ACRV
AUSTRALIAN CENTRE FOR ROBOTIC VISION

ANNUAL REPORT 2014

Australian Government
Australian Research Council

ARC CENTRE OF EXCELLENCE FOR ROBOTIC VISION

OUR VISION AND VALUES

Vision

Creating robots that see and understand their environment
that can work safely with people
to benefit society

Values

Cool Research	• Taking on the big challenges • Not just any run-of-the-mill research • Being the best in robotic vision
Impact	• Being at the forefront - international networks and recognition • Research that is relevant, timely and sustainable • Transform the world in a positive way
Creativity	• Be bold • Attempt the impossible • Imagine, create, tell the world, repeat...
Entrepreneurial Spirit	• Foster enterprise • Inspire selection of different paths • Rapid prototyping and demonstration of ideas
Community	• Respect • Openness • Embracing diversity
Supportive	• Empowering researchers to focus on what they do best • Nurturing the next generation of robotic vision researchers • Communications, conversations and crazy ideas
Fun	• Enjoying what we do • Celebrating success • Sparking motivation/ideas

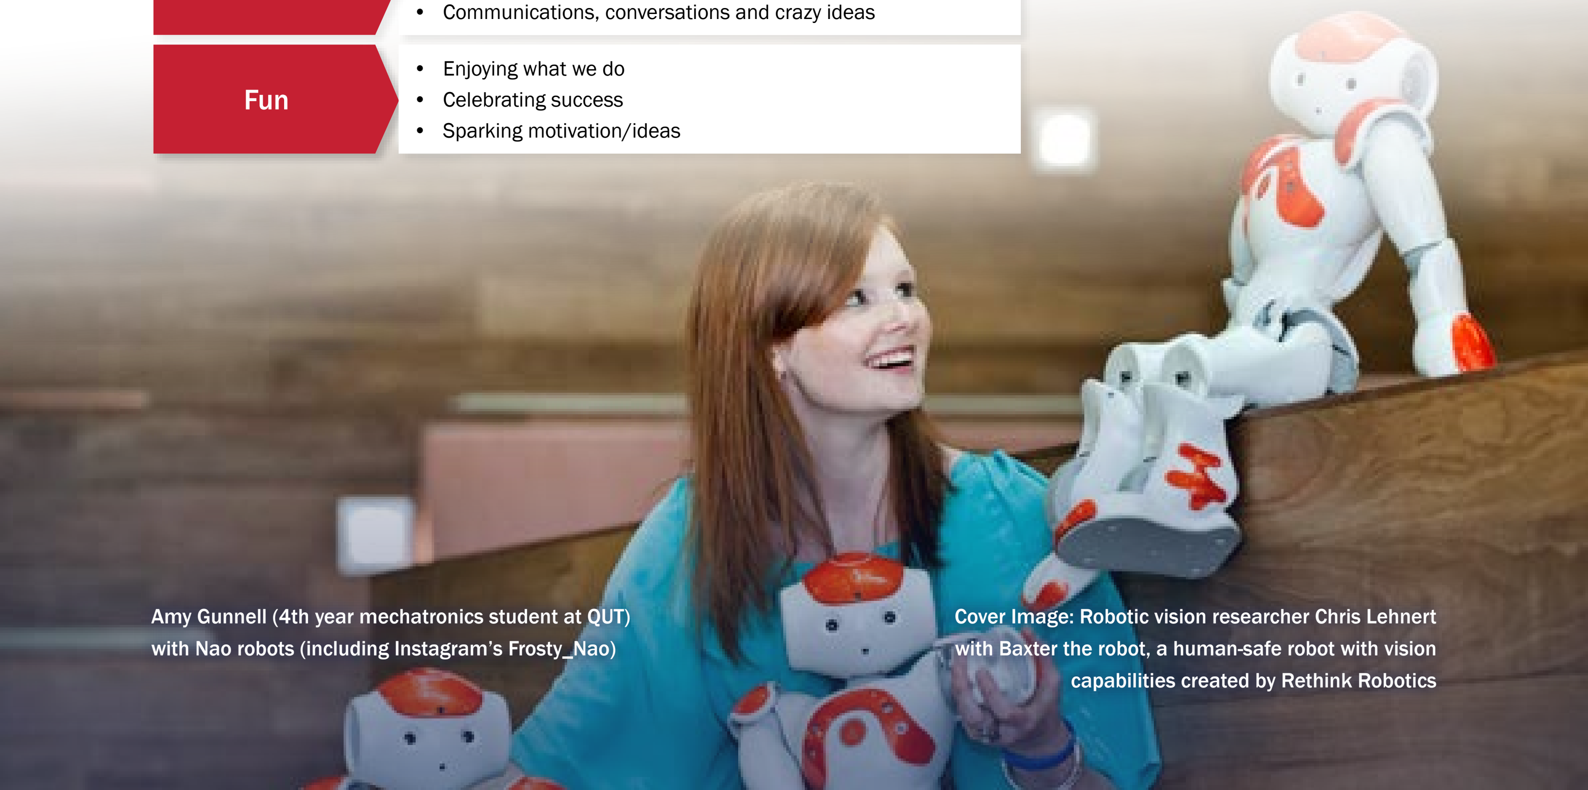

Amy Gunnell (4th year mechatronics student at QUT) with Nao robots (including Instagram's Frosty_Nao)

Cover Image: Robotic vision researcher Chris Lehnert with Baxter the robot, a human-safe robot with vision capabilities created by Rethink Robotics

TABLE OF CONTENTS

About ACRV

Who we are

Robots are changing the way we live and work. The Australian Centre for Robotic Vision (ACRV) brings together Australia's top researchers in computer vision and robotics to lead the world in robotic vision research. Enabling robots to see and understand the world, robotic vision, is the key technology that will allow robotics to transform labour-intensive industries, disrupt stagnant markets, and ensure robots become a ubiquitous feature of the modern world.

The ACRV is an unincorporated collaborative venture with funding of $25.6m over seven years to pursue an ambitious research agenda tackling the critical and complex challenge of applying robotics in the real world. $19m in public funding has been contributed from the Australian Research Council's Centres of Excellence program. To create the Centre we have assembled an interdisciplinary research team from four leading Australian research universities Queensland University of Technology (QUT), The University of Adelaide (UoA), The Australian National University (ANU), and Monash University as well as NICTA (previously known as National ICT Australia Ltd), and overseas universities and research organisations including INRIA Rennes Bretagne, Georgia Institute of Technology, Imperial College London, the Swiss Federal Institute of Technology Zurich, and the University of Oxford.

ACRV 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.

Centres of Excellence

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. Once upon a time they were closely related and investigated, separately and together, in Artificial Intelligence (AI) labs around the world. Vision has always been a hard problem, and early roboticists struggled to make vision work using the slow computers of the day — particularly for metric problems like understanding the geometry of the world. In the 1990s affordable laser rangefinders entered the scene and roboticists adopted them with enthusiasm, delighted with the metric information they could provide. Since the late 1990s, laser-based perception has come to dominate robotics, while processing images from databases, not from robots, has come to dominate computer vision. The ACRV will reprise the early partnership between robotics and vision, undertaking research that will allow vision to be considered a powerful and cost-effective sensor for robots.

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

Just as vision played a major role in the evolution of life on Earth it will underpin the the capability of robots that understand their environment, make decisions and 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 too will the capabilities of robots.

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 involving commitments from 13 Chief Investigators (CIs) and 13 Associate

Investigators (AIs) based at Australian universities (QUT, ANU, UoA and Monash), along with six Partner Investigators (PIs) from international research organisations (see People profiles p. 59), 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 and proceeded to a full proposal submitted in August 2013. After supplying a rejoinder in response to review comments, the ACRV was short-listed and interviewed in Canberra in November 2013.

On 17th December 2013, the Hon. Minister for Education and Training, Christopher Pyne, MP, approved the allocation of $285m over seven years for 12 new ARC Centres of Excellence, including the ARC Centre of Excellence for Robotic Vision. The Centre commenced on 1st January 2014 and became fully operational on 21st July 2014 following the signing of legal agreements between all collaborating organisations.

Timeline showing the Life of the Centre

Map showing location of ACRV partners

About this report

Report Description

Our report covers the activities of the ARC Centre of Excellence for Robotic Vision for the 2014 calendar year. Activities encompass research, training, business strategy, 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

The 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: QUT, The Australian National University, Monash University, and The University of Adelaide. In return for this funding the ACRV have 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 aims to outline our vision for our new Centre as well as highlighting our achievements and providing an overview of our operations for 2014. 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 (ACRV).

You will also find this report, and various other ACRV publications on our website at **www.roboticvision.org**

To provide feedback on this report please visit www.roboticvision.org/feedback

Centre Performance

Objectives

The objectives of the ACRV are to:

- create robots that see and understand their environment;
- do internationally impactful science;
- establish a vibrant national/international robotic vision community;
- develop the next generation of robotic vision experts;
- transform industry and generate wealth for the community;
- increase educational opportunities in robotic vision at high school, undergraduate, and postgraduate level;
- foster innovation, entrepreneurship, and new enterprises that convert research into useful products and services;
- connect research organisations, governments, industry, and the private sector to build critical mass in robotic vision;
- breakdown the barriers that exist to female participation in computer science and engineering, and encourage general community engagement with STEM;
- engage with people about robotic and vision technologies and the impact these technologies will have on society and the way we work.

Highlights for 2014

42 Conference Publications

10 New Honours Students

8 International Visitors

79 Briefings

40 Media Pieces

13 Awards

34,000 Website Hits

Director's Report

ACRV Director Peter Corke

It's time for robots

Robotics and computer vision have teetered on the cusp of achieving their potential for more than 50 years but to date, technologies have fallen short of our imagination — advanced machines in movies far surpass the abilities of modern robots. Until the advent of the Roomba vacuum cleaner, a successful household robotic helper, robots were mainly restricted to factory floors, where they have undergone little substantive change since their advent in the 1950s. Such industrial robots are excellent at repetitive tasks that require high speed and precision but they also require controlled environments where humans are kept separate from machines and the robots do not perceive or adapt to their environment because they are without the capability of sight.

Thanks to accelerating advances in sensors, actuators, computation and machine-to-machine communication, robots are gaining the ability to adapt to their surroundings, making it possible to deploy them safely alongside workers. Advanced robotics is predicted to be one of the top 10 disruptive technologies over the next ten years (Disruptive Technologies, McKinsey Global Institute, 2013). The goal of our Centre 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.

Building the Centre

In mid-2014, once all legal agreements were in place, we commenced recruitment for the Centre, hiring a Chief Operating Officer (Dr Sue Keay) and Centre Coordinator (Laura Vernon), as well as six Research Fellows and six PhD candidates with many more to come on board in 2015. Over the life of the Centre, our funding will be used to support 24 Research Fellows and 80 PhD candidates across our collaborating organisations. We also accepted the nominations of two additional Associate Investigators (AIs), Professor Tristan Perez (Agricultural Robotics) and Professor Jonathan Roberts (Medical Robotics), bringing our total number of AIs to 15 (see AI profiles p. 65).

We have put in place internal management processes (policies and guidelines), Confluence-based intranet to facilitate communication (science and administration) between researchers at all nodes, and launched our external web site (roboticvision.org). The Centre Executive held 21 meetings and to-date we have held two meetings of all Chief Investigators. The Centre schedules monthly research fellow and PhD candidate meetings, supports a research training toolbox on our intranet, while also developing research training and mentoring programs for all early career researchers (research fellows and PhD candidates).

We held our first annual symposium "RoboVis" at ANU (Canberra) in November (see p. 12). RoboVis

Refurbished space for the Centre of Excellence at ANU represented the first time all members of the Centre; Chief Investigators, Associate Investigators, Research Fellows, PhD candidates, and friends of the Centre have been 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.

Industry engagement is critical to the success of our Centre, and I was delighted to have twelve representatives from industry at our first symposium. RoboVis is a showcase and a celebration of the Centre’s research, an annual event that will, in the future, also involve our Centre Advisory Committee (CAC).

Centre researchers have published 15 journal papers, 42 conference papers and one book chapter in top robotics and computer vision outlets.

In terms of physical facilities the Centre enjoys new purpose-built facilities at ANU (see p. 5, 12) and a major refurbishment of space at QUT and the University of Adelaide will be completed during 2015.

In November, 2014, we had the opportunity to demonstrate the Centre’s work in agricultural robotics to Indian Prime Minister Narendra Modi of India during the G20 held in Brisbane (see Improving Food Productivity p. 35). The Centre includes a $3m program in agricultural robotics, funded by the Queensland Department of Agriculture and Fisheries. One outcome of this

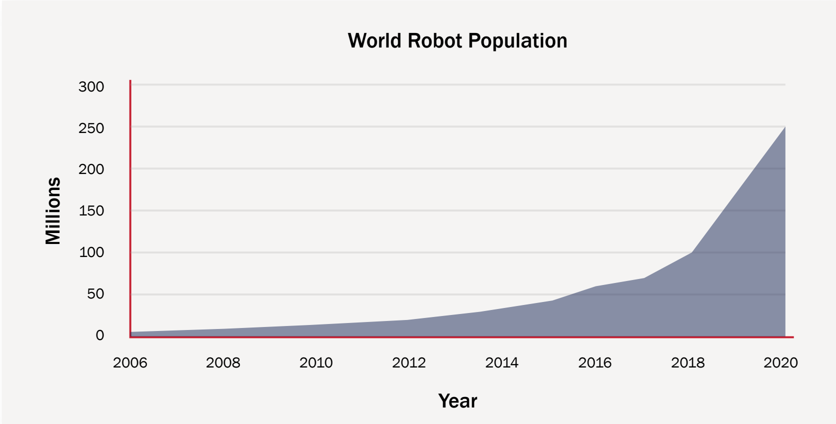

Figure 1: World robot population interpolated from global robotics industry data sourced from the International Federation of Robotics (www.ifr.org)

program is a prototype field robot called the AgBot II, which has been developed for broadacre weeding, fertilising, and horticulture by Centre-affiliated researchers at QUT.

At the end of the year Centre researchers at QUT received funding for 2015 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. We will work with the ARC Centre of Excellence in Mathematical and Statistical Frontiers (ACEMS) to interpret the data collected by the robots to further our understanding about how to control the threat these starfish pose.

Centre researchers have published 15 journal papers, 42 conference papers and one book chapter in top robotics and computer vision outlets. The number of best paper awards won by Centre researchers in 2014 illustrates the high calibre of their research efforts (see Esteem Measures p. 14-15). We have also developed two 6-week massive open online courses, MOOCs, which underwent user acceptance tests and will be launched globally in 2015.

Outlook for the Centre

2015 is going to be a busy and productive year. Importantly, we will complete our recruitment at each node. I look forward to growing our team, welcoming more research fellows and PhD candidates to the Centre and engaging them in our training and mentoring programs. In conjunction with NICTA we’re running our first summer school on Robotic Vision. Open to PhD students from the Centre and beyond, it will be a great opportunity to learn the fundamentals from experts in the field. We’ll be chairing a number of workshops at major international conferences, putting our first MOOCs online, running our annual symposium RoboVis in South Australia, starting an outreach program, ramping up connections to industry, delivering on collaborative research projects and winning new ones, holding our official launch, showing off our new-look website, and creating this, our first annual report.

With a full complement of researchers and support staff we will push the envelope on the underpinning science and technology of robotic vision that the Centre has committed to deliver. Our research is grounded by national challenges around productivity and competitiveness and we will expand Australia's knowledge base through the production of research publications, open source and protected intellectual property and the targeted communication of these results. We will convene our advisory committee (CAC) and recruit our end-user advisory board (EUAB) to provide guidance on our program and our modes of engagement. Face-to-face and video communications, as well as travel between nodes, will be a major focus for 2015 to enhance Centre effectiveness.

The Centre plans to use a variety of engagement models (see Communication, Outreach and Engagement p. 41) with enterprises from small to large: an affiliates program; active collaboration with others in the national innovation system including non-partner universities and organisations such as NICTA, CSIRO and DSTO 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.

Over the lifetime of the Centre the computing power to cost ratio is likely to increase by a factor of 1,000, 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). Vision will become the sensing modality of choice for real world robotic projects, with the ACRV well-positioned to deliver the breakthrough science and technologies that will create a new generation of robots that can visually sense, and understand, complex and unstructured real-world environments.

I would like to thank the Centre Executive, my fellow Chief Investigators, and all the researchers and staff who have helped invent this Centre from scratch. Together we have a rare opportunity to work on a truly grand challenge - creating robots that see - that will have a profound and positive impact on the future. We thank the Australian Research Council for providing us with this opportunity and believing in the potential of the Centre. We also have an important responsibility to create human capital, to nurture the next generation of innovating robotic vision experts who will go on to work in industry, government and academia. They will help translate research results in robotics and vision to existing industries and convert their research into new products, services and enterprises so as to forge whole new industries. The next twelve months of the Centre is where our vision will become a reality and I look forward to working with you to solve the challenges ahead.

Peter Corke
Director

What is a robot?

- A **robot** is an autonomous machine, moving within its environment, to perform intended physical tasks.
- A **field robot** is a mobile robot that performs useful tasks in an outdoor environment, e.g., agriculture, forestry, mining, construction, logistics etc.
- An **industrial robot** is a robot used in a factory environment. These were the first generation of robot technology in a controlled and organised manufacturing environment.
- A **service robot** is a robot (typically mobile) that performs useful tasks for humans in a non-factory environment.
- A **personal service robot** is a service robot for personal use (non-commercial) e.g., robot vacuum cleaner, entertainment robot.
- A **professional service robot** is a service robot for professional use (commercial) e.g., a surgery robot in a hospital.

A Nao robot demonstrates vision capabilities driving an electric toy car and avoiding obstacles

Market Trends

Robots are becoming cheaper, more adaptable, easier to program and capable of learning and switching between tasks. The McKinsey Global Institute (2013) estimates that the worldwide economic impact of advanced robotics by the year 2025 could be $1.7 trillion to $4.5 trillion each year (http://tinyurl.com/nmbecug). A recent study by Boston Consulting Group (http://tinyurl.com/qetyfts) suggests that robots will boost manufacturing productivity in some countries by 30%, with the price of robots set to decrease by 20% over the next decade, making them affordable to SMEs. The Baxter robot, from Rethink Robotics (see Baxter p. 21) is an example of this next generation of low-cost manufacturing robot. As the mechanical dexterity and perceptual capability of robots improves, they will move away from the assembly line and into the wider world — playing a greater role in personal and professional services where significant market growth is expected (see Figure 2).

The market for robots is booming. The Economist, a prestigious London-based business magazine with global readership had a 14-page Special Report in their March 2014 issue, “Rise of the Robots”. Robot orders and shipments in North America set new records in 2014, with shipments of robots (worth $1.6 billion) up 28% from 2013, according to Robotic Industries Association (RIA), the industry’s trade group. According to the International Federation for Robotics (www.ifr.org), the market for industrial robots will grow by 12% pa from 2015-2017 (from > 200,000 sold in 2014), while the market for professional service robots is expected to grow from $USD1.7bn pa to $USD18.9bn from 2015-2017. Demand for personal service robots is predicted to reach 23.9m units by 2017 with a value of $USD6.5bn. Apart from personal service robots, increased investment in robotics is forecast to occur for defence, agriculture, logistics, and medical applications.

Advanced robotics could improve the productivity of many Australian industry sectors, while robots themselves could form a new export industry for Australia. Although robotics hardware and software groups 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 people with expertise to lead the new era of robotics companies. The Centre anticipates playing a critical role in enabling this by improving collaboration between researchers and industry.

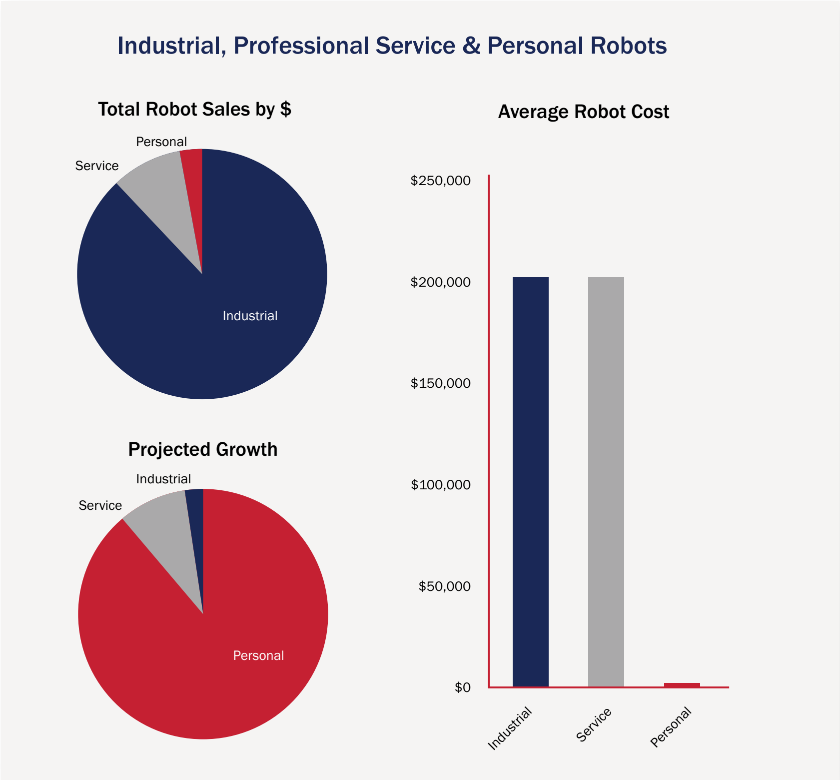

Figure 2: The projected growth of personal service robots, set to overshadow growth of professional service and industrial robots, although numbers of these robots continue to increase. Data sourced from International Federation of Robotics (www.ifr.org)

Robotics Worldwide

As the trend to adopt more robots unfolds, leading companies like Google, Facebook, Uber, and Amazon are acquiring robotics, vision and machine learning companies and the expertise they contain. In 2012 Amazon acquired Kiva systems (a company developing robots for warehouses) for $USD775m, Japan's Softbank acquired an 80% stake in Aldebaran robotics for $USD100m and in 2013 Google acquired seven robot companies including; SCHAFT Inc., Industrial Perception, Redwood Robotics, Meka Robotics, Holomni, Bot & Dolly, Boston Dynamics, as well as deep neural networks company DNNresearch Inc. Apple acquired Prime Sense, creators of the technology used in the Microsoft Kinect sensor. The trend continued in 2014 with Google acquiring machine learning and automation companies; Nest labs, Deep Mind Technologies, Titan Aeropace, Quest Visual, Dropcam, Jetpac, Dark Blue labs, Vision Factory and Revolv. Technology Review quotes Google's research director, Peter Norvig, as claiming that the company now employs more than 5% of the worlds' leading experts in machine learning. Uber recently hired 50 scientists from Carnegie Mellon, a U.S. university with a well-known strength in robotics engineering, to advance their quest to get a start in the rapidly evolving autonomous vehicle market (Feb, 2015 http://tinyurl.com/p7hyhf4).

Social Impact of Robotics

With increased adoption of robotics, the total cost of manufacturing labour in 2025 could be 16 percent lower, on average, in the world's 25 largest goods-exporting nations than they would be otherwise --- according to Boston Consulting Group (http://tinyurl.com/qetyfts). Depending on the industry and country, labour productivity could rise by an estimated 10 to 30 percent over and above productivity gains that typically come from other measures. With this increased productivity, understandably, comes a concern that robots are a threat to jobs. Oxford University Professor Michael Osborne, predicts that in 20 years time the world of work will look quite different with 47% of current jobs done by robots. According to Osborne, the most likely impact of robots on jobs are in accommodation and food services, transportation and warehousing, and real estate. The Centre will be sensitive to community concerns about the impact of robots on the workforce and about privacy issues that arise from robots equipped with cameras, and will proactively contribute to public education and debate in these areas.

