics and Automation (ICRA), 2015, Piscataway, NJ, USA.

Pham, T. T., Reid, I., Latif, Y., & Gould, S. (2015). *Hierarchical Higher-order Regression Forest Fields: An Application to 3D Indoor Scene Labelling*. Paper presented at the International Conference on Computer Vision (ICCV) 2015.

Polok, L., Lui, V., Ila, V., Drummond, T., & Mahony, R. (2015). *The Effect of Different Parameterisations in Incremental Structure from Motion*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015, Canberra, Australia.

Rezazadegan, F., Shirazi, S., Milford, M., & Upcroft, B. (2015). *Evaluation of object detection proposals under condition variations*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, Boston, Mass.

Rezazadegan, F., Shirazi, S., Sunderhauf, N., Milford, M., & Upcroft, B. (2015). *Enhancing human action recognition with region proposals*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA2015), Australian National University, Canberra.

Salas, M., Latif, Y., Reid, I. D., & Montiel, J. (2015). *Trajectory Alignment and Evaluation in SLAM: Horn's Method vs Alignment on the Manifold*. Paper presented at the Robotics: Science and Systems (RSS) 2015.

Sergeant, J., Sünderhauf, N., Milford, M., & Upcroft, B. (2015). *Multimodal Deep Autoencoders for Control of a Mobile Robot*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015.

Sokeh, H. S., Gould, S., & Renz, J. (2015). *Determining interacting objects in human-centric activities via qualitative spatio-temporal reasoning*. Paper presented at the Asian Conference on Computer Vision (ACCV) 1-5 Nov. 2014, Cham, Switzerland.

Spek, A., & Drummond, T. (2015). *A Fast Method For Computing Principal Curvatures From Range Images*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015.
Sünderhauf, N., Dayoub, F.,

McMahon, S., Eich, M., Upcroft, B., & Milford, M. (2015). *SLAM–Quo Vadis? In Support of Object Oriented and Semantic SLAM*. Paper presented at the Robotics: Science and Systems (RSS) 2015, Rome, Italy.

Sünderhauf, N., Dayoub, F., Shirazi, S., Upcroft, B., & Milford, M. (2015). *On the Performance of ConvNet Features for Place Recognition*. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015.

Sünderhauf, N., Upcroft, B., & Milford, M. (2015). C*ontinuous Factor Graphs For Holistic Scene Understanding*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

Wang, P., Shen, C., & van den Hengel, A. (2015). *Efficient SDP inference for fully-connected CRFs based on low-rank decomposition*. Paper presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, Piscataway, NJ, USA.

Wang, X., Sekercioglu, Y. A., & Drummond, T. (2015). *Self-calibration in visual sensor networks equipped with RGB-D cameras*. Paper presented at the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2015, South Brisbane, QLD.

Weberruss, J., Kleeman, L., & Drummond, T. (2015). *ORB Feature Extraction and Matching in Hardware*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015.

Zhang, F., Leitner, J., Milford, M., Upcroft, B., & Corke, P. (2015). *Towards Vision-Based Deep Reinforcement Learning for Robotic Motion Control*. Paper presented at the Australasian Conference on Robotics and Automation (ACRA) 2015.

Glossary

3DV - International Conference on 3D Vision

ACEMS - ARC Centre of Excellence in Mathematical and Statistical Frontiers

ACFR - Australian Centre for Field Robotics

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

ACRV - Australian Centre for Robotic Vision - an ARC Centre of Excellence

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

EOI - Expression of Interest

EUAB - End-User Advisory Board

FPGA - Field-Programmable Gate Array

IBVS - Image-Based Visual Servo

ICIP - IEEE International Conferenc eonf Image Processing

ICRA - IEEE International Conference on Robotics and Automation

IEEE - Institute of Electrical and Electronics Engineering

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

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.

Index

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.

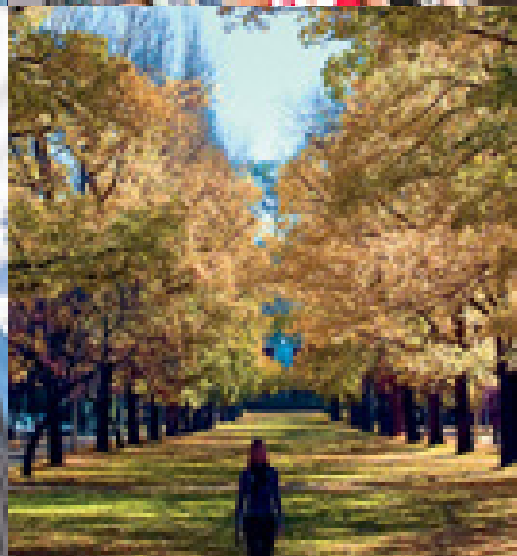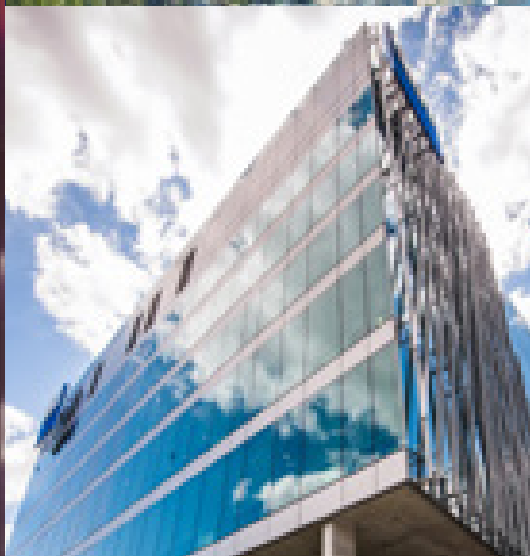The silver outer circle represents the sclera, the protective, outer layer of the eye.

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.

Design by Studio 55

Photography Greg Lee

creating robots that see