Robotronica, held in conjunction with National Science Week, is an opportunity for the community to engage with robots

Plans for 2015

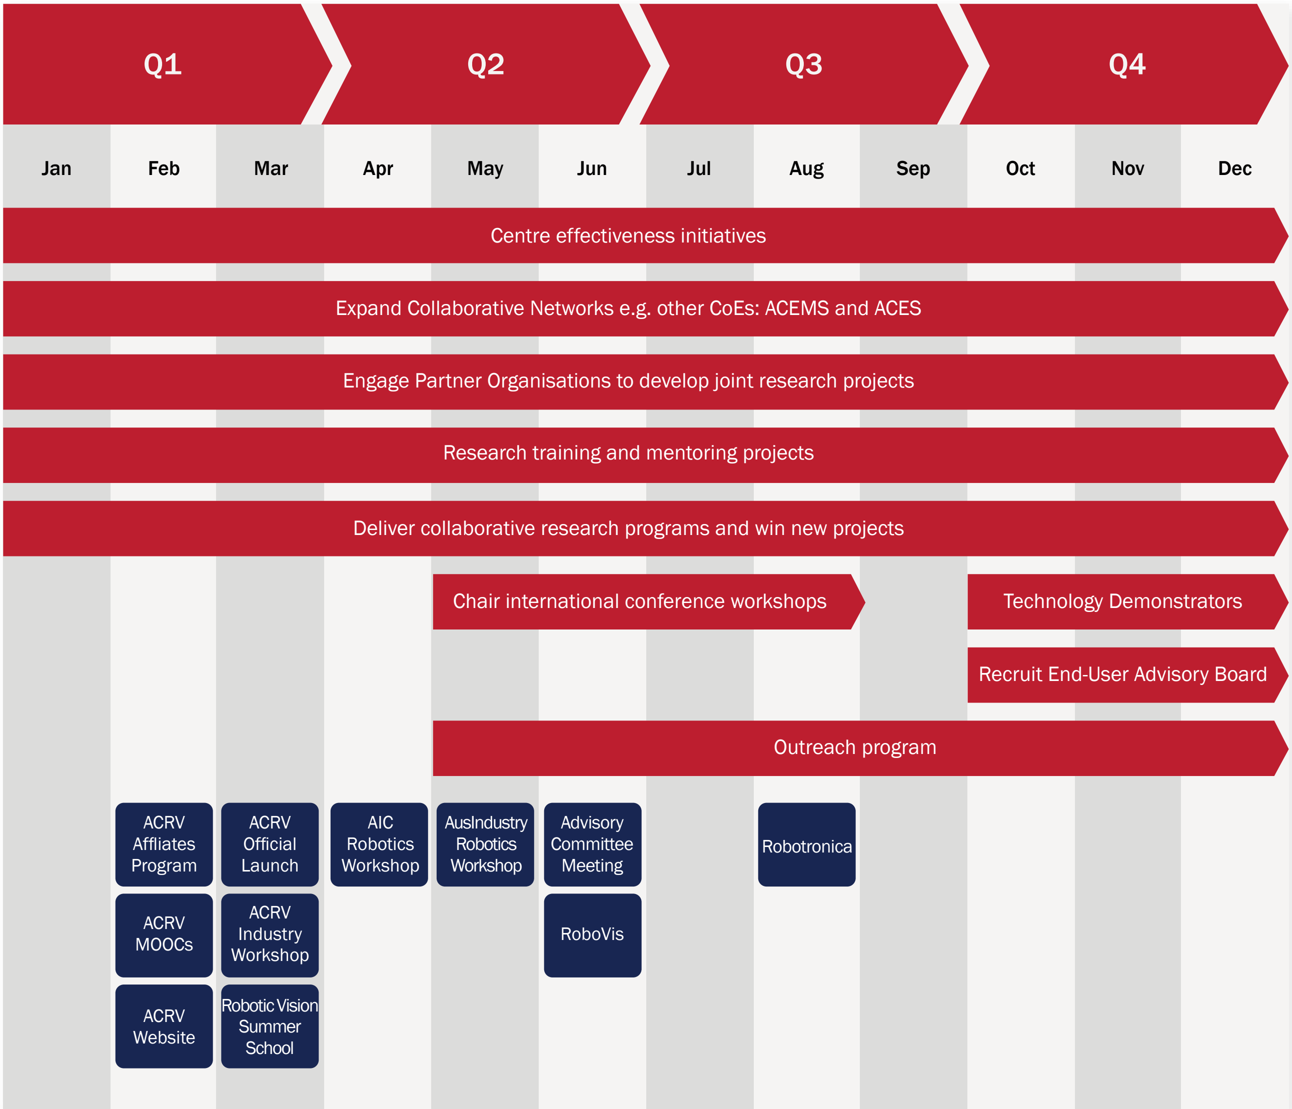

Inaugural Annual Symposium

RoboVis 2014

RoboVis2014

The Centre held its first symposium, RoboVis2014, at ANU in Canberra in November 2014. It was the first time all members of the Centre, including Chief Investigators, Associate Investigators, Research Fellows, PhD candidates, and friends of the ACRV have been together in one location, and the event represented a critical milestone in the Centre's formation.

The two-day symposium was called "RoboVis" as a symbol of what the Centre is seeking to achieve, bringing researchers from the fields of robotics and computer vision together.

RoboVis2014 was opened by Chief Defence Scientist, Dr Alex Zelinsky, Chair of the ACRV Advisory Committee, and also the Chief Executive of the Defence Science and Technology Organisation (DSTO). Alex delivered a presentation on the “Applications of Robotic Vision”. This was followed by talks from some of the Centre's new research fellows.

Professor Jenny Corbett, Pro Vice-Chancellor (Research and Training) ANU, officially opened the ANU's ACRV node. The ANU has recently completed an extensive refurbishment of space to accommodate the Centre's researchers, students and laboratories. The opening was followed by a tour of the new space complete with poster and demonstration sessions.

Industry engagement is key to the success of the Centre and an Industry Applications session was held as part of the symposium, attended by twelve Canberra-based industry representatives. Speakers included several of our industry guests.

The second day of the symposium explored the Centre’s themes of Algorithms and Architecture (AA), Robust Vision (RV), Semantic Vision (SV), and Vision and Action (VA). The Centre's Executive and Theme Leaders, Tom Drummond, Peter Corke, Ian Reid and Rob Mahony, led the sessions addressing the key scientific challenges of each theme.

The symposium wrapped up with discussions around ACRV's future work, including recruitment of PhD candidates, and hosting international conferences.

RoboVis will continue to be an annual showcase and celebration of the Centre's research with RoboVis2015 scheduled to be hosted by ACRV's University of Adelaide node.

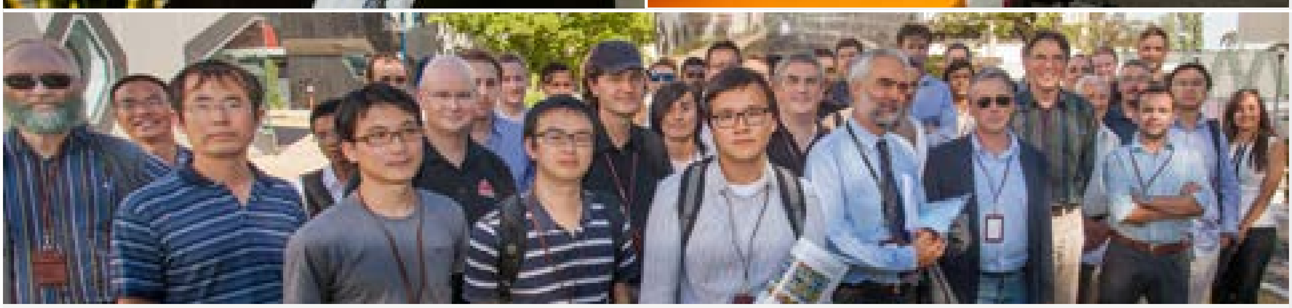

Clockwise from top left: ACRV CAC Chair Alex Zelinsky opens the inaugural ACRV annual symposium, RoboVis2014; ANU's newly refurbished space for Centre researchers; delegates at the symposium.

Research Performance

Research Themes

The Centre will achieve its aim to develop breakthrough science and technology in robotic vision by addressing four key research objectives: Robust Vision (RV), Vision and Action (VA), Semantic Vision (SV), and Algorithms and Architecture (AA). Together the four research objectives form the Centre's research themes, which serve as organisational groupings of the Centre's research projects.

Robust Vision (RV) 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.

Vision and Action (VA) 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.

Semantic Vision (SV) 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.

Algorithms and Architecture (AA) 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.

ACRV Research Themes and Projects

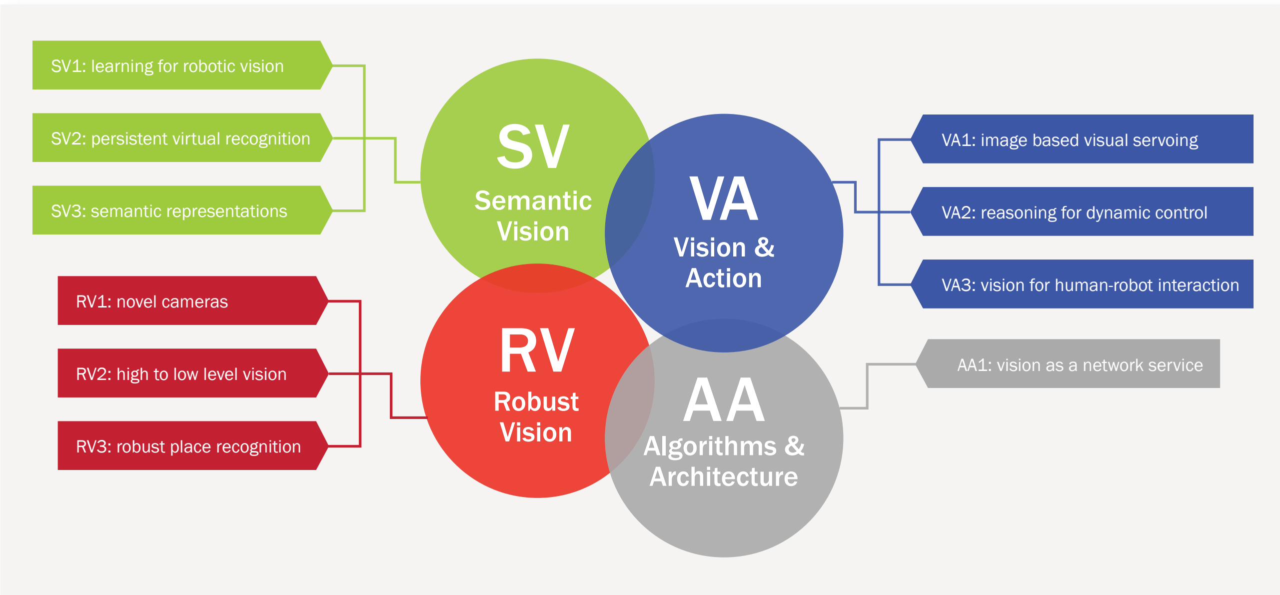

Research Performance (KPIs)

Research Findings

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014
Number of research outputs	Annually	-	-
Conference Publications		25	42
Journal Publications		12	15
Disclosures/Patents		0	0
Citation data for publications *(list web hits for online articles separately)* *Measured using Google Scholar*	At review	0	N/A

Esteem Measures

Recognising our People

The excellence of ACRV's people has been recognised this year with a number of our researchers receiving prestigious awards or gaining appointments or promotions during 2014.

CI Ian Reid and colleagues from Oxford have been awarded the Best Paper Award at the International Conference on 3D Vision (3DV), held in Tokyo, Japan in December 2014 for their paper, "3D Tracking of Multiple Objects with Identical Appearance using RGB-D Input".

The Best Paper award at the Australasian Conference on Robotics and Automation (ACRA 2014) held in Melbourne in December 2014, went to ACRV Research Fellow Donald Dansereau for his work at the Australian Centre for Field Robotics (University of Sydney). The paper was entitled 'Exploiting Parallax in Panoramic Capture to Construct Light Fields'. Three papers associated with the Centre's RV3 project were also finalists for the Best Paper award involving ACRV researchers Michael Milford, Zetao Chen, Obadiah Lam, Adam Jacobson, Ed Pepperell, Matt Dunbabin and Peter Corke. ACRA is run by the Australian Robotics and Automation Association (ARAA).

CI Hongdong Li supervised the research and co-authored a paper that won the Institute of Electrical and Electronics Engineering (IEEE) International Conference on Image Processing (ICIP) 2014 Best Student Paper Award for the paper entitled "Null Space Clustering with Applications to Motion Segmentation and Face Clustering". The paper awards for ICIP 2014 were selected from over 2,000 papers submitted to the conference.

The best paper at the International Semantic Web Conference (ISWC) 2014 was awarded to ACRV AI Wai Ho Li.

Other measures of esteem awarded to ACRV researchers include: an ARC Future Fellowship to CI Michael Milford; a Humboldt fellowship for senior researchers awarded to CI Gustavo Carneiro; CI Robert Mahony appointed as Director of ANU's Research School of Engineering; CI Gordon Wyeth appointed as

Executive Dean of QUT's Science and Engineering Faculty; and AI Jonathan Roberts awarded a CSIRO Science Excellence Medal. CI Peter Corke was recognised with a QUT Vice Chancellor's Award for Excellence in leadership, learning and teaching, research, partnerships and engagement; QUT's Centre of Excellence team (CIs Corke, Wyeth, Upcroft & Milford) was recognised with a QUT Vice Chancellor's Performance Award; and CI Peter Corke 2014 was recognised by a Vice-Chancellor's Performance Fund Award for outstanding achievements as a member of the Robotics MOOC team.

Partner Investigator Professor Paul Newman was elected a Fellow of the IEEE for his contributions to robot navigation. IEEE Fellow is a distinction reserved for select IEEE members whose extraordinary accomplishments in any of the IEEE fields of interest are deemed fitting of this prestigious grade elevation. The Board of Directors confers the IEEE Grade of Fellow and fellows comprise less than one per cent of the IEEE's membership. ACRV CIs Peter Corke and Richard Hartley are IEEE Fellows, as are AI Fatih Porikli and PIs François Chaumette and Marc Pollefeys.

There have also been many promotions amongst Centre researchers including CI Chunhua Shen and AI Tristan Perez promoted to Professor, CIs Ben Upcroft, Stephen Gould and Gustavo Carneiro promoted to Associate Professor, and AI Javen Shi promoted to Senior Lecturer.

Clockwise from top left: 3DV Best Paper Award to CI Ian Reid and colleagues, group shot of ACRV people at RoboVis2014, CI Michael Milford was awarded an ARC Future Fellowship and featured in The Australian newspaper.

Robust Vision (RV) Theme

RV Theme Leader Peter Corke

RV 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

The robot vision process starts with images acquired by cameras, but images are a complex (non-linear and non-invertible) function of: 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, by shadows, by changes due to rain or season and by extreme variation in light level. A robot 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 will rigorously investigate new sensing technologies including light-field and night-vision cameras. We will also investigate how to intelligently combine multiple low-cost sensors, to improve dynamic range, colour acuity, frame rate and field of view.

Secondly, we will develop robust visual processing algorithms that work despite distractors such as small changes in location, 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) where image change due to blur, defocus, rain or viewpoint change are significantly reduced. We will develop robust algorithms that can recognise places despite poor viewing conditions: night and day, rain or shine, summer or winter, fast moving or static.

Finally, we must infer the state of the world from images despite the image formation process being non-invertible. Current vision systems fall short of human capability when it comes to using information from multiple eyes on a body moving deliberately through the world, and using visual memory and experience, to maximise understanding of the surroundings.

Humans bring a number of tricks to bear on this problem including fusing information from two eyes, deliberate head and body motion, visual memory and a sense of visual context. Using multiple views allows the reconstruction of 3D structure, which helps resolve the ambiguities present in monocular views, e.g., to distinguish a picture of a door from a real door, and to eliminate changes due to view point.

We will develop 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 (e.g., 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 Architecture provides the computing and communications platform on which this early stage image capture and processing is built.

Project RV1

Novel camera hardware

Research Team:
Peter Corke (Project Leader)
Hongdong Li (CI)
Richard Hartley (CI)
Ben Upcroft (CI)
Andy Davison (PI)
Donald Dansereau (RF)

The Lytro ® Illum camera uses light field technology to capture the entire light field – providing much more information than a conventional camera

Modern image sensors have resolution in abundance but for robotic vision other factors are at least as important. 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. For example lightfield and random pixel cameras are becoming generally available and have potential as a depth sensor 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 indicates 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, e.g., depth from focus using high-performance variable focus optics developed for mobile phones.

This project will develop prototypes of innovative new camera technologies, explore emerging camera technologies, and quantitatively evaluate their efficacy for robotic vision. We will develop:

- innovative low-cost camera arrays to achieve one or more of these desirable characteristics using a multiplicity of commodity image sensors, with different photometric and spectral filters, placed arbitrarily on the robot. The goal is for the array to self calibrate for the location of its constituent cameras and stitch together a composite view of the world geometry, colour, and motion;
- computational photography techniques based on light-field images such as motion blur elimination, 3D reconstruction and dynamic refocus;
- evaluation methodologies to rigorously evaluate emerging camera technologies in environments typical of the Centre's application areas as well as under more controlled conditions.

The Centre will establish and equip 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, contribute to standards, support research RV2, VA1 and AA1.

Project RV2

High to low level vision

Research Team:
Tom Drummond (Project Leader)
Peter Corke (CI)
Ian Reid (CI)
Michael Milford (CI)
Ben Upcroft (CI)
Sareh Shirazi (RF)

This project will develop techniques to use semantic knowledge of the world to control how the image is sensed and processed. The common approach in computer vision is unidirectional: low-level features are found in the image and used to inform high-levels of abstraction such as place or object recognition. However, in this common model the camera settings and low-level features used are fixed and take no account of the type or condition of the scene.

We will draw on results from the Semantic Vision theme – and the potentially much richer set of low-level features available from the novel sensing technologies in RV1 – to develop the main innovation of this theme: a high to low pathway so that the system's current understanding of the environment (the semantics) and the task can optimise the sensor characteristics and specific low-level algorithms. Cortical feedback

RV2 Project Leader Tom Drummond

enables this in humans, and work in computer vision on contextual priming in single-view scene understanding points the way for the robotic vision case. We will develop the idea more fully, taking account of the constraints and also opportunities presented in robotic vision, e.g., this approach could determine that the robot is about to drive into a dark tunnel and increase the exposure time of cameras in anticipation, or that it is moving from gravel to tarmac and change low-level image processing to include detection of road markings.

We aim to develop a robotic vision system that uses semantic and affordance representations of the world developed in the SV theme to modify sensor behaviour (RV1) and low-level feature detection (AA1 and AA2), providing a 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.

Project RV3

Learning spatio-temporally robust visual representations of place

Research Team:
Michael Milford (Project Leader)
Ben Upcroft (CI)
Chunhua Shen (CI)
Ian Reid (CI)
Andy Davison (PI)
Niko Sünderhauf (RF)
Adam Jacobson (PhD)
Zetao Chen (PhD)

This 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 will approach 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 (SV2) based on the most suitable learnt features (SV1). 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 RV2, 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.

Spatio-temporally robust visual representations are essential for long-term reliable robot operation in real world environments.

The outcome of this project will be a robust system for place recognition that will work under a wide range of environment types and perceptual conditions.

RV3 Project Leader Michael Milford

Vision & Action (VA) Theme

VA Theme Leader Rob Mahony

VA 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

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 as required to inform the task, and motion in the 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. To address this, recent work has investigated 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.

Knowledge of how objects move enables prediction of their future location and where visual attention should be devoted.

The Vision and Action theme will develop modular vision-based control algorithms designed to exploit networked vision services; robust vision sensors; and learned features and semantic visual information representations developed in the AA, RV and SV themes respectively. We will develop image-based visual controllers for arbitrary robot dynamics and multiple networked cameras from AA1 (VA1); use improved image features from SV1 and scene understanding from SV2 for more intelligent control (VA2); and extend these techniques to inferring and responding to human intent in cooperative work situations (VA3).

Project VA1

Image-based visual servo control

Research Team:
Rob Mahony (Project Leader)
Peter Corke (CI)
Tom Drummond (CI)
Juxi Leitner (RF)

This project is creating the theoretical framework that allows modular vision-based control of dynamic systems based on input from an arbitrary number of cameras with varying viewpoints, sample rates and latency. The control task is to be specified independently of the cameras (from AA1) that are participating in the control; the problem is formulated in task space rather than image space. The first innovation is a modular construction in order to make use of multiple (networked) camera arrays, multiple visual cues, and multiple objectives and dynamically assigned resources. This approach will overcome limitations that are inherent in classical vision-based control schemes such as image cues leaving the field-of-view and poor visual motion perceptibility, which degrades task accuracy. The second innovation explicitly addresses performance through optimal and feed forward control strategies (such as model predictive control), which minimise a cost function based on estimated tracking error over a finite forward time interval. Key science challenges include formulating this tracking error in terms of vision data, determining the time interval based on both dynamic and computational requirements, and providing an effective framework to enable such control actions to be integrated into the modular control framework.

VA1 Project Leader Rob Mahony

We will demonstrate real-time control of complex and highly dynamic systems

We will demonstrate real-time control of complex and highly dynamic systems such as high-speed robots and fast moving objects, mediated by information from one or more networked cameras.

Baxter - bringing Vision and Action together

A robot that combines Vision and Action to work collaboratively beside humans is Baxter; a human-safe robot designed and built by Rethink Robotics, a USA (Boston) based company founded by world-renowned (Australian) roboticist, Rodney Brooks.

Standing around 1.8 metres tall, Baxter is a two-armed robot with an animated face. Unlike other industrial robots, Baxter can be "taught" tasks by simply manipulating his arms to perform the desired motion — no lines of code need to be written, which means he can be "programmed" by factory workers not software engineers. Baxter has simple object recognition capability and cameras in his wrists that allows him to be shown examples of objects that he can later find and pickup. Baxter is one of the first human-safe robots, able to detect nearby people and to stop moving if he contacts a person or other object. This gives him the ability to work beside his co-workers without the need of a safety cage.

All these characteristics make Baxter a perfect research tool for the fusion of robotics and computer vision by ACRV and Baxter has already been fitted with a Kinect sensor to improve his vision of his immediate surroundings.

Baxter, now an integral part of the ACRV family, is helping researchers to break down the technological barriers that stop robots from working safely side-by-side with humans.

"Current industrial robots are dangerous for humans to work around because they're simply not equipped to recognise and avoid unexpected obstacles," says ACRV Director Peter Corke.

"Baxter uses a range of sensors to detect movement around him, as well as flexible joints, which stop him continuing a pre-programmed movement if he meets an unexpected obstacle. We've tried to 'rush' Baxter before and can certainly vouch for his quick reflexes."

One of Baxter's party tricks is to play a game of Connect 4 but the key focus of the research at present is picking capsicums from a bush.

An industrial robot is programmed to conduct a series of repetitive tasks based on the assumption that the position of objects does not change. The task of picking a capsicum is extremely difficult for this type of robot since capsicums do not grow at a predefined position on the bush. Clearly the ability to see and adapt to what is in the world will be critical for an application like this.

Perfecting what appears to be the straightforward task of picking fruit will allow farmers to deploy robots that see, in places where labour is scarce.

"Once robots can see and understand the environment they operate in they can make decisions that allow them to work safely beside us and perform many of the tasks we take for granted" according to Peter.

ACRV's Baxter research robot uses vision to guide actions

Project VA2

High-level reasoning in vision-based dynamic control

Research Team:
Rob Mahony (Project Leader)
Ian Reid (CI)
Tom Drummond (CI)
Gordon Wyeth (CI)
Ben Upcroft (CI)

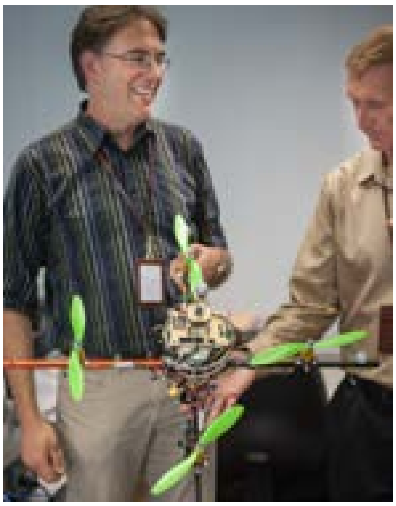

VA2 & VA3 Project Leader Rob Mahony with quadcopter used for VA experiments

VA2 builds on the work of VA1 by bringing semantic information into the control formulation. The innovation is that even partial semantic understanding of a visual scene has the potential to revolutionise vision-based control by allowing different objects in the environment to trigger different control strategies, by predicting the meaningful visual features that might be expected, and determining what control inputs might be admissible.

Initial work will determine how semantic information can be used in feature selection based on ease of visual identification, computation time and contribution to the closed-loop system sensitivity. Feature models will be provided to the control algorithm that predicts motion of the feature based on the learned structure and intent of objects. This will allow high performance predictive control to exploit scene structure while maintaining the robustness of image-based control methods. This feature processing and control will also be distributed across networked vision and computational assets using outcomes of AA1 and AA2, allowing the complex scene understanding problem to be solved in parallel with the tight closed-loop vision based control necessary to deal with highly dynamic systems.

Later work will concentrate on challenges such as what control strategy to adopt while the learning algorithm acquires enough information to understand the scene, smooth adjustment to control when new semantic information becomes available, and selection of the most appropriate control strategy given more than one high-probability hypothesis for the scene interpretation. A likely approach is to use a stochastic framework to manage a selection of control strategies and then develop probabilistic criteria to combine and switch between competing control strategies.

The results of this project will be enhanced control capability compared to VA1 that show the sensing and control strategies of robots varying according to the objects in their environment and dealing robustly with uncertain and changing labels of those objects as new information is acquired.

Project VA3

Vision for Human Robot Interaction

Research Team:
Rob Mahony (Project Leader)
Tom Drummond (CI)
Anton van den Hengel (CI)
Gordon Wyeth (CI)
Ben Upcroft (CI)

VA3 will develop 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.

The project will focus 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.

Semantic Vision (SV) Theme

SV 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

Semantic vision will allow a robot to understand what is in its environment, not just its appearance or 3D structure. This will enable the robot to interact with its environment, as semantic information will enable it to know what things are important to its task, how these things are behaving, or are likely to behave, and how these factors inform the robot's response and actions. To achieve this capability, as well as to achieve lifelong improvement and knowledge about new objects in the environment, vision systems will need to use advanced machine learning techniques. The semantics of the scene have great utility for high-level action planning, allowing reasoning about objects and how they move, which is at the heart of the perception-action loop of robotic systems.

The semantics of the scene are invariant to imaging conditions, temporal change over the short- and long-term, as well as viewpoint. Semantic Vision is therefore coupled to Robust Vision: visual cues inform semantic models, and semantic models inform low-level feature selection. Semantic Vision is also coupled to Vision and Action: action is used to gain semantic understanding, and semantics allows reasoning about how objects can move.

Semantic understanding of scenes is the "holy grail" of computer vision and the subject of considerable activity, mostly focused on segmentation and labelling of single images. There has been much less research on introducing labels and semantics into structure from motion and visual SLAM (Simultaneous Localisation And Mapping). For robots that operate in unstructured environments it is essential to develop a semantic understanding of their surroundings and to this end there have been recent efforts in segmentation and labelling of data from 3D sensors such as laser rangefinders or RGB-D sensors rather than cameras.

All of this work relies heavily on prior information about the likely appearance and geometry of regions and objects, but encoding this prior information manually is almost impossible. Recent successes have made extensive use of machine learning to automate the process of analysing large datasets to distil the empirical evidence into a prior, but typically run offline as a batch-based approach. Contextual relations (e.g., that two surface types usually occur next to each other, or that cars are more usually found on roads than on grass) can play an important role in disambiguating scene labelling. Other researchers have considered the question of the mutual context of regions and objects via a generative model that learned to distinguish different regions. The difficulty of defining and learning a generative model such as a Bayes Network *a priori* has meant that discriminative inference – e.g., using support vector machines – has met with greater success. Within the discriminative class of methods that aim to capture contextual information, the most recent and successful work has made use of the new area of structured learning in which discrete classification is generalised to finding a set of outputs that may be correlated.

SV Theme Leader Ian Reid

Important recent vision and robotics work uses this framework but there remain significant challenges. First, typically the inference is far too slow for the interactive speed required of robots. Second, these methods use a fixed set of low-level features that may not be the most effective for any given task. Third, latent variables (i.e. hidden inference layers) that permit richer descriptive classifiers, so-called deep learning, cannot currently be learned within the structured framework. Finally, these methods assume a knowledge base that is fixed and that training is a once-off rather than continual process.

The research in this theme is divided into three projects in which we: develop foundational learning techniques for robotic vision (SV1); and learn the most appropriate image features for recognition and control (SV1); learn to recognise objects and object classes reliably and in spite of change (SV2); and integrate learning and action into a framework that allows a robot to learn continuously, deal with dynamic environments and take actions to resolve uncertainty (SV3).

We place machine learning at the heart of the robot vision system to enable the system to acquire and develop semantic understanding, to detect and respond to change and to adapt its knowledge over time. We 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.

We will develop novel segmentation algorithms that can decompose a scene into its constituent parts. We will develop methods that can update and transfer knowledge between domains. Uniquely, the robotic context allows for the view to be actively selected to maximise understanding, and deliberate interactions with the environment to enhance understanding.

Project SV1

Learning for robotic vision

SV1 delivers the fundamental developments in machine learning required to support the advances in robotic vision required to create robots that can 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 to date – e.g., the strict timing constraints placed on inference, and the value of continuous and potentially active observation of the environment remain largely unexplored. Consequently problems in computer vision have been artificially compartmentalised, rather than enabling learning that is ongoing and interactive. We aim to explore new, "game-changing" learning technologies for robotic vision that address these challenges. Towards this end, we will develop: 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 efficient retraining methods and techniques for transfer of knowledge between domains, so that a new context or application can be deployed without complete redevelopment, which would discard useful experience.

First the project team will focus on feature learning. We will create a feature learning methodology that continuously estimates relevant image features from training data such that features can be shared across object classes. We will innovate using high level semantic information to drive the feature learning and inference procedures, and by using an active, on-line learning method, where the robot actively searches for object classes to optimise the feature learning objective function. Outcomes

Research Team:
Chunhua Shen (Project Leader)
Ian Reid (CI)
Stephen Gould (CI)
Richard Hartley (CI)
Gustavo Carneiro (CI)
Anton van den Hengel (CI)

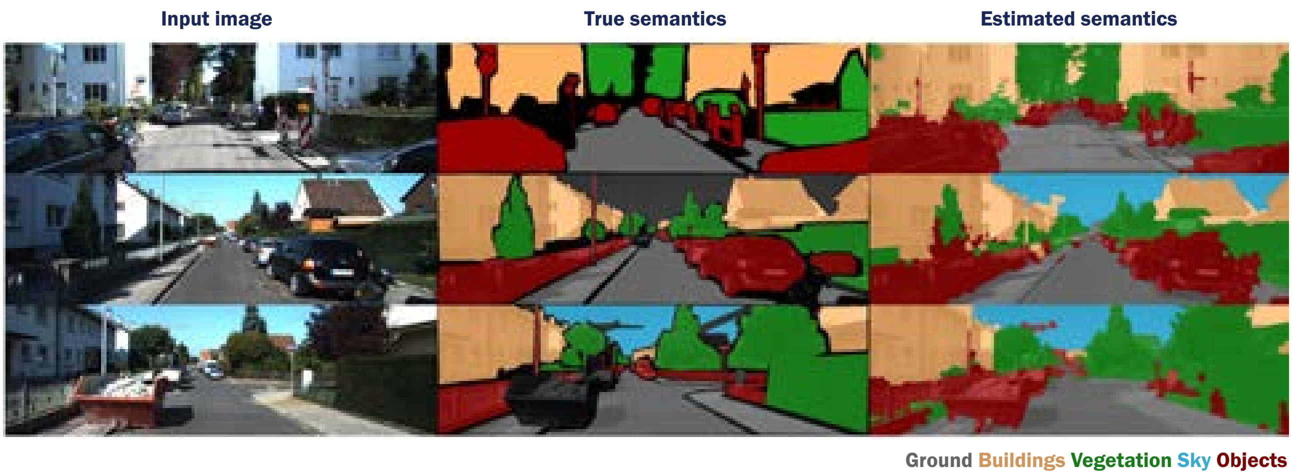

SV researchers have developed a system for scene understanding of outdoor road scenes

SV1 Project Leader Chunhua Shen

from this work will feed into RV2, RV3, to support robustness of those project outcomes to environmental conditions (e.g., sunshine or rain) for contextual priming and place recognition. The work will also feed into SV2 for general object recognition and SV3 for general scene segmentation, decomposition and understanding.

We are making substantial use of the structured output learning framework, which develops a mapping function from input data to a multivariate, or structured, output (a vector, a tree, or a graph) in which the dimensions of the multivariate output may be correlated. This approach allows modelling of contextual relationships, correlations that exist between tasks such as coarse and fine-grained recognition, and correlations between detections and object pose. We will address current limitations in the cost of training and the cost of inference associated with the use of this framework within robotic vision. A challenging issue in structured output learning is that current approaches do not admit latent internal structure in the mapping. Recently, so-called "deep networks" have achieved success because they can compactly represent more complex mappings than a simple input/output level framework. However, there is currently no work that combines rich output structures modelled using structured learning, with latent (i.e., hidden) layers of deep networks that yield more compact mappings and that may more readily generalise. We have made recent progress in training deep networks with efficient and effective approximations despite the non-convex nature of the problem.

We will develop foundational machine learning techniques that are suited to the constraints of robotic vision; that are adaptive and flexible

We will develop 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; and which underpin ACRV's other research projects.

Project SV2

Globally Persistent Visual Recognition

SV2 is developing algorithms that can efficiently and effectively detect and recognise a large and potentially ever increasing number of object classes. This outcome is a fundamental requirement for a robust, flexible robotic vision system yet remains an outstanding problem for which there is no off-the-shelf solution, despite a vast amount of literature on the subject.

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

Research Team:
Ian Reid (Project Leader)
Chunhua Shen (CI)
Gordon Wyeth (CI)
Gustavo Carneiro (CI)
Guosheng Lin (RF)
Zongyuan Ge (PhD)

SV2 & SV3 Project Leader Ian Reid

and the object may be significantly occluded. An effective system must perform rapid (real-time) classification and inference, autonomously and with high reliability.

We have adopted a learning approach with efficient and active training in order to increase reliability over time. 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?

Using the ability to learn features developed in project SV1, we will develop novel techniques for feature selection that consider the feature's value for both class-level detection and also for fine-grained recognition of subordinate classes. Historically, these problems have been treated separately but our innovation – achieved via the feature learning stage – yields gains in evaluation speed and a reduction in the memory footprint of the classifiers, as well as yielding naturally learned class hierarchies. Within numerous application domain areas of robotic vision this subordinate level classification is crucial, e.g., in a smart manufacturing environment to distinguish a variety of similar but not identical parts for assembly.

SV2 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.

Project SV3

Semantic representations for robotic vision

Research Team:
Ian Reid (Project Leader)
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)

SV3 is developing the algorithms and representations required to represent a (potentially dynamic) environment in terms of semantically meaningful entities, properties and relationships. Such representations are crucial in robotic applications to plan effective interactions such as navigation and active inspection or manipulation or other controlled actions as in VA2.

High-level scene understanding can be considered one of the grand challenges in computer vision, but progress has been limited, in part because most current research performs scene interpretation "in a vacuum", neglecting the role that a robotic vision system – actively and continuously observing – plays in such an endeavour. Meanwhile, within robotics, the geometric maps built by standard SLAM systems are useful for localisation but lack the ability to represent high-level constructs, which are more compact and expressive and required for higher-level reasoning.

Our innovation is to combine active mobile sensing with machine learning for computer vision to produce novel, meaningfully integrated semantic

Computer vision can be used to give semantic labels to objects in a scene

scene decompositions that are fit-for-purpose; i.e. they capture the objects, appearance, geometry and relationships of the scene that are relevant to the task at hand. We will achieve this by decomposing a scene into its objects, surfaces, regions i.e. the stuff (parts of the scene such as sky, floor, road surface, etc) and things (identifiable items in the scene such as cars, assembly parts, etc), and also the attributes and affordances – i.e. the action possibilities – of objects or regions in the scene. For example, stuff tends to behave in a predictable way as the robot moves but things may not. The contextual relationships between these scene elements will be important for understanding the environment for robotic interaction, and we will leverage both spatial (the layout of the environment) and temporal context (by understanding the persistent and transient elements of the environment) in recognising and distinguishing relevant from irrelevant aspects of the scene.

We are using recognition outcomes from SV1 to organise elements of a scene within a structured learning framework. This approach has become the de-facto standard over the last few years for modelling simple contextual relationships. However, in almost all existing work, inference is viewed as static – the formal optimisation problem and the algorithmic steps involved in inference are unchanged. On a robotic platform, the world is continuously changing as objects come in and out of view and more information from diverse sensors becomes available. Thus, the active robotic platform presents a unique opportunity in that the inference algorithm itself can become dynamic by adapting inferred output from the past to improve quality and efficiency of inference in real-time.

We will deliver a structured machine learning system that makes use of active robot controlled sensors to detect and label scenes using a semantic framework for representing stuff and things, together with their affordances.

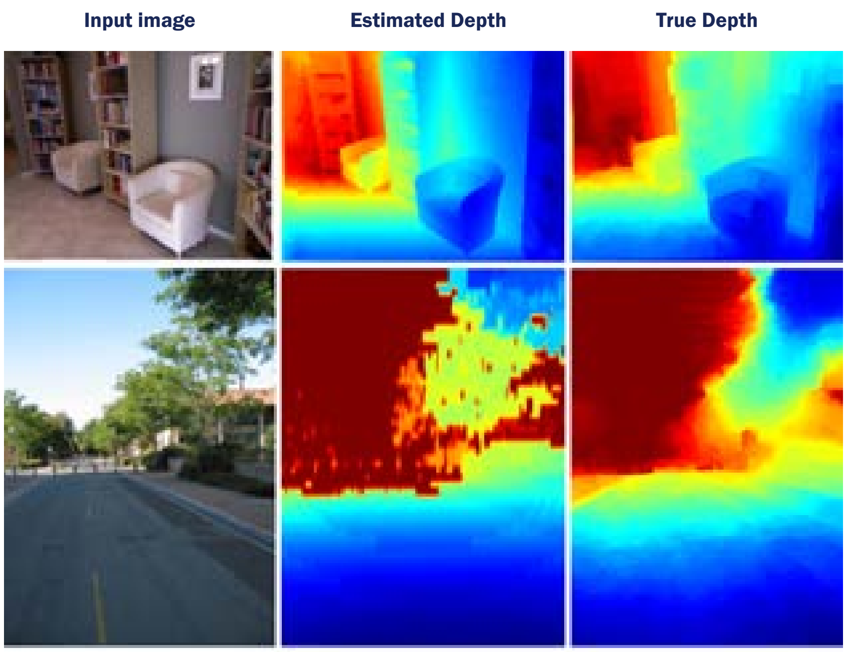

Examples of depth estimation results using deep convolutional neural fields model

Algorithms & Architecture (AA) Theme

AA Theme Leader Tom Drummond

AA 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

The Algorithms and Architecture theme will revolutionise robotic vision, moving it from system architectures, based on a tight integration of vision sensing and computational resources embedded on a robot, 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 them to be easily communicated over.

Cameras produce a prodigious amount of data, from which the essential information must be extracted.

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. A naive solution to this problem, for N points in the world, has computational complexity of $O(N^3)$. 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.

The AA theme is 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 provide pre-optimised features incorporating knowledge of scene dynamics as required by the client (AA1); an architecture that allocates computing resources as required by the task (AA2); and using results of SV2 and SV3 to apply distributed computation to scene understanding (AA3). We 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.

Project AA1

Vision as a network service

Research Team:
Tom Drummond (Project Leader)
Hongdong Li (CI)
Peter Corke (CI)
Frank Dellaert (PI)
Ahmet Sekercioglu (AI)

AA researchers are developing a Vision Operating System (VOS)

This project will transform a robot's sense of vision into a resource accessed over the network. A key aspect of this is to treat all vision sensors throughout the environment as potential resources for a robot moving in a complex environment. In this way AA1 incorporates the aims of former (from ACRV proposal) project AA4, developing dynamic environment estimates based on efficient computational implementation of forward integration of appropriate dense dynamic image models.

AA1 will develop a framework for smart, resource-aware network camera access. We 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. We will resolve this technological road-block by introducing a bi-directional communications protocol and a tailored processing architecture that adaptively specifies the vision information that is exchanged between the embedded sensor and the computing process. AA1 will integrate with outcomes from the Centre's other research projects, using both low-level robust vision algorithms and hardware developed in the RV theme as well as the high-level semantic information developed in the SV theme. In turn, the RV, SV and VA themes will benefit from the development of algorithms and architecture that are crucial to enabling robotic applications in dynamic real world environments.

An important part of developing a distributed network vision service is the ability to quickly and efficiently identify the important vision resources in an environment for a given visual task. We will develop a novel collaborative structure from motion theories that generate and maintain a compact and distributed data representation of the relative location and pose of vision resources, moving robots, and other objects of visual interest. We will also develop efficient algorithms for identification and retrieval of relevant information available to the vision network. This information will enable a robot to efficiently and robustly achieve complex and difficult tasks in dynamic and changing environments. Robots will have access to a much richer and diverse suite of visual information across a network of multiple vision resources shared between many robots.

AA1 will provide the base level communications protocols and embedded processing architectures of a Vision Operating System (VOS). The technology will be crucial in enabling areas such as flexible manufacturing or warehouse logistics operations with many interacting camera systems, multiple robots, complex and changing tasks, in which humans may be operating as well as robots, and will be presented in a Manufacturing application demonstrator.

This project will address this key technological roadblock by developing a pyramidal motion decomposition framework for dynamic image sequences. This framework will use both high-level semantic information and low-level robust vision sensors and algorithms to decompose the image information into a hierarchical motion estimation problem. The resulting computationally feasible algorithms will provide highly dynamic but low spatial-fidelity models of the environment available at high-frequency, while high spatial-fidelity models of the environment will be available at low-frequency.

The project will deliver algorithms that are crucial to enabling robotic applications in dynamic real world environments.

Application Areas

Where are all the robots?

Robots are already widely used in Australia on factory floors (automotive, food manufacturing), and in the field (autonomous mining vehicles, autonomous port container vehicles). However, robots are notably absent from many other sectors. Where are all the robots?

Without vision, a vast array of potential applications is closed to robots, for example: complex manual assembly, packing, manipulation, navigation, machine operation, fruit picking, crop spraying, remote assistance, smart homes, smarter appliances, autonomous driving, environmental surveying and monitoring. 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.

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 future challenges facing Australia including:

- Labour shortages and low productivity growth in key industries (by decoupling labour growth from population growth);
- 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; and
- Growing demands for minerals, energy and food that must be met in the context of the community's growing environmental expectations.

Concept drawing of the prototype broadacre agricultural robot, AgBot II, developed by researchers associated with the ACRV

Five application areas

We have identified five areas of economic importance to Australia that will benefit from robots with vision capabilities. Robots currently have little traction in many of these areas yet are potentially transformational. They will be important tools for improving productivity in a range of labour-intensive industries that are characterised by unstructured working environments such as; construction sites, hospitals, cities, and farms. Robots can perform useful and labour-intensive tasks like inspection, monitoring, measuring, assembly, and material transport in a range of environments (on land, in the air, in water). They can act as cost-effective and persistent environmental monitoring tools. There are also improvements that can be made to the way that robots operate on the factory floor to work more collaboratively with humans (smart manufacturing). The five areas we have identified include:

- **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.

Unmanned Aerial Vehicles (drones) are used to inspect infrastructure assets and identify maintenance requirements

Technology Demonstrators

Our research is grounded by the requirements of end-users and we will create a series of technology demonstrators for our five application areas. The technologies developed by the Centre will deliver important economic, environmental and social benefits for Australia.

Each of our research projects (apart from SV1, which is fundamental underlying research) will contribute to application demonstrators in the five application areas. The Centre aims to attract an additional $10 million in investment for prototype development and other end-user deliverables over its lifetime. Industry interaction will be encouraged through many mechanisms including industry forums and invitations to see technology demonstrators. Where significant industry demand is identified, the Centre will run continuing professional education courses in robotics, vision and/or robotic vision and has already made a start to this end by developing a series of massive open online courses covering these topics (see MOOCs p. 47).

Smart Manufacturing

Manufacturing Portfolio Lead Tom Drummond

Tom Drummond

The Australian manufacturing sector continues to represent a core source of economic prosperity for Australia – and should 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 6.5% of Australia's GDP (Gahan, The Conversation http://tinyurl.com/oz9a5xc). 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 (Pickering, Business Spectator http://tinyurl.com/kakfuhu)

Manufacturing robot technology is quite mature, with the first manufacturing robot going to work in 1961 in a GM plant, unloading die casting machines. The use of advanced industrial robots is at hand, with factories set to begin using automation at an increasing pace. For example, Amazon's robot-based business and Rethink Robotics' human-safe robot will revolutionise small batch manufacturing for SMEs. This change could boost productivity in some countries by 30 percent, according to a new study by Boston Consulting Group (http://tinyurl.com/qetyfts), which also predicts that investment in industrial robots will accelerate markedly over the next decade, from annual growth that now averages 2 to 3 percent to around 10 percent. By 2025, this would result in the productivity improvements of 10-30% in the world's 25 largest goods-exporting nations.

Our research will contribute to developing advanced robots that can use their sense of vision as smart co-workers, watching what human workers do and anticipating what comes next. To do this, robots need to have much improved levels of sensory ability than currently available.

Infrastructure monitoring and asset inspection

Infrastructure Portfolio Lead Rob Mahony

Rob Mahony

Australia has a vast and aging network of infrastructure assets that require inspection and monitoring. We need robotic vision systems that can see what needs to be seen, analyse those images and provide condition reports on the health of these assets. 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 and many have online monitoring as well. However, other 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. Safety and environmental incidents can be avoided, thus increasing the reliability of important infrastructure assets and potentially extending their life.

Agriculture, aquaculture and bio-monitoring

Agriculture Portfolio Lead Tristan Perez

There is a labour vacuum in food production. Employment in agriculture, which used to provide almost all the jobs in the pre-modern era, now accounts for only 2% of employment in the first world. With a world population projected to reach almost 9 billion by 2050, sustainability and food security worldwide are significant challenges. Australia, in particular, faces a real challenge to ensure its participation in food production is both competitive and sustainable. We need robots to enhance agricultural productivity and to achieve sustainable production of food with less land, lower inputs and fewer famers. 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 Agriculture and Fisheries (QDAF), and it is part of the Australian Centre for Robotic Vision at QUT. Such technology can be used to conduct autonomous multi-vehicle operations in applications of weed management, fertilising, and seeding. The use of multiple, relatively low cost, field robots can enable novel alternative weed destruction methods based, for example, on mechanical and thermal principles rather than herbicides. Such advances could reduce the input cost for weed management in terms of energy, labour, and improved chemical delivery by up to 40%. The robots will also have the capability to communicate, not only with each other, but also with 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 can be 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.

Tristan Perez

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.

Improving Food Productivity

Indian Prime Minister meets AgBot II during the G20 Summit

In November 2014, Brisbane hosted the G20 Leaders Summit, the principle forum for international economic cooperation. An annual event, the G20 is an opportunity for world leaders to meet and share ideas to address key international issues.

Among the dignitaries in Brisbane was the tech-savvy Indian Prime Minister Narendra Modi, who visited QUT to meet with Centre researchers working in agricultural robotics.

With the agriculture industry in transition, many modern farmers are using digital networks, sensors and autonomous devices for many agricultural functions including planting, harvesting, sorting, packaging and boxing.

Lead researcher, AI Tristan Perez says that the farm of the future will be supplied with lightweight, small, autonomous and energy efficient machines called AgBots. Each AgBot will perform a particular task but will work together to weed, fertilise, control pests and diseases, all working while collecting valuable data that can be analysed to improve farming practices.

"We are starting to see automation in agriculture for single processes such as animal and crop monitoring using unmanned aerial vehicles (drones), robotic weed management, autonomous irrigation," he said. "There is enormous potential for AgBots to be combined with sensor networks and unmanned drones to provide farmers with large amounts of data, which then can be combined with mathematical models and novel statistical techniques (big data analytics) to extract key information for management decisions — not only on when to apply herbicides, pesticides and fertilisers but how much to use and where."

Peter Corke, Director of the ACRV, and Tristan met with Prime Minister Modi to brief him on the latest development, the second generation AgBot II.

AgBot II is being developed for broad acre weeding and fertilising, and is part of a larger project funded by QDAF.

Tristan told Mr Modi that the new AgBot II would autonomously seed, weed and fertilise with trials starting in 2015.

"We have done an economic analysis for weed management that anticipates a reduction of costs of up to 40 per cent in terms of energy, labour and chemical savings," Tristan said.

Looking to ensure India's future food security Mr Modi is a strong advocate and supporter of agricultural research and development, particularly in the area of agricultural technologies.

Asked to write a message on the robot, Mr Modi wrote (translated from Gujarati):

"Research is the mother of invention. The development journey of mankind is a continuous stream of research. Science and technology is very important for agriculture. Agricultural scientists rely heavily on science and technology. Agricultural progress is very important for the welfare of humanity. I congratulate you for your efforts here and wish you the best. Narendra Modi"

Indian Prime Minister, Narendra Modi, writes a message on AgBot II

Building and construction

Construction Portfolio Lead Ian Reid

Robots with vision have strong potential to transform the building and construction industry, which has significant economic importance to Australia. The industry is characterised by a large (expensive) labour force with unskilled components, which raises issues such as safety, productivity, and quality control. Building sites appear quite chaotic, with building materials in constant flux. The challenge is getting inventory just in time and getting it to where it needs to be. The work environment is highly unstructured, which is problematic for current generation robots. There are many useful tasks robots could perform such as inspection,

Ian Reid

Peter Corke and Sue Keay visit Kane Construction's building site at Springfield to look at applications for robotic vision

measuring, assembly, material transport in a range of environments. Some construction robots do already exist, e.g., the bricklaying robot, the semi-automated mason (SAM). There is potential for increased automation using the sensing we are developing in the Centre to operate robots autonomously or semi-autonomously in ways that help humans or replace unskilled labour. For example this could be embodied in smart cranes, smart pallet trucks, smart scissor lifts that can work with the humans on site to make the industry more productive. Some areas where robots with vision may have an impact in the construction industry include:

- safety - looking for trip hazards, people entering no-go areas, or entering areas where machines are operating;
- dynamic mapping of the scene - mapping against plans over time, looking at both the geometry of buildings under construction as well as identifying objects within the map;
- augmented reality - could be used on building sites by providing virtual views of scenes to control placement of materials on site, perhaps guiding cranes;
- autonomous work platforms and vehicles - operating on building sites to follow workers with the tools they require, to move materials over night, to clean building sites, even to conduct demolition work.

Medical and healthcare

Medical and Healthcare Portfolio Lead Jonathan Roberts

There is a huge need for improved productivity in the health sector. Hospital operations require significant amounts of transport of patients, drugs, meals, linen, samples and so on. Already there are some applications of robotics to transport within hospitals, e.g., Aetheon's Tugs, which take hospital trolleys where they are needed, rather than relying on porters. Another area where robots can improve productivity is in assisting surgeons. The demand for Minimally Invasive Surgery such as arthroscopy is increasing around the World, restricted by the supply of specialised surgeons. Presently surgeons use their own vision to manipulate surgical robots to perform certain procedures. This requires the surgeon to perform like a "one-man-band" where they use both hands and feet to operate the robot. We are seeking to develop a robotic tool to take care of manipulation tasks while the surgeon takes care of the decision-making. In particular, the research will develop robotic vision systems that are capable of mapping joints in real-time via arthroscopically sourced video streams. The 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).

Jonathan Roberts

ACRV AI Jonathan Roberts (left) with members of the Medical and Healthcare robotics group, Ross Crawford (centre) and Anjali Jaiprakash (right) are looking at how robotics can be applied in arthroscopy

National Benefit

National Research and Innovation Priorities

The technologies developed by the Centre will deliver important economic, environmental and social benefits for Australia. Robots 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 which 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.39) .

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 the nation.

The Centre clearly addresses the Nation's Strategic Research 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 NICTA and DSTO 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.

ACRV will train the next generation of experts in robotics and vision who will contribute to the nation's strategic research priorities

Strategic Research Priorities (SRPs)

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

Challenge	Strategic research priorities
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

National Benefit KPIs

National Benefit			
Performance Measure	**Reporting Frequency**	**Target 2014**	**Outcome 2014**
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%
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
Number of fellowships		2	3
International grants		0	1

Communication, Engagement and Outreach

The ACRV engages with industry, business, government and the community through its Communication, Engagement and Outreach Program (CEOP). The CEOP conducts and delivers a number of engagement activities including forums, workshops and presentations. We actively promote the Centre through traditional media channels targeting local, national and international news and current affairs programs and outlets. The Centre also promotes its activities online and through social media channels including Twitter, Facebook, Google+, LinkedIn, Flickr, Instagram and YouTube. The ACRV 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 and datasets.

The general community has an enduring fascination with robots. Technology demonstrators, a key methodology for the Centre's 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, workshops for high school teachers, and talks in schools by project researchers will be encouraged. Our annual report will present the achievements of the Centre and be available in electronic form and distributed in hardcopy to selected end-users. Our research engagement 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) will be held at ANU's Kialoa campus in March 2015.

During 2014 the Centre delivered 79 government, industry and business community briefings and 40 public awareness and outreach programs. The Centre also developed "Introduction to Robotics" and "Robotic Vision" MOOCs to be delivered in early 2015. 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 (see Engagement via MOOCs p. 47).

The Centre promotes internal communication by a number of means. We have adopted a wiki-style intranet tool called Confluence, developed by Australian company, Atlassian. The Centre's Executive meets fortnightly

During 2014 the Centre delivered 79 government, industry and business community briefings and 40 public awareness and outreach programs.

via Google Hangouts. The Centre has a strong cross-platform social media presence. Research Fellows and PhD Candidates each have monthly Skype or Google Hangout meetings. Centre Director, Peter Corke keeps in regular communication with CIs via global emails (42 in 2014) and to promote Centre achievements for 2014 the Centre produced 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.

ACRV's mascot, Frosty, a Nao robot donated by Aldebaran Robotics, enjoys reading and has his own Instagram account @frosty_nao

Workshop held as part of RoboVis2014, ACRV's annual symposium, held at ANU Canberra in November 2014.

Communication, Engagement and Outreach Program Highlights

21 February 2014	Showcase of agricultural robotics work to RDCs
11 March 2014	Presentation to QUT's CEO forum on robotic vision and ACRV
April 2014	Superhuman robot navigation with a Frankenstein model…and why I think it's a great idea at the Hong Kong Convention Centre (2000 people)
14 April 2014	QUT Agricultural Robotics Workshop 1 - 2014 Agricultural Practices and Weed Management
24 April 2014	An Introduction to Artificial Intelligence at the Girls in ICT Day (ANU) (60 people)
26 June 2014	Agricultural Robotics at Queensland Agricultural Conference, Pullman Hotel (100 people)
16-21 June 2014	Machine Learning Summer School (MLSS) at Beijing, China (300 people)
28 June 2014	Tutorial on Inference and Learning in Discrete Graphical Models at IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (20-40 people)
30 June 2014	Robotics in Agriculture at RAS Winter School (QUT) (30 people)
July 2014	ACRV announced as successful in bid to host ICRA, IEEE's International Conference on Robotics & Automation (>2000 delegates expected)
28 July 2014	Agricultural Robotics at GRDC updates - Wellington NSW (30 people)
29 July 2014	Agricultural Robotics at GRDC updates - Spring Ridge NSW (70 people)
14 August 2014	Creative Industries and roboticists working together at QUT (30 people)

15 August 2014	FUTURESHAPERS: Rising stars of science research at QUT Room 360 (200 people)
19 August 2014	Melbourne GPU Users Meetup (50 people)
September 2014	Making Machine Learning Methods in Medical Image Analysis more Efficient and Less Dependent on Large Training Set at Technical University of Munich, Germany (50 researchers/students from the Technical University of Munich)
2 September 2014	ConstructionQ: Robotics in Construction at Brisbane Convention and Entertainment Centre (300 people)
3 September 2014	Robots and Performance Art (by Jonathan Roberts AI) at QUT (60 people)
9 September 2014	Guest seminar for University of Washington students at QUT (40 people)
15 September 2014	Plenary presentation on "The Quest for Robotic Vision" at IROS conference, Chicago (1500 people)
24 September 2014	Keynote address on Autonomous Systems at Future Land Force Conference, Brisbane Entertainment and Convention Centre
30 September 2014	Melbourne MathSci Meetup (20 people)
October 2014	Robotics Workshop for visiting year 10 students at Monash University (100 people)
16 October 2014	US Ambassador's Innovation Roundtable at QUT (150 people)
23 October 2014	Disruptive Innovation in Agricultural Practices at Queensland's Royal Agricultural Society (120 people)
24 October 2014	Grand Challenge Lecture - The UAV Challenge at QUT Kindler Theatre (100+ viewers on live stream)
November 2014	Making Machine Learning Methods in Medical Image Analysis more Efficient and Less Dependent on Large Training Set at Instituto Italiano di Tecnologia, Genoa, Italy (50 researchers and students from the Instituto Italiano di Tecnologia)
November 2014	Presentation to the Indian Prime Minister Modi for Agricultural Robotics at G20 Summit, QUT Science and Engineering Centre (100 people)
13 November 2014	Future of Robotics at Joint Industry Plumbing Conference 2014, Royal Pines Resort Gold Coast (150 people)
14 November 2014	G20 Visit to QUT by Indian PM at The Cube (extensive media coverage)
28 November 2014	Robotics and Infrastructure at Queensland Urban Utilities, Cauldron Innovation Series (25 people)
1 December 2014	Robotic Vision and Autonomous Systems in Complex Contexts at Defence Science Institute, Melbourne
9 December 2014	Building and using prior shape models for 3D tracking and SLAM at International Conference on 3D Vision (3DV2014), Tokyo (keynote presentation) (200 people)
10 December 2014	Towards Semantic Vision SLAM at International Conference on Control, Automation, Robotics and Vision (keynote presentation) (250 people)
11 December 2014	Robotic Vision at ShanghAI14 Lecture

Media

ACRV Media Coverage

Media Pieces			
Name	**Date**	**Distributor**	**Content**
Michael Milford	9-Dec	ABC news 24	Michael Milford explained to ABCNews24 how Nobel Prize-winning research into the brain could prove helpful for Alzheimer's and robotics research *External Link: https://www.youtube.com/watch?v=XmsL3YeI9yg*
Michael Milford	29-Sep	ABC news 24 & The Australian	The Australian ran Michael Milford's research into how rats are helping robots 'see' on the front page of its IT section *External Link: http://www.theaustralian.com.au/technology/rat-brains-enlisted-in-quest-for-spatial-certainty/story-e6frgakx-1227059462988*
Michael Milford	22-Sep	Channel Nine News	Michael Milford was on Nine News encouraging students looking for a stable and interesting career to consider robotics.
Jonathan Roberts	Sep	BBC International News	UAV Challenge: 6-minute piece on BBC International News (2014) – Audience of between 80 million and 300 million
Jonathan Roberts	29-Sep	4BC, South Burnett Times & ABC Radio	4BC, the South Burnett Times and ABC Radio widely covered the UAV Outback Challenge, co-organised by Professor Jonathan Roberts. *External Link: http://www.abc.net.au/local/photos/2014/09/24/4093623.htm?site=gippsland*
Tristan Perez	30-Jul	ABC	Tristan Perez sees a future where swarms of robots could independently control weeds in crops. *External Link: http://www.abc.net.au/news/2014-07-30/rural-nsw-weedblastingdaleks-2907/5632774*
Tristan Perez	20-Sep	ABC	Tristan Perez says robots will help agriculture become more productive *External Link: http://www.abc.net.au/news/2014-09-02/tristan-perez-ag-bot/5713938*
Tristan Perez	30-Jul	ABC	Could weed-zapping Daleks be the new farmer's friend, with robots roaming crops? *External Link: http://www.abc.net.au/news/2014-07-30/nrn-weedzapping-daleks/5633014*
Tristan Perez	2-Sep	ABC	As the global population increases, robots could help feed the world by lowering labour costs and using farm inputs more efficiently *External Link: http://www.abc.net.au/news/2014-09-02/sach-agricultural-robots-tristan-perez/5713812*

Tristan Perez	15-Jul	GRDC Media	Robots in the paddock and drones in the air – a different future in grain production is just one of the fascinating topics being discussed at the upcoming Northern Grains Research and Development Corporation (GRDC) Grains Research Updates. *External Link: https://www.grdc.com.au/Media-Centre/Media-News/North/2014/07/Robots-and-drones-back-to-the-future-for-GRDC-Updates*
Tristan Perez	25-Jul	GRDC Media	Insights into future agricultural robotic systems *External Link: http://grdc.com.au/Research-and-Development/GRDC-Update-Papers/2014/07/Insights-into-future-agricultural-robotic-systems*
Wai-Ho Li	14-Oct	Monash Media	Monash PhD student Titus Tang and his supervisor Dr Wai Ho Li, from Monash Vision Group and the Department of Electrical and Computer Systems Engineering won the Best Paper Award at the 18th International Symposium on Wearable Computers in Seattle, a major conference for specialists in the field. *External Link: http://monash.edu/news/show/monash-duo-win-international-accolade-for-wearable-computer*
Peter Corke	28-Jan	Times of India	
Peter Corke	8-Apr	ABC Radio	Professor Peter Corke's expertise was sought by ABC Sydney's Drive program following the release of a productivity report about robotics.
Peter Corke	9-Apr	Courier Mail	Professor Peter Corke's research to give robots vision was reported on by the Courier Mail's Queensland Business Monthly magazine
Peter Corke	15-August	Courier Mail	Robotic Vision research is Idea #1 on Courier Mail's 25 Ideas to Change the World series *External Link: http://www.couriermail.com.au/news/special-features/ideas-to-change-the-world-part-1-unlocking-cool-tools/story-fnoufkj0-1227025125464*
Peter Corke	23-Oct	The Australian	Powering Australia: robots workers of the future *External Link: http://www.theaustralian.com.au/business/powering-australia/powering-australia-robots-workers-of-the-future/story-fnnnpqpy-1227123685149*
Peter Corke	14-Nov		Indian PM Mr Narendra Modi visits the Cube to see the AgBot
Peter Corke	2-Dec	Qweekend	Listed in "50 Best, Brightest"

Peter Corke		ABC Radio	Driverless cars
Peter Corke	15-Nov	Australian Weekend Magazine	“My mechanical friend”, feature on robotics, including online supplement
Peter Corke		Robohub	2x interviews
Peter Corke		IJARS	Interview
Peter Corke		NOVA TV	Interview
Peter Corke	5-Aug	ABC 612 (Tim Cox)	“Autonomous cars”
Peter Corke	18-Mar	QUT media	“New centre will give robots the gift of sight”
Peter Corke	18-Mar	The Australian	“Robots given a watching brief”
Peter Corke	8-Apr	ABC 702 Sydney	Comment on SMH article re: report by the Australian Workforce and Productivity Agency
Peter Corke	9-Apr	Courier Mail Business Supplement	Interview with Harry Clarke

ACRV garnered plenty of media attention in 2014, camera crews in action getting robot footage

Engagement via MOOCs

Robots and Vision – free and online for everybody in 2015

ACRV Director Peter Corke, has developed the world's first massive open online courses (MOOCs) on robotics and robotic vision, designed for undergraduate engineering students but suitable for anyone with a strong interest in robotics, sufficient mathematical and programming knowledge, and willing to learn. The courses were developed by ACRV's administering organisation QUT.

MOOCs are free, open-access courses delivered online to an unlimited number of participants worldwide. Very few robotics MOOCS have been offered by institutions in the past - and all were designed for postgraduate-level students undertaking science, technology, engineering and mathematics (STEM) related research.

"While the MOOCs might attract some high school STEM stars and skilled armchair roboticists, I expect most of the students will be undergraduates, perhaps studying engineering or computer science at a university that doesn't itself have a strong robotics program," said Peter, the course creator, world-renowned roboticist and the Director of the ARC Centre of Excellence for Robotic Vision.

"It could also be helpful for STEM professionals looking to expand their skill set - with big players like Google, Apple and Boeing pouring billions of dollars into robotics and automation, it's an industry that'll be screaming for workers in the near future."

The Introduction to Robotics MOOC is designed to develop the fundamental mathematics and algorithm skills that underpin robotics, including representation of pose and motion, kinematics, dynamics and control. As an optional practical assignment, students with a LEGO Mindstorm kit will be able to build a simple robot arm and write the control software for it.

The Robotic Vision MOOC takes that knowledge a step further, introducing students to the rapidly developing field of computer vision, learning how images are formed, and fundamental algorithms to process images in a computer to extract information such as the colour, size, shape and position of objects in the world. As an optional practical assignment students can build an intelligent vision system that can recognise objects of different colours and shapes.

"If students did the first course and built the robot, they can connect the vision system to the robot to create a robot that can see and respond to objects in its environment," according to Peter.

Throughout the MOOCs students are quizzed to check their understanding and given weekly tests and programming assignments, which count towards the final grade. Those who pass receive a certificate of completion. Students in the MOOCs are supported by online discussion forums for sharing information and asking questions of both tutors and other students. While free, the courses are not easy.


Peter Corke takes his classes online with two new robotic vision MOOCs to be released in 2015

"They're certainly not going to be a walk in the park - both MOOCs involve theory, mathematics and programming," Peter said.

These courses herald QUT's first venture into MOOCs, in a project that was driven by Peter's passion for robotics and his sizable YouTube following.

"A few years ago I broke my knee, just days before semester started," Peter said. "I had to record my lectures at home and a colleague showed them to the class. At the end of semester I put them up on YouTube. I was really surprised by the interest those lectures generated - more than 70,000 views, with one lecture alone viewed over 30,000 times. That made me realise just how many people are genuinely interested in robotics and it got me thinking about how I might be able to deliver structured course content online. It's interesting to reflect on progress. Once upon a time we needed a lecture theatre and a lab full of hardware to teach robotics but in this digital age we don't always need that resource-intensive, bricks-and-mortar model to deliver a strong robotics course. These days we can teach mechatronics with LEGO kits at home rather than in labs, and I find that a truly exciting prospect."

Introduction to Robotics will open on February 15, 2015 with a "Getting Started" period followed by six weeks of instructional material and participant activities. Robotic Vision opens on April 13, 2015 with a "Getting Started" period followed by six weeks of instructional material and participant activities.

For more information:
www.moocs.qut.edu.au

Linkages

To be effective the ACRV needs 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. The Centre encourages and supports travel for chief investigators, research fellows, and PhD candidates to other Centre nodes, including overseas partner organisations, and to host visits by researchers from our overseas partners. We also 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. Within Australia ACRV holds an annual symposium, a three-day event that includes presentations from CIs, overseas PIs, students; poster presentations; and invited industry guests. Our annual symposium will be open to Australian researchers from non-partner universities, other research organisations and end-users. With our partner NICTA we will cooperate in areas such as graduate education and exchange of researchers, sharing of datasets, and co-sponsoring of international visitors. We are also working with DSTO to define research problems of mutual interest in robotic vision and have already hosted visits and presentations by DSTO researchers.

Our international reputation will be 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. The Centre will enhance international linkages through the effective and ongoing research collaborations that exist between the Australian and international partners. Meaningful visits between the laboratories for the investigators, research fellows and PhD candidates are already being planned, creating a transnational research community around robotic vision.

As part of building the Centre's profile, we are organising workshops on Robotic Vision at the main conferences such as International Conference on Robotics & Automation (ICRA), International Conference on Computer Vision (ICCV) and International Conference on Computer Vision Systems (ICVS). Already in 2014 robotic vision has been profiled at the international level (see p. 50). In 2015, ACRV researchers have had a number of robotic vision workshops accepted for presentation at the top robotics and computer vision conferences. In May, Michael Milford, Niko Sünderhauf and Peter Corke will be running dual Place Recognition workshops at CVPR2015, ICRA2015 and ICVS2015, the premiere large robotics and computer vision conferences. In addition, Tom Drummond and Peter Corke will be running a workshop at ICRA on Challenges to Robotic Vision, while Ian Reid and Stephen Gould will run a workshop on Semantics for Structure from Motion and Visual SLAM at CVPR. These workshops follow the very well attended inaugural Place Recognition workshop Milford and Corke ran at ICRA2014. Guest speakers at the workshops will include: computer vision and visual neuroscience expert David Cox from Harvard University; robotic mapping expert, and Centre PI, Paul Newman from Oxford University; computer vision expert 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.

Seminar presented at ACRV by Daniela Rus, Director of the Computer Science and Artificial Intelligence Laboratories at the Massachusetts Institute of Technology (MIT)

International and National Links & Networks KPIs

International, National, and Regional Links and Networks			
Performance Measure	**Reporting Frequency**	**Target 2014**	**Outcome 2014**
Number of international visitors and visiting fellows	Annually	2	8
Number of national and international workshops held/organised by the Centre	Annually	1	1
Number of visits to overseas laboratories and facilities	Annually	2	11
Examples of relevant interdisciplinary research supported by the Centre *Number of interdisciplinary projects within the Centre	Annually	0	0

End-User Links KPIs

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

Examples of Linkages

The Centre received international recognition in 2014 with presentations on robotic vision being made at significant international conferences. These included the International Conference on Intelligent Robots and Systems (IROS) and the International Conference on 3D vision (3DV).

IROS 2014 was held in Chicago, Illinois and sponsored by the Institute of Electrical and Electronics Engineering (IEEE) Robotics & Automation Society, the Robotics Society of Japan, the Society of Instrument and Control Engineers, the New Technology Foundation, and the IEEE Industrial Electronics Society. The aim of IROS is to bring together a diverse and multidisciplinary group of researchers interested in creating intelligent robotic systems and it featured research streams from the fields of Artificial Intelligence and Robotics. IROS 2014 featured three plenary speeches, including one from Centre Director Professor Peter Corke and another by PI Andy Davison; 39 session keynotes by leaders in the field; a vibrant industrial exhibition and talks from sponsors; special forums and panels on industry and entrepreneurship and government policy as it relates to robotics. 750 papers and 27 workshops and tutorials were selected for the final program, with authors from nearly 50 countries from around the world.

Peter's plenary presentation at IROS was on the "Quest for Robotic Vision", and discussed the history of robotics and vision, examined the state of the art, and discussed what may happen by bringing vision back to robotics in the future.

The International Conference on 3D Vision provides a premier platform for disseminating research results covering a broad variety of topics in the area of 3D research in computer vision and graphics, from novel optical sensors, signal processing, geometric modeling, representation and transmission, to visualisation and interaction, and a variety of applications. ACRV Deputy Director, Ian Reid, was one of only five keynote speakers invited to 3DV in Tokyo, Japan in 2014 and presented his work on, "Building and using prior shape models 3D tracking and SLAM". In addition to his keynote presentation, Ian and colleagues from Oxford won the 3DV Best Paper Award for their paper titled, “3D Tracking of Multiple Objects with Identical Appearance using RGB-D Input”.

The ACRV will continue to promote robotic vision in major international conferences and conventions in 2015 and beyond and has won the right to host ICRA, IEEE's International Conference on Robotics and Automation, in Brisbane in 2018. 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.

To find out about the Quest for Robotic Vision:

http://robohub.org/robots-podcast-quest-for-computer-vision-with-peter-corke/

Dr Andre Barczak from Massey University, New Zealand, visited the ACRV for six weeks in 2014

ACRV Director Peter Corke was invited to present a plenary, "The Quest for Robotic Vision", at IROS 2014 in Chicago

People

The 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). The team includes six IEEE Fellows; ARC Laureate, Future and DECRA Fellows; and two Microsoft Research Faculty Fellows.

To this group of talented individuals, the Centre has actively recruited five research fellows at QUT, two research fellows at UoA and offers have been made to research fellows to commence at Monash and ANU in 2015. We also have our first PhD candidates, with four at QUT and two at ANU, with more to follow in 2015. Two new Associate Investigators joined the ACRV in 2014 and have already taken active roles in the Centre's leadership, Professor Tristan Perez (Portfolio Lead Agricultural Applications) and Professor Jonathan Roberts (Portfolio Lead Medical and Healthcare Applications). More information about all of ACRV's people follows in our profile sections (see p. 59).

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	15	-	0 : 15
Professional Staff	2	2	2 : 0
Research Fellows	6	6	1 : 5
PhD Candidates	6	6	1 : 5

Research Training

The Centre recruited six PhD students in 2014; Zetao "Jason" Chen, Zongyuan "Tony" Ge, Adam Jacobson, Fahimeh Rezazadegan, Ilya Magurov, Yi "Joey" Zhou (see new PhDs p. 56). The Centre will train 24 postdoctoral fellows and 80 PhD students over its lifetime. To support this, the Centre will have a strong program in undergraduate and graduate training, in which all CIs (and where appropriate PIs) will be involved. At the undergraduate level, the Centre will create material that can be incorporated into programs in the fields of Artificial Intelligence, Computer Vision, Robotics, Distributed Computing and Feedback Control. This will include MOOCs as well as textbooks, and will build on the work of Centre Director Peter Corke who is currently creating two MOOCs (see p. 47). Educational hardware, firmware and software, developed within the Centre will be open sourced allowing undergraduate courses to develop state-of-the-art laboratory experiments. Honours students will be actively recruited to work on research projects within the Centre for thesis projects and as summer scholars. The Centre will run a national competitive process to select honours students who have an interest in PhD study to attend our annual Robotic Vision Summer School (RVSS) for graduate students. These undergraduate initiatives are important tools for post-graduate student recruitment and to develop the researchers of the future. With major nodes in the Centre across four capitals in Australia it is expected that the educational impact, over the full seven-year life of the Centre, will be significant.

The Centre will expose students to leading international researchers in the field who visit the Centre's nodes or by students travelling to the overseas partner nodes – overseas research experience is a valuable part of research training, which facilitates research and knowledge transfer and establishes strong international personal networks. ACRV CIs have a demonstrated track record in research training and will mentor their PhD students and postdoctoral fellows.

Research Training and Professional Development KPIs

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

Gender Diversity

The ACRV investigators are all male, and this reflects gender balance issues that are acute in computer science, and only slightly better in engineering. The Centre is developing strategies to remediate this as much as possible and has appointed CI Ben Upcroft to be the Gender Diversity Lead for the Centre and its Male Champion for Change. The ACRV is looking at ways to address unconscious bias in the advertising and recruitment of research fellows and PhD candidates, 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 women's group will be formed based on the fledgling group that meets every 6 weeks at ACRV's QUT node.

Leadership Development

As part of the human capital development and planning the Executive Committee has identified successors for the key roles of Director and Research Theme leaders. These individuals will be mentored in areas of leadership and take on responsibilities such as participating in Executive Committee meetings when the incumbent is on leave. Each Theme Leader, and the Director, has a nominated deputy who fulfils their role during short-term absences. The Executive Committee has also appointed CIs to portfolio leadership roles (supported by the operations and administration group) to achieve the Centre's goals (see Governance p. 79).

Mechatronics undergraduate student and Nao specialist Amy Gunnell with QUT's Teacher-in-residence, Anne Brant, and blue Nao robot

Mentoring

All ACRV researchers will be appointed a mentor located at another node. The goal is to provide all ACRV researchers with a mentor who will act as a trusted advisor. Mentors will be selected from a node that the mentee intends to collaborate with. Mentoring relationships are flexible and intended 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.

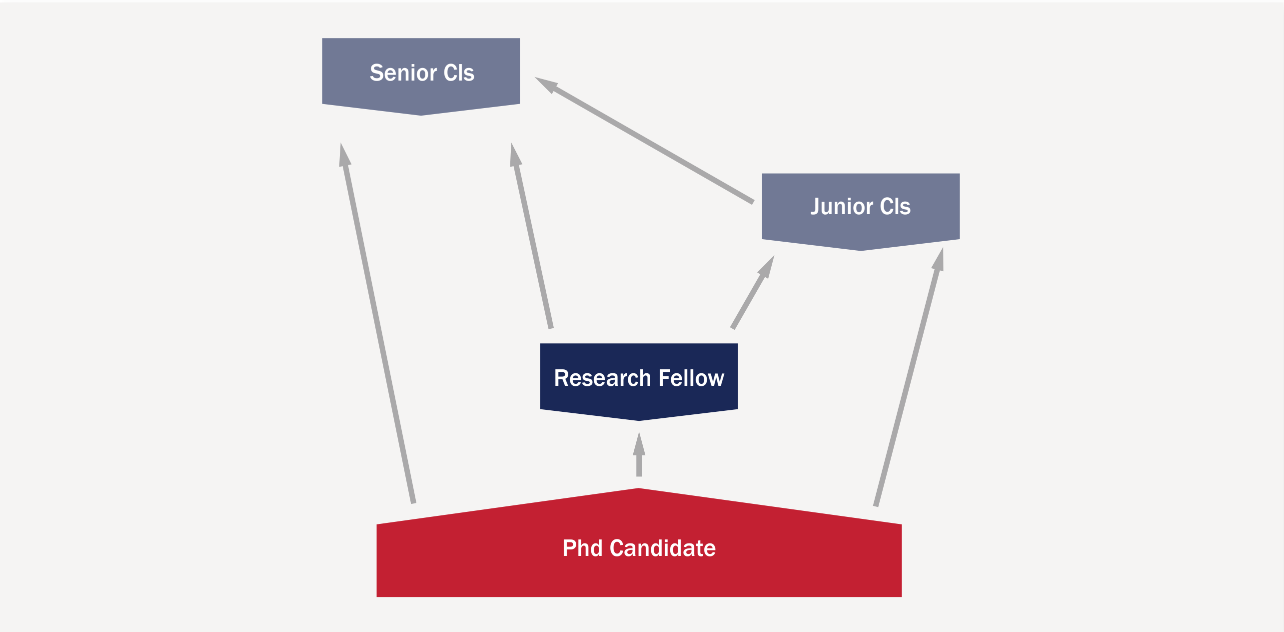

NEW PhDs

There's something special about being part of an enterprise right from the beginning. We are welcoming new people to the Centre all the time. Of these new people, here are the profiles of three of our new PhD candidates, who started with the Centre in 2014.

Fahimeh Rezazadegan

PhD title: "Vision-based Human Action Recognition Using Temporal Information".

Fahimeh received her Bachelors of Science in Electronic Engineering from the Isfahan University of Technology (IUT), Esfahan, Iran in 2006.   After completing her degree she worked at Behineh niroo Espadan Engineering Company and the IUERC (Isfahan University Engineering Research Center) until 2014 where she led the design and simulation of electronic circuits and developed microcontrollers for monitoring electronic devices. She continued her study in control engineering and obtained a Masters Degree in Iran in 2013. Her research area was the problem of trajectory tracking control of mobile robots. She focused on design of the adaptive back-stepping controller for an autonomous underwater vehicle, AUV, in the presence of parametric uncertainties. Fahimeh is currently working on deep learning methods using vision for human action recognition in a mobile robot. She is based at QUT and her principal supervisor is Ben Upcroft and associate supervisor is Michael Milford.

Zetao "Jason" Chen

PhD title: "Brain-based sensor fusion for navigating robots"

Zetao obtained his bachelors degree in Electric Engineering in 2009 from South China University of Technology and his masters degree in Artificial Intelligence in 2012 from University of Groningen. He is currently pursuing his PhD at Queensland University of Technology under the supervision of Michael Milford, Gordon Wyeth and Peter Corke. In his PhD, he is working on using deep learning techniques, metric learning and biology-inspired methods to perform place recognition over changing environments. Zetao was awarded the Huygens Scholarship from the Dutch Government for his masters study, and received the Best Paper Award from the Australasian Conference on Robotics and Automation (ACRA) in 2013 and the Best Paper Finalist from the Australasian Conference on Robotics and Automation in 2014. He has published (and co-authored) papers in ACRA, ICRA, IROS, ISRR and JFR.

Yi "Joey" Zhou

PhD title: "Dense Visual Odometry in Dynamic Environment"

Joey is a PhD candidate at CECS, ANU. He is doing research under the supervision of Hongdong Li and Laurent Kneip. His research interests are Visual SLAM and Stereo Matching. Joey studied a Bachelor of Mechanical Engineering and Automation, BeiHang University (BUAA), China in 2012. Before he came to ANU, he did research as a PhD candidate in iTR Lab, Institute of Robotics, BUAA, supervised by Prof. Tianmiao Wang. His publications include: "Bearing-only Visual SLAM for Small Unmanned Aerial Vehicles in GPS-denied Environments" International Journal of Automation and Computing 10(5), October 2013; "Research on Monocular Visual FastSLAM for a Small Unmanned Helicopter"; High Technology Letters, 2013; "Structural Target Recognition Algorithm for Visual Guidance of Small Unmanned Helicopters", 2012; IEEE International Conference on Robotics and Biomimetics, ROBIO, "An Implementation of RPV Control System for Small Unmanned Helicopters", 2012; IEEE International Conference on Robotics and Biomimetics, ROBIO; 'Model Adaptive Gait Scheme Based on Evolutionary Algorithm" 2013; IEEE Conference on Industrial Electronics and Applications (ICIEA).

Centre Culture

As we build the Centre we are also all contributing to a Centre culture, with a set of pervasive values, beliefs and attitudes that characterise the Centre. At the CI retreat in July 2014, everyone put forward some ideas about the Values of the Centre, which we collected in the following Wordle and have distilled in the Values listed at the beginning of this report.

Values are the essence of ACRV's identity. They are deeply held beliefs that shape the way the Centre sees the world and how we all act. Our values will be timeless, sustainable in the longer term, and reflect the values of our people. While the Centre's strategy may change through time our core values will remain the same. Values underpin our culture, brand, and strategy, forming the framework for Centre decision-making. Our values will continue to be developed collaboratively by all Centre personnel to serve as an affirmation of the intrinsic and unique value of each member of the ACRV community and as a guide to the way we interact with others.

As people join the Centre, as well as personal greetings, they also receive a welcome kit containing information about 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. In this way we can support the full cycle of the innovation value chain.

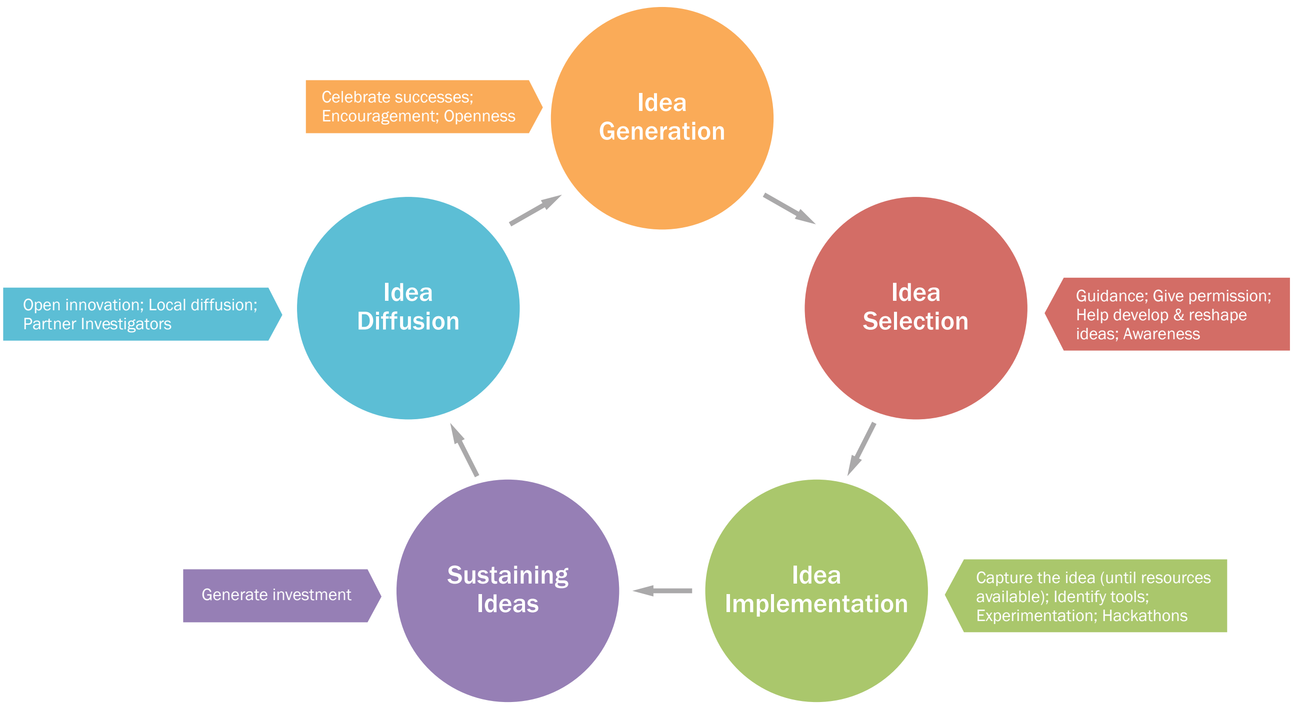

The ACRV Innovation Value Chain

CI Profiles

Peter Corke

Peter is the Director of the ACRV, leads our Robust Vision research theme, and is Professor of Robotic Vision in QUT's School of Electrical Engineering and Computer Science. His research spans topics including visual servoing, high-speed hardware for machine vision, field robotics, particularly for mining and environmental monitoring, and sensor networks. He is well known for his seminal papers in visual servo control, and work in field robotics, robotic vision and sensor networking. He has written two books: "Robotics, Vision & Control" (2011) and "Visual Control of Robots" (1997); developed the Robotics and Machine Vision Toolboxes for MATLAB; was Editor-in-Chief of the IEEE Robotics and Automation magazine (2010-13); was a founding editor of the Journal of Field Robotics (2006); is a member of the editorial boards of the International Journal of Robotics Research and the Springer STAR series; holds 5 international patents and is a Fellow of the IEEE (2007). He received a Bachelor of Engineering (Electrical), Masters and PhD, all from The University of Melbourne. (Over 14,000 citations and an h-index of 52).

Ian Reid

Ian is Deputy Director of the ACRV and leads our Semantic Vision research theme. Ian is also an ARC Laureate fellow, who joined the University of Adelaide's School of Computer Science in September 2012. Prior to his recent move to Adelaide, Ian led the highly successful Active Vision Group at University of Oxford (2005-2012). The group is recognised as the leading facility in the world for work in visual SLAM and one of the top groups for work in visual tracking and visual surveillance, resulting in a Best Paper award at CVPR 2008, and a seminal paper in visual SLAM with over 1059 citations. Ian went to Oxford as the Western Australian Rhodes Scholar in 1988, completing his DPhil in 1991 and went onto win a number of postdoctoral research grants and fellowships. His work received national UK attention in the mid-1990s when he and an Oxford colleague showed, through an early application of visual geometry analysis (prior to Hawkeye™), that the controversial third goal of the winning 1966 World Cup England soccer team, should not have been allowed. His research includes combining visual SLAM with radar to enable highly accurate mapping and surveillance of moving targets in a marine situation, which was used in security systems for the Weymouth sailing events at the 2012 Olympics. (Over 10,000 citations and an h-index of 48).

All citation data from Google Scholar

Rob Mahony

Rob is a member of the Centre Executive, leads our Vision and Action research theme and is a Professor in the Research School of Engineering at ANU. He is currently Director of the Research School and leads the Computer Vision and Robotics Group. Rob is a former Logan Fellow at Monash, and known for his work in sensor fusion, vision-based control and haptic interfaces between humans and robots. He obtained a science degree majoring in applied mathematics and geology from ANU in 1989. After working for a year as a geophysicist processing marine seismic data he 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 at 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. Rob's 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. (Over 6,000 citations and an h-index of 37).

Tom Drummond

Tom is a member of the Centre Executive Committee, leads the Algorithms and Architecture research theme, and is Professor in Engineering at Monash UniversityPreviously at Cambridge University, he is recognised for his breakthroughs in real-time vision architectures and has developed a number of widely used algorithms such as the FAST feature detector. Tom 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. His 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, and assistive technologies for visually impaired users as well as medical imaging. (Over 7,000 citations and an h-index of 33).

Chunhua Shen

Chunhua is an ARC Future Fellow (awarded in 2012) and Professor of Computer Science at UoA who will contribute expertise in Machine Learning to the challenges in the Semantic Vision theme. He studied at both Nanjing University and ANU, and received his PhD degree from The University of Adelaide. Chunhua worked at NICTA in their computer vision program for about six years. From 2006 to 2011, he held an adjunct position at the College of Engineering and Computer Science at the ANU. He moved back to Adelaide in 2011. As well as being the ACRV's deputy node leader for UoA, Chunhua is also a Project Leader (machine learning for robotic vision). He is also involved in the Data to Decisions CRC (D2DCRC), in particular on the projects of large-scale image classification and text analysis. His research interests are an intersection of statistical machine learning, large-scale optimisation, and computer vision. (Over 2,000 citations and an h-index of 21).

Hongdong Li

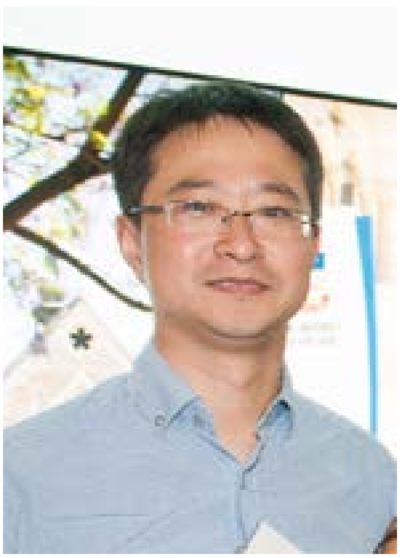

Hongdong is an Associate Professor at the Research School of Engineering in the College of Engineering and Computer Science at ANU. As well as being the deputy node leader for ANU, he is also contributing to both Robust Vision and Algorithms and Architecture themes. Hongdong was previously a Fellow with the RSISE at ANU and a Senior Research Scientist with NICTA Canberra Labs. He taught at Zhejiang University in China before moving to Australia. His main research interests are 3D Computer Vision, Image and Video Processing, Mathematical Optimisation, Pattern Recognition and Computational Visual Geometry. He was a Member of Bionic Vision Australia (BVA) and an Associate Investigator on the ARC funded Australia Bionic Eye Project special research initiative grant. Hongdong was supervisor and coauthor of a paper awarded the best student paper at ICIP2014 and has previously won the 2012 CVPR best paper award – the most prestigious in the computer vision community – for his work in 3D reconstruction. (Over 1,000 citations and an h-index of 18).

Ben Upcroft

Ben is an Associate Professor and the leader of QUT's Robotics and Autonomous Systems Discipline. He is also QUT's HDR portfolio lead and ACRV's Gender Diversity lead. 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 is an Associate Professor in the School of Electrical Engineering and Computer Science at QUT. 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). (Over 900 citations and an h-index of 14).

Stephen Gould

Stephen is an Associate Professor in the Research School of Computer Science in the College of Engineering and Computer Science at the ANU. He is the portfolio lead for higher degree research (HDR) students and will contribute to the Semantic Vision theme of the Centre. Stephen is formerly an ARC Postdoctoral Fellow and Microsoft Faculty Fellow and is a Visiting Researcher in the Machine Learning Group at NICTA. He has broad interests in computer and robotic vision, machine learning, probabilistic graphical models, and optimisation. Stephen's main research focus is on the application of machine learning techniques (specifically, conditional Markov random fields) to geometric and semantic scene understanding. He is also interested in seeing research being applied. Stephen co-founded Sensory Networks, a network security start-up company, and has developed freely available software libraries for machine learning and scene understanding. He holds eight international patents. (Over 1,500 citations and an h-index of 17)

Gustavo Carneiro

Gustavo is a Humboldt fellow, and a former Marie Curie International Incoming Fellowship (IIF) fellow who will contribute expertise in Deep Learning and classifiers to the Semantic Vision theme as well as being UoA's HDR portfolio lead. He is an Associate Professor of the School of Computer Science at the University Adelaide, joining UoA as a senior lecturer in 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. In 2005, he was a post-doctoral fellow at the the University of British Columbia with Professor David Lowe and at the University of California San Diego with Professor Nuno Vasconcelos. Gustavo received his Ph.D. in computer science from the University of Toronto under the supervision of Professor Allan Jepson in 2004. Gustavo has received a number of awards including: the 50,000th Marie Curie Fellowship Award (European Commission), Outstanding Reviewer Award (2010) for Computer Vision and Pattern Recognition (CVPR) and the Siemens Corporate Research Outstanding Achievement Award (2008). (Over 1,800 citations and an h-index of 18).

Michael Milford

Michael is an ARC Future Fellow, and ACRV deputy node leader, who builds on his award-winning work in robust vision for localisation (ICRA Best Vision Paper 2012) to contribute as a project leader in the Robust Vision theme with connections to the Semantic Vision theme. Formerly an ARC DECRA fellow and Microsoft Faculty Fellow, Michael holds a PhD in Electrical Engineering and a Bachelor of Mechanical and Space Engineering from the University of Queensland (UQ). After a brief postdoc in robotics at UQ, he 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 he became a Microsoft Faculty Fellow and lived in Boston on sabbatical working with Harvard and Boston University. His research interests include, vision-based mapping and navigation, computational modeling of the rodent hippocampus and entorhinal cortex, especially with respect to mapping and navigation, computational modeling of human visual recognition, biologically inspired robot navigation and computer vision, simultaneous localisation and mapping (SLAM). (Over 1,200 citations and an h-index of 19).

Gordon Wyeth

Gordon is Executive Dean of QUT's Science and Engineering faculty and a Professor of Robotics who contributes to the Semantic Vision theme of the Centre. He also provides connections from Semantic Vision into the Vision and Action theme, and contributes to the development of robotic systems in the Centre's Application areas.

Gordon holds a PhD and a Bachelor of Engineering degree (with honours) in Computer Systems Engineering. He is the President of IEEE Control Systems, Robotics and Automation Queensland chapter, former president of the Australian Robotics and Automation Association and has served in various leadership positions in the RoboCup International Federation. He serves in various editorial positions for leading international robotics journals and conferences. Gordon's team has designed and constructed more than twenty types of robots, including flying robots, wall-climbing robots, high performance wheeled robots, legged robots, manipulators and a humanoid robot. His robot soccer team, the RoboRoos, have been runners-up three times in the RoboCup World Cup of robot soccer. Gordon's research is internationally recognised for building practical and useful robots that exploit, explain and expand models of living systems. (Over 2,000 citations and an h-index of 22).

Richard Hartley

Richard is renowned as the founder of the field of multi-view geometry in computer vision – his text has received over 12,500 citations. He is contributing to the Centre's RV1 and SV1 projects. Richard has been at the ANU since January 2001. He is also the Program Leader for the Autonomous Systems and Sensor Technology Program of NICTA. Richard worked at the General Electric Research and Development Center from 1985 to 2001. He became involved with Image Understanding and Scene Reconstruction working with GE's Simulation and Control Systems Division. This division built large-scale flight-simulators. Richard's projects in this area were in the construction of terrain models and texture mosaics from aerial and satellite imagery. From 1995 he was GE project leader for a shared-vision project with Lockheed-Martin involving design and implementation of algorithms for an AFIS (fingerprint analysis) system being developed under a Lockheed-Martin contract with the FBI. This involved work in feature extraction, interactive fingerprint editing and fingerprint database matching. In 2000, he co-authored (with Andrew Zisserman) a book for Cambridge University Press, summarising the previous decade's research in this area. (Over 34,000 citations and an h-index of 57).

Anton van den Hengel

Anton is a Professor of Computer Science and the Director of the Australian Centre for Visual Technologies (ACVT) at The University of Adelaide. He will contribute his experience in computer vision to the Semantic Vision theme and the Human Machine Interaction project. Anton conducts research in a variety of areas within Computer Vision, ranging from fundamental mathematical analysis to multiple technologies that have been successfully commercialised. Industry research collaborators include Google, Microsoft, Holopoint, MonkeyStack, Landmark, Carbon Planet, Apogee, Rising Sun Pictures, BAES, BHP Billiton, Sola Optical, Champion Data, Tenix, Sydac, and Bayer Cropscience. He has had IP licensed to companies, sold outright, patented, and released into the public domain. He has formed 2 start-up companies to commercialise his research, both of which have been successful thus far. He has received a series of awards recognising the quality of his research including the Pearcey Award for Innovation in ICT, 2010, the Innovic Next Big Thing People's Choice Award 2009 for the most exciting innovation (awarded for Videotrace), and the Research and Development category of the South Australian iAwards, 2010 (for Videotrace). The iAwards are Australia's premier technology innovation awards program, and are presented by The Australian Information Industry Association (Over 2,000 citations and an h-index of 22).

AI Profiles

Tat-Jun Chin

Tat-Jun received a BEng (Mechatronics) from the University Technology Malaysia (UTM) in 2003. From 2003 to 2004 he worked for Agilent Technologies' Sensor Solutions Division as a Test Engineer. In 2004, he received the Endeavour Australia-Asia Award (AUD$100,000) to pursue his PhD in computer vision at Monash University. He was a Research Fellow at the Institute for Infocomm Research (I2R) in Singapore from 2007–2008. In 2008 he was a Postdoctoral Research Fellow, then appointed as a Lecturer at the School of Computer Science,The University of Adelaide, South Australia. His research interests include robust estimation and statistical learning methods in Computer Vision. (Over 600 citations and an h-index of 16)

Anthony Dick

Anthony is a Senior Lecturer at The University of Adelaide's School of Computer Science. He holds a Bachelors of Mathematics and Computer Science (Hons) from the University of Adelaide and received his PhD from the University of Cambridge in 2001. Anthony is also the Deputy Director for the Australian Center for Visual Technologies. Anthony's interest areas include computer vision: that is, the problem of teaching computers how to see. He is interested in tracking lots of people or objects at once, in more than one camera. Anthony is currently working on a project to try and track Australian Rules football players during a game, in order to provide a complete picture of where each player is at each moment in time. (Over 1,500 citations and an h-index of 20).

Anders Eriksson

Anders is a QUT Vice-Chancellor's Research Fellow and was formerly a senior research associate in computer vision at the School of Computer Science, The University of Adelaide. He received his Masters of Science degree in Electrical Engineering in 2000 and his PhD in Mathematics in 2008 from Lund University, Sweden. His research areas include optimisation theory and numerical methods applied to the fields of computer vision and machine learning. In 2010 his work on robust low-rank matrix approximation won the best paper award at the 23rd IEEE Conference on Computer Vision and Pattern Recognition, San Francisco, USA. (Over 450 citations and an h-index of 9).

Clinton Fookes

Clinton is an Associate Professor with the Speech, Audio, Image and Video Technologies group within the Science and Engineering Faculty at QUT. He holds a BEng (Aerospace/Avionics), an MBA with a focus on technology innovation and management, and a PhD in the field of computer vision. Clinton actively researches in the fields of computer vision and pattern recognition including video surveillance, biometrics, human-computer interaction, airport security and operations, command and control, and complex systems. Clinton has attracted over $12.5M of cash funding for fundamental and applied research from external competitive sources and has published over 120 internationally peer-reviewed articles. He has been the Director of Research for the School of Electrical Engineering & Computer Science. He is currently the Head of Discipline for Vision and Signal Processing. He is the Technical Director for the "Airports of the Future" collaborative research project. He is a Senior Member of the IEEE, and a member of other professional organisations including the APRS. He is also an Australian Institute of Policy and Science Young Tall Poppy and an Australian Museum Eureka Prize winner. (Over 1,000 citations and an h-index of 19).

Jason Ford

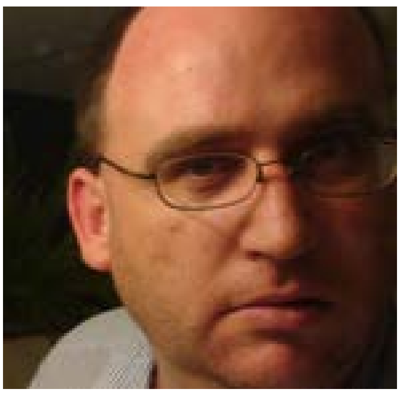

Jason received his Bachelor of Science and Bachelor of Engineering degrees from the Australian National University, Canberra, in 1995. He received his PhD from the Australian National University in 1998. Jason's first research position was as a research engineer for the Cooperative Research Centre for Robust and Adaptive Systems, Canberra, during 1996. He was appointed a research scientist at the Australian Defence Science and Technology Organisation in 1998, and then promoted to senior research scientist in 2000. In 2004, he was appointed as a research fellow at the University of New South Wales at the Australian Defence Force Academy. In 2005, he was appointed a research fellow at Queensland University of Technology, Brisbane. He was appointed to his current academic position at QUT in 2007, where he is a senior lecturer in the School of Electrical Engineering and Computer Science. His current research interests include advanced signal processing and control algorithms for aerospace applications. He has had academic visits to the Information Engineering Department at the Chinese University of Hong Kong in 2000, and to the University of New South Wales at the Australian Defence Force Academy from 2002 to 2004. (Over 600 citations and an h-index of 14).

Jonghyuk Kim

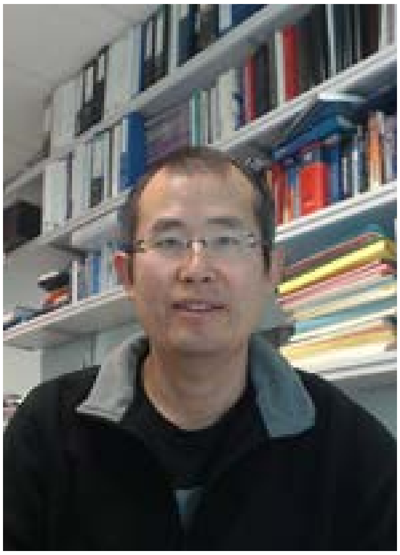

Jon obtained his PhD degree in Field Robotics at the University of Sydney in 2004, pioneering the area of airborne simultaneous localisation and mapping (SLAM). He studied mechanical engineering at KAIST, South Korea and received his BS and MS degrees in electronics/control engineering at Chungnam National University, South Korea in 1997 and 1999 respectively. He was a postdoctoral fellow at the Centre for Autonomous Systems (CAS) in Sydney before he joined ANU. He is the recipient of the Charles Sharpe Beecher Prize and Award from IMechE, UK, 2005 for his contributions to aerial robotics. He co-chaired ACRA (Australasian Conference in Robotics and Automation) in 2008 and served in Associate Editor roles for IEEE-IROS (2008) and IEEE ICRA (2010). His key contribution, published in IEEE Transactions on Aerospace and Electronics Systems together with its companion conference paper, has received over 300 citations showing significant impacts in aerial robotics. (Over 1,345 citations and an h-index of 17).

Wai Ho Li

Wai Ho is a Chief Investigator for signal processing at the Monash Vision Group (MVG), which is an $8M funded ARC research centre working on a cortical visual prosthesis ("bionic eye"). He conducts research and development on fast image processing, wearable computer vision, cortical simulations and simulated prosthetic vision user studies. He is also a core software developer for the Multiple View Geometry (MVG) open source package, writing C++ code for embedded platforms and MATLAB for desktop-based simulations. His work at Monash Vision Group has resulted in several patents and peer reviewed publications in top IEEE and ACM outlets such as ISMAR, ICRA and ISWC. The focus of his research with the ACRV is on the development of fast algorithms and real world systems that enable human-robot interaction. These algorithms allow a robot to sense people, recognise and manipulate objects as well as map and navigate the environments. The same algorithms are useful to humans in an assistive context. Wai Ho has worked with industry partners on machine learning, automation and image processing problems, such as automated asset management and pedestrian detection. (Over 200 citations and an h-index of 9).

Luis Mejias Alvarez

Luis studied electronic engineering at UNEXPO (Venezuela) obtaining the degree of Electronic Engineer in November 1999. In 2000, he joined the Masters program at ETSIT-Universidad Politécnica de Madrid obtaining a Masters of Science in Network and telecommunication systems in 2001. After completing his Masters, he moved to DISAM (ETSII) Universidad Politécnica de Madrid where he obtained his doctorate. While completing his PhD he gained extensive experience with Unmanned Aerial Vehicles and in particular Autonomous Helicopter platforms investigating in computer vision techniques for control and navigation of UAVs. He is currently a senior lecturer and researcher at QUT's Australian Research Centre for Aerospace Automation (ARCAA). (Over 900 citations and an h-index of 15).

Tristan Perez

Tristan is ACRV's Portfolio Lead in the area of Agriculture Applications. Tristan was appointed Professor of Robotics at the School of Electrical Engineering and Computer Science at QUT where he is working on agricultural robotics, bio-inspired guidance and motion control of unmanned aircraft and energy-based modeling and motion control of field robots. In 2004, he was a Research Fellow at the Mechatronics Centre of the University of Wales, Newport, where he worked on fault diagnosis and control of propulsion systems for underwater vehicles. He was also a Senior Research Fellow at the CoE for Ships and Ocean Structures (CeSOS) where he conducted research on mathematical modeling of ship dynamics for maneuvering in a seaway, ship roll stabilisation, and control allocation for ship dynamic positioning. In 2007, he joined the ARC-COE for Complex Dynamic Systems and Control (CDSC) at the University of Newcastle (UoN) where he developed a research program on dynamics and motion control of marine and aerospace vehicles. From 2009 to 2012, Tristan was appointed Adjoint Associate Professor of Ship Dynamics at NTNU. (Over 1,000 citations and an h-index of 17).

Fatih Porikli

Fatih is an IEEE Fellow and a Professor in the Research School of Engineering, Australian National University (ANU), Canberra. He is also acting as the Computer Vision Group Leader at NICTA, Australia. He received his PhD from New York University (NYU), New York in 2002. Previously he served as a Distinguished Research Scientist at Mitsubishi Electric Research Laboratories (MERL), Cambridge, USA. Before joining MERL in 2000, he developed satellite imaging solutions at Hughes Research Laboratory, Malibu CA, and 3D display systems at AT&T Research Laboratories, Middletown, NJ. Fatih was the recipient of the R&D 100 Scientist of the Year Award in 2006. He has won four best paper awards at premier IEEE conferences including the Best Paper Runner-Up at IEEE CVPR in 2007, the Best Paper at IEEE Workshop on Object Tracking and Classification beyond Visible Spectrum (OTCBVS) in 2010, and the Best Paper from IEEE AVSS in 2011, and the Best Poster Award at IEEE AVSS in 2014. Fatih has authored more than 130 publications and invented 61 patents. He is the co-editor of two books: Video Analytics for Business Intelligence and Handbook on Background Modelling and Foreground Detection for Video Surveillance. Fatih is the sole author of a pioneering paper on fast histogram computation, which is one of the most-cited papers in this area with 600+ citations in the last 6 years. (Over 6,000 citations and an h-index of 21).

Jonathan Roberts

Jonathan is ACRV's Portfolio Lead in the area of Medical and Health Applications. Jonathan is Professor in Robotics at Queensland University of Technology (QUT). His main research interest is in the area of Field Robotics and in particular making machines operate autonomously in unstructured environments. He was the Deputy Director of ARCAA until May 2013 when the Joint Venture came to an end and he was the Research Program Leader of Autonomous Systems at CSIRO's Division of Computational Informatics. His research interests encompass the area of field robotics and include autonomous aerial vehicles, autonomous ground vehicles, underwater robots and mining robots. He has over 100 publications in refereed journals and conferences and serves on the editorial board of the International Journal of Field Robotics. At the end of 2013, he was the General Chair of the International Conference on Field and Service Robotics and he is currently co-leading a special issue on UAVs in the Journal of Field Robotics. Jonathan is a co-founder of the UAV Challenge and is a past president of the Australian Robotics & Automation Association. Jonathan also has an interest in the use of robotic technology for recording cultural heritage places and experiencing museums and galleries from a distance. (Over 1,900 citations and an h-index of 36).

Ahmet Sekerciougl u

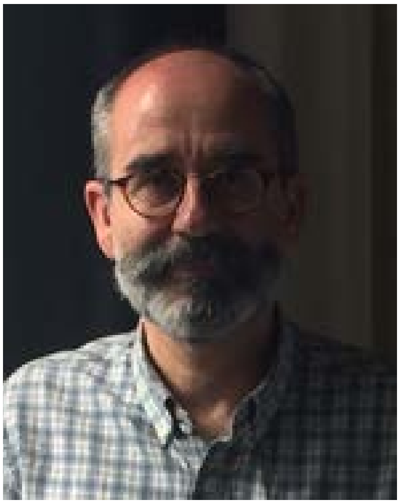

Ahmet is a senior lecturer at the Department of Electrical and Computer Systems Engineering of Monash University, Melbourne, Australia. He was the leader of the Applications Program of the Australian Telecommunications CRC until the completion of the centre's research activities (December 2007). He completed his PhD degree at Swinburne University of Technology, BSc and MSc degrees (all in Electrical and Electronics Engineering) at Middle East Technical University, Ankara, Turkey. He lectured at Swinburne University of Technology, Melbourne, Australia for 8 years. Prior to his academic career, he held numerous positions as a research engineer in private industry. Ahmet's recent research is in distributed algorithms for self-organisation in wireless networks. He is also working in the application of intelligent control techniques for multi-service networks as complex, distributed systems. (Over 1,500 citations and an h-index = 20).

Qinfeng Shi

Qinfeng is a a senior lecturer at the School of Computer Science, The University of Adelaide. He is also a member of The Australian Centre for Visual Technologies. He was an ARC Discovery Early Career Researcher Award (DECRA) Fellow between 2012-2014 (awarded in 2011). Qinfeng received a PhD in computer science in 2011 at The Australian National University (ANU) in Machine Learning. At ANU, he was under the supervision of Alex J. Smola (primary advisor, now in CMU and Google), S. V. N. Vishwanathan (now in Purdue), Li Cheng (now in A-Star), Tiberio Caetano (principal advisor, ANU and NICTA), Richard Hartley (Panel, ANU and NICTA). Before that, he completed his Bachelor and Masters study in computer science and technology in 2003 and 2006 at The Northwestern Polytechnical University (NPU). His research interests include structured learning, probabilistic graphical models, Pac-Bayes bounds analysis, and probability inequalities. (Over 800 citations and an h-index of 14).

David Suter

David holds a BSc (Applied Maths and Physics) and Diploma of Education, from The Flinders University of SA; Grad. Diploma Computing from RMIT; and a PhD in Computer Science from La Trobe University. His appointments include Lecturer (1988-1991) at La Trobe University Dept. of Computer Science and Computer Engineering; Senior Lecturer (1992-2000), Associate Professor (2001-2005), and Professor (2006-2008) at Monash University Dept. Electrical and Computer Systems Engineering; Professor (2008-) at The University of Adelaide School of Computer Science. He also was a member of the Australian Research Council College of Experts (2008-2010). His administrative positions include: Member of Monash University Council (2003-2004); Associate Dean Research and Development Faculty of Engineering Monash University (2005); Vice-President of Academic Board Monash University (2006-2008); Head of School of Computer Science, The University of Adelaide (2009 -). David is a Senior Member of the IEEE, member of Australian Computer Society, and a member of the Australian Pattern Recognition Society (having held the Vice-Presidency of that Society in previous years). He was a General co-Chair of Asian Conference on Computer Vision (ACCV2002, Melbourne Australia) - a major international computer vision conference. He was General co-Chair of the International Conference on Image Processing (ICIP2013, Melbourne Australia) - one of the premiere and largest image processing conferences, sponsored by the IEEE Signal Processing Society. (Over 3,500 citations and an h-index of 32).

Jochen Trumpf

Jochen earned a Vordiplom (Bachelor degree) in Mathematics in 1993, a Vordiplom (Bachelor degree) in Computer Science in 1994, a Diploma (Masters degree with a one year thesis project) in Mathematics in 1997 and a PhD in Mathematics in 2002, from the Department of Mathematics and Computer Science at the University of Wurzburg, Germany. From 1997-2001 he was a Research Fellow in the Institute for Applied Mathematics and Statistics at the University of Wurzburg (Germany). From 2001-2002 Jochen was a Research Fellow in the Mathematical Institute at the University of Wurzburg (Germany). During Spring 2003 he was a Visiting Fellow in the Center for Applied Mathematics at the University of Notre Dame (USA). Since April 2003, Jochen has been a Research Fellow at the Research School of Engineering, The Australian National University, Canberra (Australia). Jochen's research interests lie in applying advanced mathematical theory to engineering problems. He concentrates mostly on observer theory and design and on optimisation on manifolds or other spaces with nontrivial structure. Application areas he is interested in include computer vision, robotics and telecommunication. (Over 800 citations and an h-index of 17).

PI Profiles

François Chaumette

François will contribute to the Vision and Action theme of the ACRV and is currently the Senior Research Scientist ("Directeur de Recherche"), the head of the Lagadic group, which is a common group to Inria Rennes Bretagne Atlantique and IRISA. He graduated from Ecole Nationale Supérierure de Mécanique, Nantes, in 1987. He received a Ph.D. degree and "Habilitation à Diriger des Recherches" in computer science from the University of Rennes in 1990 and 1998 respectively. Since 1990, he has been with IRISA/INRIA, Rennes. His research interests include robotics, computer vision, and especially the coupling of these two research domains (vision-based control, active vision and purposive vision). François received the AFCET/CNRS Prize for the best French thesis in automatic control, in 1991. He was an Associate Editor of the IEEE Transactions on Robotics and Automation and of the IEEE Transactions on Robotics from 2001 until 2005. François was an Associate Editor of the International Journal of Optomechatronics from 2007 till 2012. Since 2008 he has been on the Editorial Board of the International Journal of Robotics Research. (Over 12,000 citations and an h-index of 45).

Andy Davison

Andy will contribute to the Robust Vision and Semantic Vision themes of the ACRV and holds the position of Professor of Robot Vision at the Department of Computing and leads the Robot Vision Research Group and the Dyson Robotics Laboratory at Imperial College, London. He is working in computer vision and robotics: specifically his main area of research has concerned SLAM (Simultaneous Localisation and Mapping) using vision, with a particular emphasis on methods that work in real-time with commodity cameras. This is technology that can provide low-cost and robust real-time localisation and scene understanding for domestic robots, humanoid robots, wearable sensors, game interfaces or other devices. Alongside his academic research, he has a longstanding relationship with Dyson Ltd. in the UK, having worked for them as a consultant on robot vision technology since 2005. This collaboration led to the creation in 2014 of the Dyson Robotics Laboratory at Imperial College of which he is the Director and founder. (Over 9,500 citations and an h-index of 36)

Frank Dellaert

Frank will contribute to the Algorithms and Architecture theme of the ACRV and is an Associate Professor in the School of Interactive Computing at the Georgia Institute of Technology. He is also affiliated with the RIM@GT center and is well known for contributions to Robotics and Computer Vision. He attended the Catholic University of Leuven, in Belgium, from 1984 to 1989 and received a degree in Electrical Engineering. He attended the Case Western Reserve University from 1993 to 1995 and received a master's degree in Computer Science and Engineering. In 1995 he began studying at Carnegie Mellon University where he worked as a Research Assistant and received his Ph.D. degree in Computer Science in 2001. In August of that same year, he joined the faculty of Georgia Institute of Technology. Frank holds interests in the areas of robotics and computer vision, including Bayesian inference and Monte Carlo approximations and how to attain efficiency with approximation methods. (Over 12,500 citations and an h-index of 47)

Paul Newman

Paul will contribute to the Semantic Vision theme of the ACRV and is BP Professor of Information Engineering at the University of Oxford and head of the Mobile Robotics Group. He obtained a Masters of Engineering Science from Oxford University, Balliol College in 1995. He then undertook a PhD in autonomous navigation at the Australian Center for Field Robotics, University of Sydney, Australia. In 1999 he returned to the United Kingdom to work in the commercial sub-sea navigation industry. The navigation software he wrote then was used to repair the Deep Sea Horizon leak in 2010. In late 2000 he joined the Department of Ocean Engineering at M.I.T. where, as a post-doc and later a research scientist, he worked on algorithms and software for robust autonomous navigation for both land and sub-sea agents. He was appointed to a University Lectureship in Information Engineering and became a Fellow of New College in 2005, Professor of Engineering Science in 2010 and BP Professor of Information Engineering and Fellow of Keble College in 2012. (Over 9,000 citations and an h-index of 41)

Marc Pollefeys

Marc will contribute to the Robust Vision and Vision and Action themes of the ACRV and is a full professor and head of the Institute for Visual Computing of the Department of Computer Science of ETH Zurich, which he joined in 2007. He leads the Computer Vision and Geometry lab. Before that he was a postdoctoral researcher at the Katholieke Universiteit Leuven in Belgium, where he also received his M.S. and Ph.D. degrees in 1994 and 1999, respectively. His main area of research is computer vision. One of his main research goals is to develop flexible approaches to capture visual representations of real world objects, scenes and events. He is a regular reviewer for most of the major vision, graphics and photogrammetry journals. Marc was on the Editorial Board of the IEEE Transactions on Pattern Analysis and Machine Intelligence, the International Journal of Computer Vision, Foundations and Trends in Computer Graphics and Computer Vision and several other journals. He is an IEEE Fellow. (Over 13,500 citations and an h-index

Philip Torr

Philip will contribute to the Semantic Vision theme of the ACRV and is a Professor in Computer Vision and Machine Learning at Oxford University. He completed his PhD at the Robotics Research Group of the University of Oxford. He left Oxford to work for six years as a research scientist for Microsoft Research, first in Redmond USA in the Vision Technology Group, then in Cambridge UK founding the vision side of the Machine learning and perception group. He has won several awards including the Marr prize (the highest honour in vision) in 1998, and is a Royal Society Wolfson Research Merit Award Holder. More recently he has been awarded best science paper at BMVC 2010 and ECCV 2010. He was involved in the algorithm design for Boujo released by 2D3. Boujou has won many industry awards, including Computer Graphics World Innovation Award, IABM Peter Wayne Award, and CATS Award for Innovation. The segmentation work of his group was used in the Sony Wonderbook. He continues to work closely with 2D3 as well as other companies such as Sony and Sharp. (Over 14,000 citations and an h-index of 61)

Research Fellow Profiles

Niko Sünderhauf

Niko received his PhD from Chemnitz University of Technology, Germany in 2012. In his thesis, Niko focused on robust factor graph based models for SLAM and general probabilistic estimation problems and developed the concept of switchable constraints. After his PhD, his research interests have been in the area of robust place recognition in changing environments. After two years postdocing in Chemnitz, Niko joined QUT in Brisbane in March 2014. In the ACRV Niko is continuing robust visual representations for place recognition under extreme appearance changes and will address this problem by utilising and extending various deep learning techniques.

Sareh Shirazi

Sareh completed her PhD programme in Computer Vision within the school of Information Technology and Electrical Engineering (ITEE) at the University of Queensland (UQ) and NICTA (Australia's Information Communications Technology (ICT) Research Centre of Excellence) in March 2014. Her thesis topic was mainly focused on using the non-Euclidean geometry of Riemannian manifolds to model the video sequences and it addressed three predominant tasks pertaining to computer vision applications, e.g., visual recognition, clustering and visual tracking. As a part of ACRV, Sareh will work on high to low-level vision using high-levels of abstraction such as place recognition to optimise the low-level features as well as working on robust methods for place recognition under varying lighting, viewpoints and etc.

Donald Dansereau

Donald completed his doctoral thesis on light field imaging in challenging environments, with the Australian Centre for Field Robotics at the University of Sydney. His fieldwork includes modelling a submerged Bronze Age city in Greece, hydrothermal vent exploration near Sicily, and mapping marine habitats off the coast of Tasmania. Donald completed B.Sc. and M.Sc. degrees in electrical and computer engineering at the University of Calgary in 2001 and 2004, receiving the Governor General's Gold Medal for his Master's work. Following some time in industry, he returned to research with an eye to simplicity and robustness as enabled by plenoptic imaging. At the ACRV, Donald will focus on robust vision through the application and extension of computational imaging techniques in projects RV1 and AA1. By treating visual sensing as a communication problem, he will explore the information-theoretic limits of visual perception, and tackle the challenges of distributed visual processing under limited bandwidth and computational resources.

Trung Than Pham

Trung is a research fellow based at the UoA. Trung secured his PhD in Computer Vision from the Australian Centre for Visual Technologies, the UoA, in June 2014. His thesis topic mainly focused on innovating robust statistical algorithms that effectively and efficiently discover multiple objects/structures from vision data (e.g., images). Trung was awarded a Google PhD Fellowship in 2012, and was one of only 200 young researchers selected worldwide to attend the Heidelberg Laureate Forum in 2013. As part of the ACRV, Trung is working on semantic scene segmentation. Specifically, he will develop the algorithms and representations for the environments dynamically captured by robots in terms of semantic entities such as doors, walls, cars, streets, trees, etc., as well as their properties and relationships.

Guosheng Lin

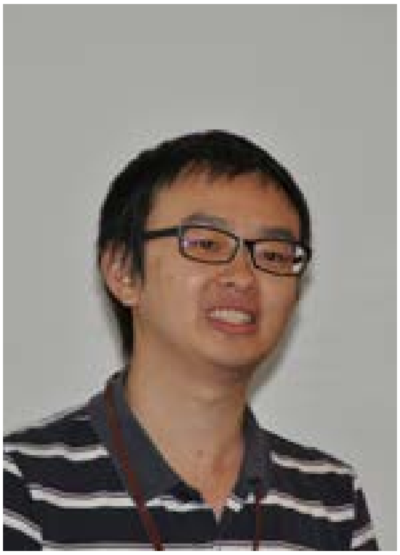

Guosheng is a research fellow with the ACRV based at the UoA on project SV2. Guosheng completed his PhD, on machine learning for computer vision, at the UoA, primarily supervised by Chunhua Shen, David Suter and T-J Chin. He has an outstanding publication record with papers in CVPR 2014, ECCV 2014, ICCV 2013, T-PAMI, ICML 2013 (oral), ICIP 2013 and ACCV 2012. He has been working on several topics on computer vision and machine learning, including image retrieval, image segmentation, object detection, structured prediction and boosting methods. Guosheng's future research with the ACRV will be on large scale recognition problems at the coarse and fine-grained level. Guosheng is a joint appointment between ACRV and the Data to Decisions CRC, which explores the use of big data in defence and national security.

Jürgen "Juxi" Leitner

Juxi is a research fellow at the QUT node. He is working in the Vision and Action (VA) and Algorithms and Architecture (AA) themes. Previously Juxi worked on making the iCub humanoid robot see and interact with the world at the Robotics Lab of the Dalle Molle Institute for AI (IDSIA). Prior to that he worked with the European Space Agency's Advanced Concepts Team. His background includes a Joint European Master in SpaceScience and Technology (SpaceMaster) and a Bachelors degree in Computer Science from the Vienna University of Technology.

Professional Staff

Sue Keay

Sue is the Centre's Chief Operating Officer, responsible for the effective leadership, management and development of the ACRV's research, education and knowledge exchange programs. A university medallist and Jaeger scholar, Sue has more than 20 years experience in the research sector, managing and ensuring impact from multidisciplinary R&D programs and teams. She has a PhD in Earth Sciences from the Australian National University and was an ARC post-doctoral fellow at The University of Queensland, before escaping the lab and moving first to science communication and then into research management, research commercialisation and IP management. She has managed national investments in water recycling research for a Centre of Excellence funded by the Department of Environment, and was the Director of UQ Business School Commercial. Sue has worked with physical scientists, engineers, social science researchers and economists and loves the challenge of managing research across distributed sites. Sue is a graduate of the Australian Institute of Company Directors and Chairs the IP and Commercialisation Committee for the Board of the CRC for Optimising Resource Extraction. She is currently completing her MBA with UQ Business School and mentors three female-led start-up companies.

Laura Vernon

Laura is ACRV's Centre Coordinator, responsible for day-to-day operations of the Centre spanning event organisation, human resources and coordination of Centre activities across nodes. She joined the team as Centre Coordinator in August 2014. Laura has worked in the University sector for almost seven years with her previous role being Executive Assistant to the Dean in the Faculty of Arts and Design at the University of Canberra. Previous to this, she was Executive Assistant to the COO at the Illawarra Health and Medical Research Institute in the University of Wollongong. She is also currently studying, part-time, a Bachelor of Commerce majoring in Human Resource Management and is due to graduate at the end of 2015.

Governance

Our Organisational Structure

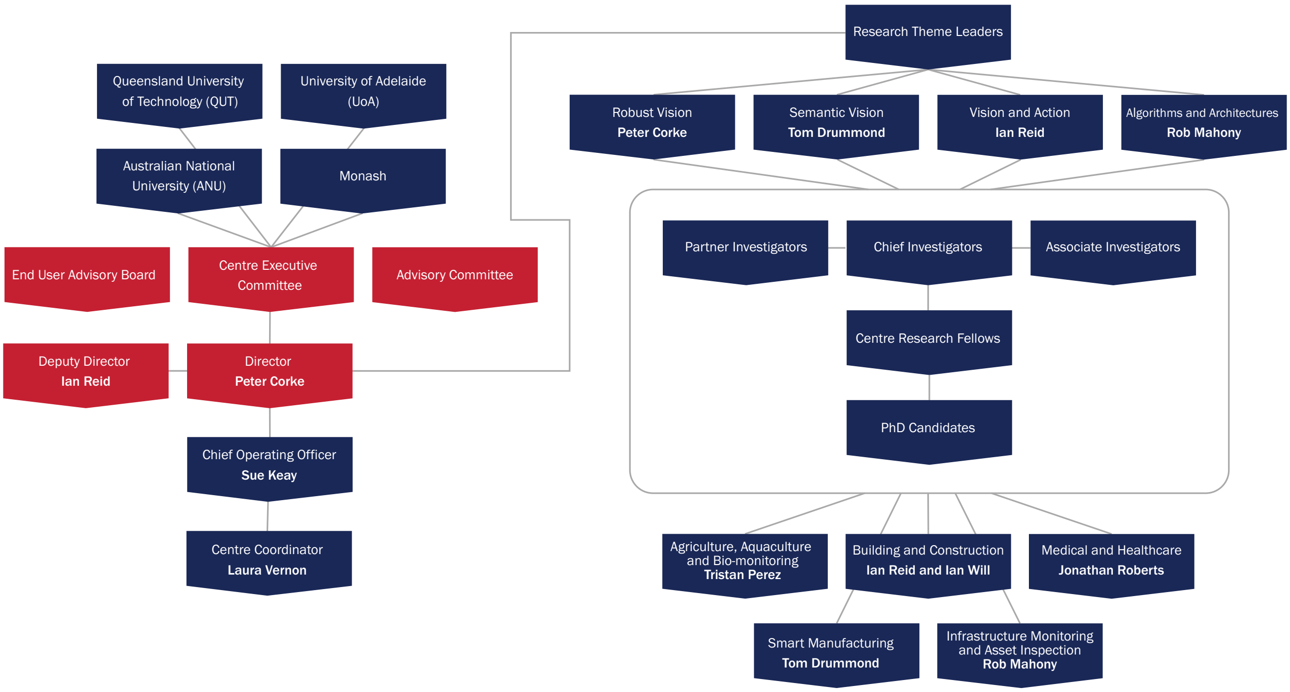

Centre Advisory Committee (CAC)

The Centre Advisory Committee comprises six people with expertise in the relevant science and with a track record in technology commercialisation. The CAC will meet annually with the Executive Committee, ideally at the Centre's annual symposium to afford the chance for the Advisory Committee members to meet with researchers and PhD candidates. The CAC will provide frank advice on the research directions of the Centre, evaluate progress toward goals and KPI achievement, and provide introductions and opportunities from their own extensive networks.

Centre Executive Committee (CEC)

The Centre's Executive Committee comprises the Centre Director, the three other Research Theme Leaders, and the Chief Operating Officer. The CEC governs the Centre. The Committee meets fortnightly via Google Hangouts: 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 and notes are recorded using the Centre's Confluence-based intranet.

Centre Roles and Responsibilities

Centre Director

Centre Director	Peter Corke

Peter is responsible for ensuring that the Centre achieves its goals in the areas of research, education, end-user engagement, capacity building, research training and development, cohesiveness and impact. He works with the Centre Executive Committee through regular fortnightly meetings to achieve these goals. Peter oversees the creation of an annual operational plan and ensures that Centre milestones and performance targets are achieved. Peter has also convened a meeting of the Centre's Advisory Committee (CAC) in 2015 and will ensure that the CAC's advice is incorporated into the Centre's strategic planning.

Research Theme (Node) Leaders

AA Theme / Monash node	Tom Drummond
SV Theme / UoA node	Ian Reid
RV Theme / QUT Node	Peter Corke
VA Theme / ANU node	Rob Mahony

Deputy Leaders

AA Theme	vacant
Monash node	Wai Ho Li
SV Theme	Stephen Gould
UoA node	Chunhua Shen
RV Theme	Hongdong Li
QUT Node	Michael Milford
VA Theme	Ahmet Sekercioglu
ANU node	Hongdong Li

Research theme leaders represent geographic nodes on the Centre Executive Committee. They are responsible for ensuring the excellence of theme research and for achieving agreed goals. Theme leaders are also responsible for mentoring, capacity building, ensuring cross-centre connectivity, communication and engagement with end-users, and identifying protectable IP. Each theme is supported by a geographically diverse group of Chief Investigators: but the choice of researchers in themes has been driven by domain expertise rather than geographical location. Each theme leader is responsible for delivery of research that will take place across more than one of the Centre's four nodes.

Each Research Theme Leader is located at a node in the Centre and plays a role in the administration of the node, principally as the contact point for institutional administrative matters. The node administration role is secondary to the theme leadership role. This is a deliberate strategy to foster strong inter-node collaboration in the Centre. Theme Leaders will be assisted by part-time administration support at each node, yet to be recruited at UoA, ANU and Monash.

As a deliberate strategy to foster leadership development amongst younger CIs, a number of portfolio roles have been created within the Centre.

Deputy Node Leaders

Are responsible for filling in for node leaders at CEC meetings as required and also for organising regular node meetings.

Deputy Theme Leaders

Are responsible for supporting and filling in for Research Theme Leaders as required. They are also responsible for organising regular theme meetings, and for recording meetings on Confluence.

HDR Portfolio Leads

Have oversight of PhD recruitment strategies, arrange robotic vision summer schools as appropriate, and are responsible for recording HDR lead meetings on Confluence. The HDR portfolio is to be split across nodes with one "leader" each year filled by rotation.

ANU	Stephen Gould – HDR Lead
Monash	Tom Drummond
QUT	Ben Upcroft
UoA	Gustavo Carneiro

Research Training

Responsible for the development and oversight of training activities for the Centre's early career researchers, as well as oversight of mentoring program, and responsibility for recording meetings and events on Confluence.

Research Training	Sue Keay

Outreach & Education

Responsible for oversight of the Centre's Communication, Outreach and Engagement strategy and for coordinating outreach activities across nodes.

Engagement	not appointed in 2014

Applications Lead

Responsible for introducing robotics to the application area and encouraging application of Centre research.

Agriculture, Aquaculture and Bio-monitoring	Tristan Perez (QUT)
Infrastructure and Asset Lead	Rob Mahony (ANU)
Smart Manufacturing	Tom Drummond (Monash)
Building and Construction	Ian Reid (UoA)
Medical and Healthcare	Jonathan Roberts (QUT)

Gender Diversity

Responsible for the development and oversight of Gender Diversity initiatives, to advocate for women in STEM.

Male Champion for Change	Ben Upcroft

Governance KPIs

Governance			
Performance Measure	Reporting Frequency	Target 2014	Outcome 2014
Breadth, balance and experience of the members of the Advisory committee	At review		
Frequency, attendance and value added by Advisory Committee meetings	At review	1 meeting of Advisory Committee	0
Vision and usefulness of the Centre strategic plan	At review	Annual Review by Committee	0
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
• Number of new joint research projects		0	
Capacity building of the Centre through scale and outcomes	At review	Early career staff will be developed through supervisor mentoring, a Centre-run mentoring program, and opportunities for supervision (of PhD and undergraduate students), teaching and grant writing. PhD students will be developed through supervisor mentoring, a Centre-run mentoring program, and opportunities for supervision of undergraduate students.	

Members of the Centre Advisory Committee

Dr Alex Zelinsky (Chair)

Alex is the Chief Defence Scientist and head of the Defence Science and Technology Organisation (DSTO). Before joining DSTO 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 currently a Scientia professor and Laureate Fellow at the University of New South Wales. She obtained a double degree in physics and chemistry and was awarded a PhD in Physics from Durham University, UK in 1992. Her postdoctoral position was as a Research Fellow in quantum electronics at the Cavendish Laboratory in Cambridge, UK, where she gained an international reputation for her work in the discovery of the '0.7 feature' and metallic behaviour in 2D Gas hole systems. In 1999, she was awarded a QEII Fellowship and came to Australia where she was a founding member of the Centre of Excellence for Quantum Computer Technology. She established a research group dedicated to the fabrication of atomic-scale devices in silicon and germanium using the atomic precision of a scanning tunneling microscope. She has published more than 360 papers in refereed journals with an h-index of 40 including 27 Physical Review Letters and papers in Nature, Science, Nature Materials, Nature Physics and Nature Nanotechnology. She has been awarded two Federation Fellowships and is currently an ARC Laureate Fellow and Director of an ARC Centre of Excellence - the Centre for Quantum Computation and Communication Technology (CQC2T).

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), and also a Fellow of Oriel College Oxford. He served as the Chief Executive Officer of National ICT Australia Limited (NICTA) from December 2010 to November 2014.

Professor Sir Michael Brady

Mike was recently appointed as 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.

Key Guidelines and Policies

The Centre has developed a number of key Policies and Guidelines since starting in July 2014.

Guidelines	
Acknowledgments and Affiliations	Mentoring
ACRV Position Titles	Recruitment of PhD students
Authorship of Publications (Determining)	Sponsorship
CI Travel	Support for AI Travel
Conference Attendance – Support from ACRV	Travel - between nodes and to partner organisations
Examples of How to Apply the ACRV Publication Policy as an Author	

Policies
Distribution of Joint Centre Costs
Publication Policy
Travel Policy

Finance

The following table provides a summary of the ACRV's financial performance for the 2014 calendar year. The ACRV receives 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 found via industry engagement and relevant industry projects.

Expenditure by the Centre is dominantly on personnel, with smaller amounts allocated towards travel, equipment and operating expenses (see Figure 3). A significant carry-forward has been accumulated for 2015 as ARC funds were released to the ACRV 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 has been low. As the Centre is actively recruiting personnel at all nodes, this under-spend is not expected to continue in 2015, although a carry-forward may continue over the life of the Centre reflecting its "late" start.

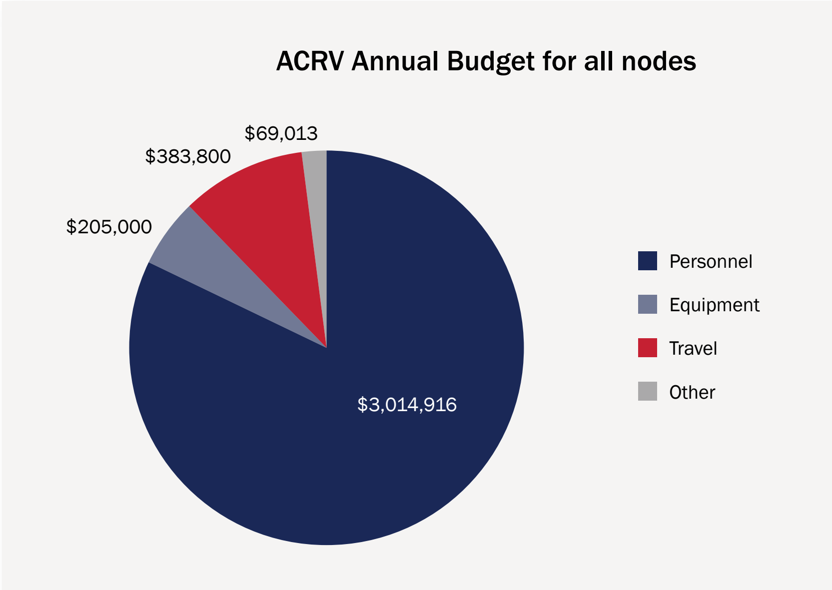

Figure 3: Allocation of Centre Resources

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 ACRV 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	UoA	Monash	Oxford	ICL	GT	INRIA	ETHZ	NICTA
Cash	$2.45m	$1.75m	$1.82m	$840k	$0	$0	$0	$0	$0	$0
In-kind	$2.13m	$2.47m	$1.23m	$1.14m	$224k	$127k	$126k	$101k	$185k	$210k

Finance KPIs

Organisational support			
Performance Measure	**Reporting Frequency**	**Target 2014**	**Outcome 2014**
Annual cash contributions from Administering and Collaborating Organisations	Annually		
QUT		$350,000	$350,000
Monash		$116,400	$116,400
ANU		$230,000	$230,000
Adelaide		$244,400	$244,400
Annual in-kind contributions from Administering and Collaborating Organisations	Annually		
QUT		$309K	$434K
Monash		$151K	$151K
ANU		$332K	$114K
Adelaide		$153K	$200K
Annual cash contributions from Partner Organisations *list each Organisation separately	Annually	Nil	Nil
Annual in-kind contributions from Partner Organisations	Annually		
Georgia Tech		$18K	$3.6K
INRIA		$14.5K	$0K
Imperial College		$18.1K	$3.6K
NICTA		$30K	$5.8K
Swiss Federal Institute		$26.4K	$0K
Oxford		$32K	$3.2K
Other research income sourced by Centre *End User (industry, public sector, ARC Linkage and Discovery in non-core areas, CRC	Annually	$0	$1M
Number of new organisations collaborating with, or involved in, the Centre	Annually	2	2
Level and quality of infrastructure provided to the Centre	At review		

Financial Statement

ARC Centre of Excellence for Robotic Vision

Statement of Operating Income and Expenditure for year ended 31 December 2014

Income			
ARC Income			
	ARC Centre Grant distributed as follows		2,714,290
	Monash University	332,360	
	University of Adelaide	720,120	
	Australian National University	692,420	
	Queensland University of Technology	969,390	
Collaborating Organisational Income			
	Monash University		116,929
	University of Adelaide		244,400
	Australian National University		230,000
	Queensland University of Technology		360,000
Other Income			
	ARC Indexation		82,027
Total Income			3,747,646

Expenditure		
	Purchased Equipment	42,256
	Travel - Regular meetings of Centre staff	3,404
	Travel - Conferences and workshops (Dir, COO, CI's)	11,549
	Travel - Conferences and workshops (Postdocs and students)	7,260
	Travel - Visits to nodes (Dir, COO, CI's)	1,043
	Travel - Visits to nodes (Postdocs and students)	1,020
	Maintenance (IT and lab)	8,350
	Salaries	373,061
	Other	38,002
Total Expenditure		485,946
Surplus/Deficit		3,261,700
	Balance brought forward from 2013	-
Total carry forward surplus*		3,261,700

*Surplus from 2014 half year of trading
** Based on combined data from all nodes. Each nodes's data is compliant with its own policy and procedures

Outputs

Publications

Book Chapters

Bazin, J., Seo. Y., Hartley, R., & Pollefeys, M. 2014. Globally Optimal Inlier Set Maximization With Unknown Rotation and Focal Length. Computer Vision-ECCV 2014, Springer International Publishing. pp803-817.

Carneiro, G. & Lazaro, D. Projecto PRINTART: Sistema computacional para apoio ao estudo da azulejaria portuguesa. Chapter in Book “A Herança de Santos Simões Novas Perspectivas para o Estudo da Azulejaria e da Cerâmica” (in Portuguese). ISBN: 978-989-689-400-9. 2014.

*Dick, A., van den Hengel, A., & Detmold, H. 2014. “Large Scale Camera Topology Mapping: Application to Re-Identification”. In Person Re-Identification, Springer, pp.391-411 2014.

Nascimento, J., Carneiro, J., &Freitas, A. Tracking and Segmentation of the Endocardium of the Left Ventricle in 2D Ultrasound using Deep Learning Architectures and Monte Carlo Sampling. Chapter in Book “Biomedical Image Segmentation: Advances and Trends”, edited by Dr. Ayman El-Baz, University of Louisville.

Shen, C. and van den Hengel, A. 2014. “Semidefinite Programming”, in Computer Vision: A Reference Guide, Springer US, pp. 717-719.

Journal Articles

*Ajanthan, T., Hartley, R., Salzmann, M., & Li, H. Iteratively Reweighted Graph Cut for Multi-label MRFs with Non-convex Priors. arXiv preprint arXiv:1411.6340, 24 Nov 2014.

Bazin, J. C., Seo, Y., Hartley, R., & Pollefeys, M. “Globally Optimal Inlier Set Maximization With Unknown Rotation and Focal Length”, in Computer Vision–ECCV 2014, pp. 803-817.

Chen, Y., Li, X., Dick, A., & Hill, R. Ranking Consistency for Image Matching and Object Retrieval. Pattern Recognition, Vol 47, Issue 3, 2014

*Dai, Y., Li, H., & He, M. A Simple Prior-Free Method for Non-rigid Structure-from-Motion Factorization. International Journal of Computer Vision 107(2): 101-122 (2014)

*Edussooriya, C., Dansereau, D., Bruton, L., & Agathoklis, P. 2015. “Five-Dimensional (5-D) Depth-Velocity Filtering for Enhancing Moving Objects in Light Field Videos.” IEEE Transactions on Signal Processing (TSP).

Gould, S. & He, X. “Scene Understanding by Labeling Pixels”. In Communications of the ACM, November 2014.

Gould, S. “Learning Weighted Lower Linear Envelope Potentials in Binary Markov Random Fields”. In IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 2014.

*Harandi, M., Hartley, R., Lovell, B., & Sanderson, C. Sparse Coding on Symmetric Positive Definite Manifolds using Bregman Divergences.Neural Networks and Learning Systems, IEEE Transactions on (Volume:PP, Issue: 99) 27 January, 2015.

*Harandi, M., Hartley, R., Shen, C., Lovell, B., & Sanderson, C. Extrinsic Methods for Coding and Dictionary Learning on Grassmann Manifolds. Neural Networks and Learning Systems, IEEE Transactions on (Volume:PP, Issue: 99), 27 Jan 201.

Hartley, R. “Projective Reconstruction,” in Computer Vision: A Reference Guide, 2014, pp640-651

*Isaksson, M. & Eriksson, A. “A Method for Extending Planar Axis-Symmetric Parallel Manipulators to Spatial Mechanisms”, in Mechanism and Machine Theory, vol. 83, pp. 1-13, 2014.

*Jayasumana, S., Hartley, R., Salzmann, M., Li, H., & Harandi, M. “Kernel Methods on Riemannian Manifolds with Gaussian RBF Kernels”, IEEE Transactions on Pattern Analysis & Machine Intelligence, 2015

Kusakunniran, W., Wu, Q., Zhang, J., Li, H., & Wang, L. Recognizing Gaits Across Views Through Correlated Motion Co-Clustering. IEEE Transactions on Image Processing 23(2): 696-709 (2014)

Li, X., Hu, W., Shen, C., Dick, A., & Zhang, Z. Context-aware hypergraph construction for robust spectral clustering. IEEE Transactions on Knowledge and Data Engineering. volume: 26, number: 10, pages: 2588–2597. 2014.

Li, Y., Jia, W., Shen, C., & van den Hengel, A. Characterness: An indicator of text in the wild. IEEE Transactions on Image Processing. volume: 23, number: 4, pages: 1666–1677. 2014.

Liu, F., Zhou, L., Shen, C., Yin, J. Multiple kernel learning in the primal for multi-modal Alzheimer’s disease classification. IEEE Journal of Biomedical and Health Informatics. 2014. [Online published at IEEE: 10 October 2013]

Liu, J., Sridharan, S., Fookes, C., & Wark, T. “Optimal Camera Planning under Versatile User Constraints in Multi-camera Image Processing Systems”, IEEE Transactions on Image Processing, Vol. 23, No. 1, pp.171-184, January 2014. [2013 IF: 3.111; 5Yr IF: 3.925]

Lu, Y., Wang, L., Lu, J., Yang, J., & Shen, C. Multiple kernel clustering based on centered kernel alignment. Pattern Recognition. volume: 47, number: 11, pages: 3656–3664. 2014.

*Lu, Z., Carneiro, G., & Bradley, A. P. An Improved Joint Optimization of Multiple Level Set Functions for the Segmentation of Overlapping Cervical Cells. Accepted for Publication by IEEE TIP in DEC 2014.

*Mcfadyen, A., Corke, P., & Mejias, L. “Visual predictive control of spiral motion,” Robotics, IEEE Transac on, vol. PP, no. 99, pp. 1–14, 2014.

Milford, M., & Schulz, R. “Principles of goal-directed spatial robot navigation in biomimetic models”, in Philosophical Transactions of the Royal Society B: Biological Sciences, vol. 369, no. 1655, pp20130484, 2014.

Milford, M., Vig, E., Scheirer, W., & Cox, D. “Vision□based Simultaneous Localization and Mapping in Changing Outdoor Environments”, in Journal of Field Robotics, vol. 31, no. 5, pp780-802, 2014.

Neubert, P., Sünderhauf, N., Protzel, P. (2014). Superpixel-based Appearance Change Prediction for Long-Term Navigation Across Seasons. Robotics and Autonomous Systems Journal.

Niu, Y., Dick, A., & Brooks, M. Compass Rose: A Rotational Robust Signature for Optical Flow. Computation IEEE Transactions on Circuits and Systems for Video Technology, Vol 24, Issue 1, 2014.

Paisitkriangkrai, S., & Shen, C. “Large-margin learning of compact binary image encodings”, in Image Processing, IEEE Transactions on, vol. 23, no.9, pp.4041-4054, Sept 2014.

Paisitkriangkrai, S., Shen, C., & van den Hengel, A. “A Scalable Stagewise Approach to Large-Margin Multiclass Loss-Based Boosting,” in Neural Networks and Learning Systems, IEEE Transactions on, vol. 25, pg. 5, pp. 1002-1013. 2014.

Paisitkriangkrai, S., Shen, C., & van den Hengel, A. Asymmetric pruning for learning cascade detectors. IEEE Transactions on Multimedia. volume: 16, number: 5, pages: 1254–1267. 2014.

Paisitkriangkrai, S., Shen, C., & van den Hengel, A. Large-margin learning of compact binary image encodings. IEEE Transactions on Image Processing. volume: 23, number: 9, pages: 4041–4054. 2014.

Paisitkriangkrai, S., Shen, C., Shi, Q., & van den Hengel, A. “RandomBoost: simplified multiclass boosting through randomization,” Neural Networks and Learning Systems, IEEE Transactions on, vol. 25, no.4, pp.764-779, 2014.

Pham, T. T., Chin, T. J., Schindler, K., & Suter, D. “Interacting geometric priors for robust multi-model fitting”, in IEEE Transactions on Image Processing (TIP), 23(10):4601-4610, Oct 2014.

Pham, T. T., Chin, T. J., Schindler, K., & Suter, D. Interacting geometric priors for robust multi- model fitting. IEEE Transactions on Image Processing (TIP), accepted 24 July 2014

Pham, T. T., Chin, T. J., Yu, J., & Suter, D. “The random cluster model for robust geometric fitting”, in IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 36(8):1658-1671, Aug 2014.

Pham, T. T., Chin, T. J., Yu, J., & Suter, D. “The random cluster model for robust geometric fitting”, in IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 36(8):1658-1671, Aug 2014.

Purkait, P., Chin, T. J., Suter, D., & Ackermann, H. “Clustering with hypergraphs: The case for large hyperedges,” (ECCV2014) in Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2014, vol. 8692 LNCS, pp. 672– 687.

Ren, C. Y., Prisacariu, V., & Reid, I. Regressing local to global shape properties for online segmentation and tracking. International Journal of Computer Vision. Vol 106. No 3. Pages 269-281. 2014

Rezatofighi, S., Gould, S., Vo, B., Mele, K., & Hartley, R. Multi-Target Tracking with Time-Varying Clutter Rate and Detection Profile: Application to Time-lapse Cell Microscopy Sequences. Medical Imaging, IEEE Transactions on , vol. PP , no. 99, p1, 12 Jan 2015.

Shen, C., Kim, J., Liu, F., Wang, L., & van den Hengel, A. “Efficient dual approach to distance metric learning,” in Neural Networks and Learning Systems, IEEE Transactions on, vol. 25, no. 2, pp. 394-406, 2014.

Shen, C., Kim, J., Liu, F., Wang, L., van den Hengel, A. Efficient dual approach to distance metric learning. IEEE Transactions on Neural Networks and Learning Systems. volume: 25, number: 2, pages: 394–406. 2014.

*Shen, C., Lin, G., & van den Hengel, A. StructBoost: Boosting methods for predicting structured output variables. IEEE Transactions on Pattern Analysis and Machine Intelligence. volume: 36, number: 10, pages: 2089–2103. 2014.

Shen, F., Shen, C., Hill, R., van den Hengel, A., & Tang, Z. Fast approximate l_∞ minimization: Speeding up robust regression. Computational Statistics and Data Analysis.

volume: 77, pages: 25–37. 2014.

Shen, F., Shen, C., Hill, R., van den Hengel, A., Tang, Z. “Fast approximate L∞ minimization: Speeding up robust regression”, in Computational Statistics & Data Analysis, vol.77, pp.25-37. 2014.

*Strelow, D., Wang, Q., Si, L., & Eriksson, A. General, Nested, and Constrained Wiberg Minimization. To appear, under revision. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2014.

Szpak, Z. L., Chojnacki, W., & van den Hengel, A. “Guaranteed Ellipse Fitting with a Confidence Region and an Uncertainty Measure for Centre, Axes, and Orientation”, in Journal of Mathematical Imaging and Vision, 27 Jan 2014

*Szpak, Z. L., Chojnacki, W., Eriksson, A., & van den Hengel, A. “Sampson distance based joint estimation of multiple homographies with uncalibrated cameras,” in Computer Vision and Image Understanding, vol. 125, pp200-213, 2014,

Tran, Q. H., Chin, T. J., Chojnacki, W. & Suter, D. “Sampling minimal subsets with large spans for robust estimation”, in International Journal on Computer Vision (IJCV), 106(1):93-112, Jan 2014.

Wang, L., Zhou, L., Shen, C., Liu, L., & Liu, H. A hierarchical word-merging algorithm with class separability measure. IEEE Transactions on Pattern Analysis and Machine Intelligence. volume: 36, number: 3, pages: 417--435. 2014.

*Wang, X., Sekercioglu, Y. A., & Drummond, T. 2014. “Vision-Based Cooperative Pose Estimation for Localization in Multi-Robot Systems Equipped with RGB-D Cameras. Sensors Journal (Coordination of Robotics Systems special issue).

*Warren, M., Corke, P., & Upcroft, B. “Long-range stereo visual slam for extended altitude flight of UAVs,” J. Robotics Research, (accepted).

Yan, Y., Shen, C., & Wang, H. Efficient semidefinite spectral clustering via Lagrange duality. IEEE Transactions on Image Processing. volume: 23, number: 8, pages: 3522--3534. 2014.

Yan, Y., Wang, H., & Suter, D., “Multi-subregion based correlation filter bank for robust face recognition,” PATTERN RECOGNITION, vol. 47, no. 11, pp. 3487–3501, Nov. 2014.

*Yu, J., Eriksson, A., Chin, T. J., & Suter, D. “An adversarial optimization approach to efficient outlier removal”, in Journal of Mathematical Imaging and Vision (JMIV), 48(3): 451-466, Mar 2014.

Zaragoza, J., Chin, T. J., Tran, Q. H., Brown, M. S., & Suter, D. “As-projective-as-possible image stitching with Moving DLT”, in IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 36(7):1285-1298, July 2014.

Zhong, B., Yao, H., Chen, S., Ji, R., Chin, T. J., & Wang, H. “Visual tracking via weakly supervised learning from multiple imperfect oracles”, in Pattern Recognition (PR), 47(3):1395-1410, Mar 2014.

Conference Papers

*Anantharajah, K., Ge, Z., McCool, C., Denman, S., Fookes, C., Corke, P., Tjondronegoro, D., & Sridharan, S. “Local inter-session variability modelling for object classification.”, IEEE Winter Conference on Applications of Computer Vision (WACV), 2014.

*Anh Ngo, T. & Carneiro, G. “Fully Automated Non-rigid Segmentation with Distance Regularized Level Set Evolution Initialized and Cosntrained by Deep-structured Inference.” CVPR 2014.

Berkvens, R.; Jacobson, A., Milford, M., Peremans, H., & Weyn, M. “Biologically inspired SLAM using Wi-Fi”, in Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on, pp1804-1811, 2014.

Cadena, C., Dick, A., & Reid, I.. A Fast Modular Scene Understanding System using Context-Aware Object Detection, accepted for International Conference on Robotics and Automation, 2015

*Chen, Q., & Carneiro, G. “Artistic Image Analysis using the Composition of Human Figures.” ECCV 2014 Workshop. Springer ISBN 978-3-319-16177-8

Chen, Y., Dick, A., Li, X. Context Based Re-ranking for Object Retrieval. Proc. Asian Conf. Computer Vision, 2014.

*Chen, Z., Lam, O., Jacobson, A., & Milford, M., “Convolutional Neural Network-based Place Recognition”, 2014 Australasian Conference on Robotics and Automation

*Chen, Z.; Jacobson, A., Erdem, U. M., Hasselmo, M. E., Milford, M. “Multi-scale bio-inspired place recognition”, in Robotics and Automation (ICRA), 2014 IEEE International Conference on, pp.1895-1901, 2014.

*Chi, Y., & Porikli, F,. Classification and boosting with multiple collaborative representations, IEEE Transaction on Pattern Recognition and Machine Intelligence (PAMI), August 2014

Dai, Y., & Li, H. Rank Minimization or Nuclear-Norm Minimization: Are We Solving the Right Problem? DICTA 2014: 1-8

English, A. R., Ross, P., Ball, D., & Corke, P. “Vision based guidance for robot navigation in agriculture Proc. IEEE Int. Conf. Robotics and Automation, pp. 1693–1698, 2014.

*Eriksson, A., & Isaksson, M. Pseudoconvex Proximal Splitting. Conference on Computer Vision and Pattern Recognition. CVPR, Columbus, USA, 2014.

*Eriksson, A., Isaksson, M., & Chin, T-J. High-Breakdown Bundle Adjustment, IEEE Winter Conference on Applications of Computer Vision WACV 2015.

*Faraki, M., Harandi, M., & Porikli, F. Material classification on symmetric positive definite manifolds, IEEE Winter Applications and Computer Vision Conference (WACV), 2015

Ford, J. J., Molloy, T. L., & Hall, J. L. “Compressed sensing using hidden Markov models with application to vision based aircraft tracking”, in Proc. of FUSION 2014.

Gould, S., Zhao, J, He, X., & Zhang, Y. “Superpixel Graph Label Transfer with Learned Distance Metric”. In Proceedings of the European Conference on Computer Vision (ECCV), 2014.

Guo, F, Fookes, C., Denman, S, Mejias, L, & Sridharan, S., “Automatic UAV Emergency Landing Site Detection using Support Vector Machine”, IEEE International Conference on Digital Image Computing Techniques and Applications, 2014.

*Harandi, H., Salzmann, M., & Porikli, F. Bregman divergences for infinite dimensional covariance matrices, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014

*Harandi, M. T., Salzmann, M., & Hartley, R. “From manifold to manifold: geometry-aware dimensionality reduction for SPD matrices”, in Computer Vision–ECCV 2014, pp17-32, 2014.

Harandi, M. T., Salzmann, M., Jayasumana, S., Hartley, R., & Li, H. Expanding the Family of Grassmannian Kernels: An Embedding Perspective. ECCV (7) 2014: 408-423

Harandi, M. T., Salzmann, M., Jayasumana, S., Hartley, R., Li, H. 2014. Expanding the family of grassmannian kernels: An embedding perspective. Computer Vision-ECCV 2014. Springer International Publishing. pp408-423

Hartley, R. “Keynote lecture 2: Riemannian manifolds, kernels and learning”, in 2014 11th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), Aug. 2014.

He, X., & Gould, S., “An Exemplar-based CRF for Multi-instance Object Segmentation”. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.

*Isaksson, M., & Eriksson, A. Analysis of the Inverse Kinematics Problem for 3-DOF Axis-Symmetric Parallel Manipulators with Parasitic Motion IEEE International Conference on Robotics and Automation ICRA, Hong Kong, 2014.

*Jayasumana, S., Hartley, R., Salzmann, M., Li, H., Harandi, M. “Optimizing over radial kernels on compact manifolds”, in 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June, 2014.

*Ji, P., Li, H., Salzmann, M., & Dai, Y. “Robust Motion Segmentation with Unknown Correspondences.” ECCV (6) 2014: 204-219

*Ji, P., Salzmann, M., & Li, H. “Efficient dense subspace clustering.” WACV 2014: 461-468

*Ji, P., Zhong, Y., Li, H., & Salzmann, M. “Null space clustering with applications to motion segmentation and face clustering”, IEEE International Conference on Image Processing 2014 (ICP 2014), Paris, France.

*Kneip, L., & Li, H. “Efficient Computation of Relative Pose for Multi-camera Systems.” CVPR 2014: 446-453

*Kneip, L., Li, H., & Seo, Y. “UPnP: An Optimal O(n) Solution to the Absolute Pose Problem with Universal Applicability.” ECCV (1) 2014: 127-142

*Lam, O., Dayoub, F., Schulz, R., & Corke, P. “Text recognition approaches for indoor robotics : a compari” in Australasian Conference on Robotics & Automation (ACRA2014), (University of Melbourne, Melbo VIC), December 2014.

Li, B., Dai, Y., He, M., & van den Hengel, A. “A relaxation method to articulated trajectory reconstruction from monocular image sequence,” in Signal and Information Processing (ChinaSIP), 2014 IEEE China Summit & International Conference on, pp389-393, 2014.

*Li, H., Li, Y., & Porikli, F. DeepTrack: learning discriminative feature representations by convolutional neural networks for visual tracking, (multiple CNNs), British Machine Vision Conference (BMVC), 2014

*Li, H., Li, Y., & Porikli, F. Robust online visual tracking with an single convolutional neural network, Asian Conference on Computer Vision (ACCV), 2014

*Lin, G., Shen, C., & Wu, J. “Optimizing ranking measures for compact binary code learning”. European Conference on Computer Vision (ECCV’14). 2014.

*Lin, G., Shen, C., Shi, Q., van den Hengel, A., & Suter, D. “Fast supervised hashing with decision trees for high-dimensional data”. IEEE Conference on Computer Vision and Pattern Recognition (CVPR’14). 2014.

Liu, B., He, X., & Gould, S., “Joint Semantic and Geometric Segmentation of Videos with a Stage Model”. In IEEE Winter Conference on Applications of Computer Vision (WACV), 2014.

*Liu, L., Shen, C., Wang, L., van den Hengel, A., Wang, C. “Encoding high dimensional local features by sparse coding based Fisher vectors.” Advances in Neural Information Processing Systems (NIPS’14). 2014.

Lowry, S., Milford, M. J., Wyeth, & G. F. “Transforming morning to afternoon using linear regression techniques”, in Robotics and Automation (ICRA), 2014 IEEE International Conference on, pp.3950-3955, 2014.

Lowry, S., Wyeth, G. & Milford, M., “Unsupervised Online Learning of Condition-Invariant Images for Place Recognition”, 2014 Australasian Conference on Robotics and Automation

Lowry, S., Wyeth, G. F., & Milford, M. J. "Towards training-free appearance-based localization: probabilistic models for whole-image descriptors", in Robotics and Automation (ICRA), 2014 IEEE International Conference on, pp711-717, 2014.

Lowry, S., Wyeth, G., & Milford, M. "Biologically inspired SLAM using Wi-Fi Unsupervised online learning of condition-invariant images for place recognition", in Proceedings of the Australasian Conference on Robotics and Automation 2014

Martin, S., & Corke, P., "Long-term exploration and tours for energy constrained robots with online pro ceptive traversability estimation," in Proc. IEEE Int. Conf. Robotics and Automation, pp. 5778–5785, 2014

*Milan, A., Gade, R., Dick, A., Moeslund, T. B., & Reid, I.. Improving Global Multi-target Tracking with Local Updates. Workshop on Visual Surveillance Re-Identification, held at European Conf. Computer Vision, 2014

Milford, M., & George, A. (2014) Featureless visual processing for SLAM in changing outdoor environments. In Field and Service Robotics: Results of the 8th International Conference [Springer Tracts in Advanced Robotics, Volume 92], Springer, Matsushima, Japan, pp. 569-583.

*Milford, M., Firn, J., Beattie, J, Jacobson, A., Pepperell, E., Mason, E., Kimlin, M., & Dunbabin, M. "Automated Sensory Data Alignment for Environmental and Epidermal Change Monitoring", 2014 Australasian Conference on Robotics and Automation

Milford, M., Scheirer, W., Vig, E., Glover, A., Baumann, O., Mattingley, J., & Cox, D. "Condition-invariant, top-down visual place recognition", in Robotics and Automation (ICRA), 2014 IEEE International Conference on, pp.5571, 2014.

Molloy, T. L., & Ford, J, "Short-data Recursive HMM Parameter Estimation For Rapid Vision-based Aircraft Heading Estimation"", accepted to appear in Proc. of Australian Control Conference (AUCC 2014), Nov. 2014.

Molloy, T. L., Ford, J & Mejias, L. Looming Aircraft Threats: Shape-based Passive Ranging of Aircraft from Monocular Vision, in Proc. of Australia Conference on Robotics and Automation (ACRA 2014), Dec. 2014.

Molloy, T., & Ford, J. J. "Asymptotic Minimax Robust and Misspecified Lorden Quickest Change Detection For Dependent Stochastic Processes", in Proc. of FUSION 2014.

Morris, T, Dayoub, F., Corke, P., & Upcroft, B, "Simultaneous localization and planning on multiple hypotheses," in Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference pp. 4531–4536, Sept 2014.

Morris, T, Dayoub, F., Corke, P., Wyeth, G., & Upcroft, B, "Multiple map hypotheses for planning navigating in non-stationary environments," in Proc. IEEE Int. Conf. Robotics and Automation, pp. 2765–2 2014.

Moussavi, F., Yu, W., Lorenzen, P., Oakley, J., Russakoff, D., & Gould, S., "A Unified Graphical Models Framework for Automated Human Embryo Tracking in Time Lapse Microscopy". In Proceedings of the International Symposium on Biomedical Imaging (ISBI), 2014.

*Nascimento, J. & Carneiro, G. "Non-rigid Segmentation using Sparse Low dimensional Manifolds and Deep Belief Networks". CVPR 2014.

*Nasihatkon, B., Hartley, R., & Trumpf, J. "On Projective Reconstruction In Arbitrary Dimensions", in Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, pp477-484, 2014.

*Ni, J., Marks, T. K., Tuzel, O., & Porikli, F. Detecting 3D geometric boundaries of indoor scenes under varying lighting, IEEE Winter Applications and Computer Vision Conference (WACV), 2014

*O'Sullivan, L. & Corke, P. "Empirical modelling of rolling shutter effect," in Proc. IEEE Int. Conf. Rob and Automation, (Hong Kong), pp. 5277–5283, 2014.

*Olsson, C., Ulen, J., & Eriksson, A. Local Refinement for Stereo Regularization International Conference on Pattern Recognition ICPR, Stockholm, 2014

*Paisitkriangkrai, S., Shen, C., & van den Hengel, A. "Strengthening the effectiveness of pedestrian detection with spatially pooled features." European Conference on Computer Vision (ECCV'14). 2014.

Parra Bustos, A., Chin, T. J., & Suter, D. "Fast rotation search with stereographic projections for 3D registration" In Computer Vision and Pattern Recognition (CVPR) 2014.

Pepperell, E., Corke, P., & Milford, M. "Towards Vision-Based Pose- and Condition-Invariant Place Recognition along Routes", 2014 Australasian Conference on Robotics and Automation

Pepperell, E., Corke, P., & Milford, M. "All-environment visual place recognition with smart," in Proc. I Int. Conf. Robotics and Automation, pp. 1612–1618, 2014.

Purkait, P., Chin, T. J., Ackermann, H., & Suter, D. "Clustering with hypergraphs: the case for large hyperedges" In European Conference on Computer Vision (ECCV) 2014.

Ren, C., Prisacariu, V., Murray, D., & Reid, I. 3D Tracking of Multiple Objects with Identical Appearance using RGB-D Input, 3DV, Dec 2014

Ross, P., English, A, Ball, D., & Corke, P., "A method to quantify a descriptor's illumination variance Australian Conference on Robotics and Automation, 2014.

Ross, P., English, A. R., Ball, D., Upcroft, B, Wyeth, G., & Corke, P., "Novelty-based visual obstacle dete in agriculture," in Proc. IEEE Int. Conf. Robotics and Automation, pp. 1699–1705, 2014.

*Sa, I., Hrabar, S., & Corke, P. "Inspection of pole-like structures using a vision-controlled vtol uav and sh autonomy," in Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference pp. 4819–4826, Sept 2014.

Segal, A. & Reid, I. Hybrid Inference Optimization for Robust Pose Graph Estimation, International Conference on Intelligent Robots and Systems 2014

*Sherazi, S. Object Tracking via non-Euclidean geometry: A Grassmann approach, IEEE Winter Conference on Applications of Computer Vision (WACV) 2014

*Shi, Q., Reid, M., Caetano, T., van den Hengel, A., & Wang, Z. "A Hybrid Loss for Multiclass and Structured Prediction", IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI), 37(1): 2-12, 2015.

*Shu, X., Porikli, F,. & Ahuja, N. Robust orthonormal subspace learning: efficient recovery of corrupted low-rank matrices, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014

*Singh, A., Porikli, F., & Ahuja, N. Super-resolving noisy images, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014

Sokeh, H., Gould, S., & Renz, J. "Determining Interacting Objects in Human-Centric Activities via Qualitative Spatio-temporal Reasoning". In Proceedings of the Asian Conference on Computer Vision (ACCV), 2014.

*Soni, J. H., & Porikli, F. Recycled linear classifiers for multiclass classification, IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2014

Sünderhauf, N., McCool, C., Upcroft, B., Perez., T. (2014). Fine-Grained Plant Classification Using Convolutional Neural Networks for Feature Extraction. Proc. of ImageCLEF Workshop, International Conference of the CLEF Initiative, Sheffield, UK.

*Tao, L., Porikli, F., & Vidal R. Sparse dictionaries for semantic segmentation, European Conference on Computer Vision (ECCV), 2014

*Topkaya, S., Erdogan, H., & Porikli, F. Counting people by clustering person detector outputs, IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS), 2014 (Best Poster Award)

*Wang, Y., Jodoin, P. M., Porikli, F., Konrad, J., Benezeth, Y., & Ishwar, P. CDnet 2014: an expanded change detection benchmark dataset , IEEE Change Detection Workshop (CDNet) in conjunction with CVPR, 2014

Yang, J., Li, H., & Jia, Y. Optimal Essential Matrix Estimation via Inlier-Set Maximization. ECCV (1) 2014: 111-126

*Zhu, G, Porikli, F,. Ming, Y, & Li, H. Lie-Struck: affine tracking on Lie groups using structured SVM, IEEE Winter Applications and Computer Vision Conference (WACV), 2015

***Centre outputs (not including outputs from other ARC-funded projects)**

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 Reseach 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
k	Thousand
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
SRPs	Strategic Research Priorities
UoA	University of Adelaide
VOS	Vision Operating System

Definitions

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.
SLAM (Simultaneous Localisation 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 class from those of another class.
Visual servoing	the use of information from one or more cameras in real-time to guide a robot in performing a task.

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 centre 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 (ACRV) is to give robots the gift of sight, our logo incorporates the most important elements of the eye.

We (ACRV) sit 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 set the spark 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 too will the capabilities of robots.

ACRV
AUSTRALIAN CENTRE
FOR ROBOTIC VISION

ACRV Partners

Design by Studio 55

Photography by Tony Phillips, David Burke and Paul Fenn

Australian Centre for Robotic Vision
Queensland University of Technology
2 George Street, Brisbane QLD 4000 AUSTRALIA
P. +61 7 3138 7549
E. info@roboticvision.org
W. www.roboticvision.org

ACRV can also be found on: