

Annual Report 2017



AN ARC CENTRE
OF EXCELLENCE



Australian Government
Australian Research Council

Who We Are

At the Australian Centre for Robotic Vision, we are striving to lead the world in the new discipline of robotic vision by applying computer vision to the domain of robotics. We are the largest robotic vision group in the world with more than 200 people and \$25.6m funding over seven years.

Our Vision

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



Our Mission

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

We have created a Centre with an exciting high-energy **culture** that supports us to do impactful **science**, by **integrating** robotics and computer vision. We **engage** with people about robotic vision technologies and their impact, and we will **transform** the world by solving critical innovation challenges.



Robotic Vision **expands the capabilities** of robots, allowing them to see and understand the world in which they are working. We believe it is the key technology that will allow robotics to change the way we live and work.



CODE-A-BOT image courtesy of QUT media.

While robotics is about machines that **interact** with the **physical world**, computer vision is about **analysing and understanding** the world through images.

Already we are looking to apply our technologies to **solve real challenges** in the monitoring and protection of the natural and built **environments**, the provision of **healthcare** in hospitals and in the home, sustainable food production, and efficiently harnessing our natural **resources**.



Values

Our values represent our culture and the way we do things

Create

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



Collaborate

We work together and partner to solve grand challenges



Empower

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



Impact

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



About this report

REPORT DESCRIPTION

Our report covers the activities of the ARC Centre of Excellence for Robotic Vision for the 2017 calendar year. It forms part of our official reporting and accounting requirements to the Australian Research Council. Our activities encompass research, training, outreach, industry engagement, operations and finance.

AIMS OF THE REPORT

As a publicly funded organisation, our report demonstrates to our stakeholders and the broader community, the impact we are having in the new field of robotic vision. We provide an overview of our operations in 2017 and our progress on meeting key performance indicators (KPIs), specifically with reference to:

- / research performance p. 44
- / research training and professional development p. 73
- / international, national and regional links and networks p. 64
- / end-user engagement p. 57
- / financial performance p. 94
- / governance p. 90
- / contributions to national benefit p. 52

As well as reporting on our performance, we also report on the elements that make our Centre a leading innovator in the field of robotic vision including our culture, our people and our achievements.

Our History

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.


Centres are hubs of strategic collaboration between universities, publicly-funded research organisations, other research bodies, governments and businesses in Australia and overseas. We strive to produce outstanding research that has social, economic and scientific significance.

Distinguished Professor Peter Corke (QUT) and Professor Robert Mahony (ANU) identified the opportunity to bring the

disciplines of robotics and computer vision together when the ARC invited applications in 2012 for new Centres of Excellence. A team of chief investigators formed in 2013 and plans were developed for a Centre for Robotic Vision. After finalising our collaborative partner agreements in mid-2014 the Australian Centre for Robotic Vision was born. We formally launched in 2015, and have grown to involve more than 200 people across the world.


The Australian Centre for Robotic Vision is an unincorporated collaborative venture between interdisciplinary researchers from leading Australian and international research organisations shown below.

World map highlighting partner locations




Contents

SECTION 1: About Us.....	1
Director's Report	3
Centre Performance.....	7
Recognising our People	9
SECTION 2: Research Performance.....	13
Research Structure	14
Research Management	14
Activity Plan	15
SENSING – Robust Vision (RV).....	19
UNDERSTANDING – Semantic Representations (SR).....	27
ACTING – Vision & Action (VA).....	31
LEARNING – Visual Learning (VL)	37
TECHNOLOGY – Algorithms & Architecture (AA).....	39
SECTION 3: Research Impact	45
Where we live	46
How we live	46
Means to live	47
SECTION 4: National Benefit.....	48
National Benefit.....	49
Science and Research Priorities	50
SECTION 5: Engagement.....	54
End-User Engagement	55
Communication & Media Engagement.....	59
Commercialisation	61
SECTION 6: People	65
Research Training and Professional Development	67
Gender Diversity.....	70
Mentoring.....	72
Education.....	72
Our People	74
Staff Profiles	76
SECTION 7: Governance.....	82
Committees	83
SECTION 8: Finance and Operations	91
Financial Performance.....	92
OUTPUTS	95
2017 Publications.....	95
Glossary	99
Meet our Robot Family.....	100
The Story of Our Logo.....	101



Front cover: Centre researchers competing in Japan to win the Amazon Robotic Challenge

SECTION 1 | About Us



▲ Centre researchers with CartMan, the Amazon Robotics Challenge winning Cartesian Manipulator robot.

Australia to get its first robotics roadmap

Australia is set to have its very first robotics and computer vision roadmap with the Australian Centre for Robotic Vision driving this important initiative.

“Australia’s first robotics and computer vision roadmap is about more than just making industries more automated,” said Dr Sue Keay, the Centre’s Chief Operating Officer. “The strength of our robotics and computer vision technologies will drive the transformation of existing industries and create whole new industries. This has significant impact for the future of Australia’s workforce as well as making sure we have vibrant, competitive and sustainable industry sectors.

“Australia has a real advantage in robotics and computer vision at this point in time. As a national research centre and the world’s first research centre specialising in robotic vision, we want to make sure Australia really benefits from these new technologies,” Sue said.

The Centre has identified six sectors that are fundamental to the Australian economy: resources, the built and natural environments, manufacturing, services, healthcare, and defence. From October 2017 to January 2018, the Centre ran a series of workshops across these sectors, in addition to calling on submissions from interested Australians and Australian companies, to contribute to the robotics and computer vision roadmap.

“Australia is currently ranked 30 in global automation according to the International Federation of Robotics, despite our high levels of niche manufacturing. So, on the surface, while we are leaders in many areas of robotics, it appears Australia is lagging on this measure and that’s of concern.

“These are technologies that can solve many of the key challenges facing Australia in terms of productivity and sustainability,” Sue said. “Our robotics and computer vision roadmap will give us a clearer idea of what the industry looks like in Australia and how best to support take-up by existing companies and to create new companies based on these technologies.”



The USA released its third roadmap for robotics in **2016**, while the UK, EU, Korea, Japan and Singapore also have robotic roadmaps in place. The Centre will release Australia’s first robotics and computer vision roadmap in **mid-2018**.

◀ The Centre ran 6 national consultation workshops for the robotic vision roadmap.



Director's Report

This has been a great year for the Centre. Our profile, nationally and internationally, continues to rise – driven by our scientific output, widespread engagement in various communities and significant accolades.

In September, our funding agency, the Australian Research Council (ARC), conducted a mid-term review of our Centre. We are proud of the results and feedback received, and that our funding will continue through to the end of 2020. We are also pleased to welcome a new partner to the Centre, The University of Toronto, Canada. Australia can learn a lot from the emergence of robotics in Canada. We are both large countries with small population densities and the Canadian government is strategically investing in robotics and Artificial Intelligence to good effect.

SCIENCE

In our 2016 report, I projected 2017 as the year our research would get real traction, and thanks to the focus and commitment of our researchers, our professional staff and our partners, we have been able to deliver on these expectations.


Prior to our mid-term review, the international standing of our research was assessed by independent reviewer Professor Henrik Christensen from UC San Diego, an expert in the field of robotics and vision. He stated in his review that, “The number of publications in high impact archival journals and at the best scientific conferences is very impressive. In particular in the robotics field there is a clear recognition of the brand/identity provided by the Centre.” Professor Christensen went on to say, “The research performed by the Centre is of a very high

quality. The research is published at the best conferences (ECCV, ICCV, CVPR, ICRA, IROS and RSS) and there is a good transition to publications in archival journals. Slowly joint publications, across participating institutions, are emerging from the Centre. There is no doubt the research is disseminated very well to the broader science community.”

The ARC review panel also commended our research environment, including our accommodations, infrastructure and equipment.



Centre Researchers at the 2017 International Conference on Robotics and Automation (ICRA) held in Singapore. Pictured L to R: PhD researcher Fahimeh Rezazadegan (QUT), Centre Director Peter Corke (QUT), Research Affiliate Chris McCool (QUT), Chief Investigator Tom Drummond (Monash), Chief Investigator Michael Milford (QUT), Research Affiliate Chris Lehnert (QUT), Research Fellow Juxi Leitner (QUT), QUT Research Fellow David Hall, Research Fellow Feras Dayoub (QUT), Research Fellow Alireza Khosravian (Adelaide), Research Affiliate Leo Wu (QUT). *Image courtesy of Michael Milford*



We have worked hard since our inception to create an environment that supports our people, and enables high-quality research and innovation. Our culture encourages collaboration and communication, exploring challenges from all sides and perspectives to find transformational solutions, and developing knowledge leaders for industry and academia.

CULTURE

Gender Diversity is one imbalance we are addressing in our field. We have been working to encourage more women to engage academically and professionally in STEM (Science, Technology, Engineering and Mathematics). At our annual conference we launched our Robotic Vision Gender Equity Plan (available on our website) and are fortunate to have a number of great role models in this area (see Gender Diversity p. 70).

I am keenly aware of the challenges we face in competing for talent, attracting high-calibre researchers and retaining them in our team. We know that a strong and supportive culture is an important strategy to address these challenges.

INTEGRATE

The “Integrate” component of our Centre’s strategy is about bringing the disciplines of robotics and computer vision together to create new robotic vision technologies.

An example of our Integrate approach was our win at the Amazon Robotics Challenge. We integrated several novel centre-developed technologies to the problem of robot picking. These included an innovative few-shot learning approach, a novel robot, and a very effective dual-mode grasping tool.

We integrate people and organisations as we strive to build a critical mass of robotic

vision capability within the Centre. We also endeavour to create international networks and reputation by building an international research community in this new area.

Strategies include our annual summer school (RVSS) with an increasing global contingent; high profile and well attended workshops at top international conferences; and exchange visits with our international research partners.

ENGAGE

We continued our industry engagement activities in 2017, through conferences, partnerships and exchange opportunities.

We worked with specialist communications consultants, STEMmatters, on a year-long media campaign. With 60 releases contributing to more than 181 media articles we promoted the Centre, our researchers and our work to a wide audience.

Our Centre headquarters at QUT hosts more than 1,000 visitors every year. This year we welcomed some very distinguished guests at all the Centre's nodes.

Dr Rodney Brooks, CTO of ReThink Robotics visited our QUT node while in Brisbane for the MIT Global Entrepreneurship Bootcamp.

Professor Ron Arkin, from our research partner, Georgia Tech in the US, spent a sabbatical at QUT from July to September. Our ANU node hosted Professor Gregory Chirikjian from John Hopkins University and Dr Guillermo Gallego from ETH Zurich. Monash University welcomed Professor Anthony O'Neill from Newcastle University. Our University of Adelaide node hosted Professor Stefan Roth from TU Darmstadt and a visit from the Hon. Christopher Pyne MP, as the Minister for the Defence Industry.


We have hosted many high-school visits and, at QUT, we created an interactive digital game CODE-A-BOT and participated in Robotronica, a biennial robotics festival, which gives the general public the chance to learn more about robotics. 2017 also saw the launch of the Robot Academy, a platform to deliver robotics education to the world.

Public engagement is increasingly important as misinformation, uncertainty and concern about technology in general, but robotics, AI and automation in particular, continues to grow. Our strategy is to engage with

everybody we can: primary school students, business owners, visiting politicians, academics and enthusiasts. We share with them our work, our vision for the future and how they can be a part of the journey to create robots that see for the benefit of all.

TRANSFORM

2017 saw the launch of our first two start-ups, Nuarda and AlphaOne.AI. Start-ups are an important alternative means to transfer our research outcomes into the market place and create transformational change in the real world. Centre researchers are leading an international collaboration that is creating a new class of medical robotics that will make keyhole surgery safer. This new robotic imaging system makes it easier for surgeons to see inside the tiny spaces in the human body, allowing them to visualise soft tissue in real time and in 3D, and ultimately delivering better results for the patients (see p.18).



Our strategy is to engage with everybody we can: primary school students, business owners, visiting politicians, academics and enthusiasts. We share with them our work, our vision for the future and how they can be a part of the journey to create robots that see for the benefit of all.



Image courtesy of the Queensland Museum

Other work across the Centre includes:

- / a partnership between **ANU researchers and an asparagus farm in Cowra** to develop an autonomous harvesting robot for green asparagus
- / a range of externally funded projects from **nVidia, DST Group, Mining 3, Caterpillar, Ergon and others.**
- / a project involving social scientists where we **apply robotic vision to Softbank Robotics, Pepper social robot platform,** supported by a \$1.5 million Advance Queensland grant. See p. 42.

THE YEAR AHEAD

We are well-placed as the world leader in robotic vision and the coming year is about building on this foundation. Our work to the end of 2020 will be focussed on tangible applications of robotic vision. We have incorporated feedback from our 2017 reviews into our project portfolio, changing the way we operate to achieve increased impact.

Next year we will complete, and launch, Australia's first national roadmap for robotics and computer vision. We are also heavily involved in the organisation of the IEEE International Conference on Robotics and Automation (ICRA). The conference will be held in Brisbane in May 2018 – the first time ever in the southern hemisphere – and should attract over 2,500 top international robotics researchers.

Thank you to all our colleagues – within the Centre and beyond - for your contributions to ensure the Centre achieves impact. We are scaling into an enterprise that is delivering on our mission to develop new robotic technologies that expand the capabilities of robots. We must ensure these technologies continue to positively contribute to the sustainable well-being of people and the environments we all live in.

Peter Corke // Centre Director



Centre Performance



We measure our performance by how well we:



generate internationally impactful science



build new research capacity in robotic vision



develop highly skilled knowledge leaders



create a robotic vision community



devise technologies that will transform important Australian industries



engage with people about our technologies



solve innovation challenges and transform the world



We envisage a world where robots can:

understand tasks and learn how to **adapt to change**

repeat a task after having watched a person perform it once

recognise people, understand what they are doing and **help with what they need next**

fly underneath large infrastructure assets such as bridges, inspect them autonomously and monitor changes

handle goods capably using **'superhuman' hand-eye coordination**

reduce the cost of healthcare, through **low-cost devices**

operate in all weather and lighting conditions and recognise where they are and what is going on around them

navigate using signs designed for humans

recognise complex objects after having seen them just once and remember how to interact with them

engage in **lifelong continuous learning** about their immediate environment and the world

share not just physical work but **what they see and know about the world**

be deployed on very large scales, automatically discovering and using the sensory and computational resources around them



Since we launched the Centre in 2014, we have:

- / **placed first in the 2017 Amazon Robotics Challenge**, using vision to complete a complex robotics task while competing against 15 other research groups from around the world
- / **created an underwater robot that culls crown of thorns starfish**, addressing a looming ecological crisis, without the costs and risks associated with using human divers (*COTSbot*)
- / **created an agricultural robot that can visually identify the difference between weeds and crop plants** precisely applying the correct dose of the appropriate herbicide directly onto the weed while reducing herbicide use by up to 70% (*AgBot II*)
- / **created a fruit-picking robot** that harvests the right fruit, the right way, 24/7 (*Harvey*)
- / **placed first** in the international Pattern Analysis, Statistical Modelling and Computational Learning (PASCAL) Visual Object Classes (VOC) challenge in 2017
- / **launched free university-level courses in robotic vision**, with participation to date from more than 100,000 students, based in over 100 countries
- / **held robotic vision workshops** at top robotics and computer vision conferences each year to raise the profile of this new field
- / **actively participated** in the debate about robots, jobs and society



▲ Centre CI Jon Roberts introduces Queensland Minister for Science, the Hon Leeanne Enoch, to Softbank Robotics Pepper robot at Robotronica.

Recognising our People

This has been our most successful year yet. We welcomed four new Chief Investigators: Professor Jonathan Roberts, Dr Matthew Dunbabin, Dr Niko Sünderhauf and Associate Professor Tat-Jun Chin.

Alex Zelinsky, Chair of our advisory committee, was admitted as an Officer of the Order of Australia (AO) in the Queen's Birthday Honours in June 2017. This recognises his remarkable achievements throughout the course of his 30+ year career. Alex is currently Australia's Chief Defence Scientist and head of the Defence Science and Technology Group.

Dr Alex Zelinsky receiving the Order of Australia from The Governor-General, His Excellency General the Honourable Sir Peter Cosgrove AK MC (Retd).
Image courtesy of Eddison Photographic Studios ▶



Hugh Durrant-Whyte, a member of our advisory committee, was awarded the prestigious MA Sargent Medal by Engineers Australia's Electrical College Board. The award is named in honour of Dr Michael Anthony Sargent, an outstanding Australian electrical engineer. Hugh is the Chief Scientific Advisor, Ministry of Defence, United Kingdom. He is on leave from his role as a Professor and ARC Federation Fellow at The University of Sydney.

◀ Professor Hugh Durrant-Whyte receiving the MA Sargent Medal at the Engineers Australia Awards Ceremony. *Image courtesy of Engineers Australia*

Associate Investigator Dr Anders Eriksson was awarded an ARC Future Fellowship valued at over \$800,000 over a four-year period, to undertake research in the fields of computer vision and optimisation.

Anders' research will have a significant impact on the design of numerical algorithms for solving a wide range of problems in computer vision, virtual reality and robotic navigation.




Image courtesy of QUT

Partner Investigator Professor Andrew Davison,

from Imperial College London, was elected to the Fellowship of the Royal Academy of Engineering (RAEng), one of the highest honours an engineer can receive in the UK.

Andy and his team at Dyson Robotics Lab are leaders in the field of real-time 3D computer vision.



Image courtesy of Imperial College London



In December, **Centre Director Peter Corke**

received the 2017 Australian Teacher of the Year award. This recognises his transformational contributions to robotics education at QUT, and his highly accessible online learning activities and resources which have led to a real-world, project-based curriculum, open-source software platform, a series of Massive Open Online Courses (MOOCs), publication of a textbook, delivery of popular online lectures and the establishment of the innovative online *Robot Academy*. Peter was also admitted as a Senior Fellow of the Higher Education Academy (HEA), an independent, not-for-profit organisation committed to world-class teaching in higher education. In November, Peter was elected as a Fellow of the Australian Academy of Technology and Engineering (ATSE) in recognition of his many contributions to computer vision and robot navigation.

◀ Centre Director Peter Corke receives the “2017 Australian University Teacher of the Year” award from the Australian Government’s Department of Education and Training.
Image courtesy of Phillipa Carr



Dr Sue Keay, Chief Operating Officer.
Image courtesy of Tim Macuga

Centre COO Dr Sue Keay founded Australia's first Women in Robotics group and was named a Superstar of STEM, promoting female role models in the field of Robotics.

Centre Research Affiliate

Anjali Jaiprakash was acknowledged by Robohub as one of the “World’s 25 Women in Robotics you should know about”. She has also been featured on MIT Technology Review’s 2017 list of ten outstanding innovators under 35 Asia Pacific. The list highlights the region’s most creative and bold young inventors, entrepreneurs and researchers, celebrating the rising stars that have dedicated their careers to purpose-driven innovation.

Anjali was also named Queensland Young Tall Poppy, honouring her for excellence in the field of medical robotics, where she is developing a vision-based robotic leg-manipulation system for knee arthroscopy and a retinal diagnostic system to detect the early onset of blindness.




Image courtesy of QUT

Dr Damien Teney and Professor Anton van den Hengel

of the University of Adelaide, along with Centre PhD researcher **Peter Anderson** from ANU and their team, placed first in the VQA 2.0 challenge. The challenge gives participants an image and an accompanying natural language question about it, which their AI system must provide an accurate natural language answer.

Anton and his team at the Australian Centre for Visual Technologies also won the South Australian Science Awards' Excellence in Research Collaboration prize, for their work in developing a world leading intelligent medical device, with LBT Innovations Ltd. The awards showcased the critical importance of science and research for the development of industry and our society.



Professor Anton van den Hengel

Other esteem measures garnered by Centre researchers include:

- / **CI Michael Milford** was promoted to Professor
- / **Centre Director Peter Corke** was awarded the title of Distinguished Professor by QUT, recognising his truly outstanding achievements in the emerging field of Robotics. QUT's Vice-Chancellor **Professor Peter Coaldrake** acknowledged Peter's accomplishments as an enthusiastic and innovative teacher of undergraduates, through to an influential researcher on the international stage.
- / Research Fellow and SR1 Project Leader, **Dr Basura Fernando** won the Best Reviewer Award at 2017 International Conference on Computer Vision (ICCV), held in Venice, Italy.
- / Research Fellow, **Dr Anoop Cherian** received the Outstanding Reviewer Award at the 2017 Conference on Computer Vision and Pattern Recognition (CVPR), held in Honolulu, Hawaii.


SECTION 2

Research Performance



Research Structure

Our Centre will achieve breakthrough science and technology in robotic vision by addressing five key research programs. Each program operates a portfolio of projects all focused on developing and integrating robotic vision into a viable technology that will enhance our lives and extend our capabilities.



Research Management

The Centre's research projects are each headed by a Chief Investigator or Senior Research Fellow. Each individual project has a clear objective, with milestones and deadlines. A project team comprises a leader, with a staff of Research Fellows (RF), Associate Investigators (AI), PhD researchers, and Chief Investigators (CI). The litmus test for any Centre project is whether it will develop a technology that demonstrates the capability of robotic vision that can be demonstrated by a robot, ensuring it will deliver an outcome that is relevant to the Centre's vision.

All new projects are established through our Research Committee, which determines the project's duration, staffing and other resourcing matters. Project leaders represent their project teams in Research Committee meetings, presenting results and updating the committee on progress. The Research Committee's role is to ensure project progress, resolve resourcing issues and determine when projects should continue or close.

Activity Plan for 2018

CULTURE

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



STRATEGIC OBJECTIVES

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

KEY TASKS

- 1 Maintain a full complement of Research Fellows
- 2 Aim for 90% retention of PhD enrolments
- 3 Provide knowledge leadership training to all early career researchers to help in their career development
- 4 Hold one annual symposium each year to build Centre culture
- 5 Run one Robotic Vision Summer School (RVSS) each year as an international PhD recruitment event and to train our new PhD researchers. Continue building the international profile of this event
- 6 Hold monthly meetings of the Research Committee and quarterly face to face meetings
- 7 Centre Executive meet monthly via videoconference plus quarterly face to face meetings, held in conjunction with the Research Committee meeting
- 8 Visit international partners and host visits of international researchers.

TRANSFORM

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



STRATEGIC OBJECTIVES

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

KEY TASKS

- 1 Engage Centre's Advisory Board at least twice yearly
- 2 Provide knowledge leadership training to all early career researchers to help in their career development
- 3 Present on robotic vision at key industry events
- 4 Launch robotic vision roadmap for Australia.

SCIENCE

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



STRATEGIC OBJECTIVES

/ Deliver internationally recognised research in robotic vision

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

KEY TASKS

RESEARCH PROJECT: MANIPULATION AND VISION

- 1 Achieve continuous integration of visual feedback during grasping in table-top settings. Involving machine learning but simple end effectors (pincher or suction), ie. limited object generality (around 30 types)
- 2 Utilise force-driven interaction for deformable objects on simple gripper
- 3 Assess Reinforcement Learning + Heuristic learning of how to interact with novel objects.

RESEARCH PROJECT: ROBOTIC VISION VIRTUAL EVALUATION TESTBED

- 1 Object detection and semantic segmentation in open-set conditions
- 2 Uncertainty estimation
- 3 Active learning, continuous learning.

RESEARCH PROJECT: ROBOTIC VISION DEPLOYMENT AND EVALUATION

- 1 Implement a systematic study to describe the discrepancy between the current dataset-based evaluation metrics of learning-based predictive models and the true performance of these models onboard real robots. This includes testing state-of-the-art classification, detection and segmentation models onboard robots in realistic settings
- 2 Design methodologies to quantify the performance of the above systems in a variety of real-world environments.

RESEARCH PROJECT: LEARNING

- 1 Use self-supervised learning of 3D models from 2D, 2.5D images
- 2 Achieve uncertainty in 0-shot, open-set and few-shot learning (version 1 of dataset)
- 3 Develop weakly supervised instance segmentation
- 4 Trial methods to compress/simplify Deep Learning models to run in embedded system.

RESEARCH PROJECT: ROBOTS, HUMANS AND ACTION

- 1 Achieve state-of-the-art results on recognising human activities on public benchmarks
- 2 Develop models that compose complex activities in terms of discrete events and short-term actions
- 3 Ability to track object and agents (humans) across time to link activities.

RESEARCH PROJECT: SCENE UNDERSTANDING

- 1 Robustify and properly integrate segmentation with object-based mapping to allow for previously unseen objects
- 2 Incorporate physical (static) relations (supporting, inside, etc)
- 3 Extend static scene understanding to dynamic objects including people (coarse)
- 4 Demonstrate CNN-based single-view depth and pose estimation integrated into full localisation and mapping capability.

RESEARCH PROJECT: VISION AND LANGUAGE

- 1 Grounded visual conversation: Agent trained to participate in a grounded conversation by incorporating progress made towards a fixed task in the reward
- 2 Learning to learn to talk: Use meta-learning tech to train an agent to learn to exploit the capabilities it has and the data available to it. Similar to the Neural Turing Machine but for robots.

RESEARCH PROJECT: FAST VISION-BASED NAVIGATION IN UNSTRUCTURED ENVIRONMENT

- 1 Adaptive perception for visual navigation in unstructured environments
- 2 Techniques for integration of free space affordance (learning based methods) with optical flow to enhance control capabilities.

INTEGRATE

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



STRATEGIC OBJECTIVES

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

KEY TASKS

- 1 Organise robotic vision workshops at key international conferences
- 2 Hold an annual robotic vision summer school targeted at international attendees
- 3 Continue to build network of industry contacts and maintain a customer relationship management system.

ENGAGE

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



STRATEGIC OBJECTIVES

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

KEY TASKS

- 1 Promote our Robotic Vision resources hub
- 2 Promote the Robot Academy
- 3 Host visits and tours of robotic vision facilities by government, industry and the community (including school groups)
- 4 Include popular science articles on Centre work in targeted publications
- 5 Create 4 x media releases by partners related to robotic vision
- 6 Create 4 x public lectures on robotic vision.

New Class of Medical Robotics to Make Keyhole Surgery Safer

An international collaboration (headed by a partnership between QUT and the Centre) is endeavouring to develop new robotic imaging systems that will make keyhole surgeries safer and improve patient recovery time.

Project leader and renowned orthopaedic surgeon Professor Ross Crawford said the new robotic imaging system will allow surgeons to track the position of soft tissue in real time, and in 3D for the first time.

Unlike current medical imaging tools, which can only track the position of bone and medical tools, the new system will combine state-of-the-art miniaturised stereo cameras, 4D ultrasound sensing and artificial intelligence that will enable the surgeon to also navigate tendons and ligaments.

“Combining vision and ultrasound can provide us with more information, which we can use to identify and differentiate the objects within the surgical site,” Ross said.

“Working with a dynamic 3D model rather than a flat image on a screen is a real game-changer in terms of accuracy—it will give the surgeon precise knowledge of how deep the objects in the surgical site really are, (and) will vastly improve both visualisation and access issues, making keyhole surgery more accurate than ever before.”

“It’s the patients throughout the world who will be the biggest beneficiaries, through better outcomes and better access to healthcare.”

The technology is first being trialled on knee arthroscopy surgeries, but the hope is that it can be modified to assist with other surgeries including hip, shoulder, abdominal, and eventually even heart.

The research is supported by an AU\$996,000 Australia-India Strategic Research Fund grant from the Australian Government. QUT will collaborate with researchers from the Indian Institute of Technology-Madras, All India Institute of Medical Sciences-Delhi, Perfint Healthcare Indian Institute of Technology-Kharagpur, Manipal University, and the University of Adelaide.



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

Sensing: Robust Vision

The Robust Vision program is about “Sensing”, that is, creating robots that can see in all conditions. The key question we are addressing is, “How can innovations in existing computer vision, robotic vision techniques and vision sensing hardware enable robots to perform well under the wide range of challenging conditions they will encounter and how can we apply computer vision technology in the real world?”

To answer this, the Robust Vision program is developing robotic vision algorithms and novel vision hardware for robots to perceive the environment around them in a way that will allow them to act purposefully and perform their tasks, in a variety of conditions including high glare, fog, smoke or dust, rain, sleet, wind or heat.

The Robust Vision program is also working with innovative sensing hardware, developing it further to support robots operating under challenging viewing conditions such as poor light, or through partial obscuration and hyper-spectral cameras.


RV projects

RV1—ROBUST ROBOTIC VISUAL RECOGNITION FOR ADVERSE CONDITIONS

Research leader: Michael Milford (CI)

Research team: Hongdong Li (CI), Jonathan Roberts (CI), Chunhua Shen (CI), Niko Sünderhauf (CI), Abigail Allwood (NASA JPL), Nigel Boswell (Caterpillar), David Thompson (NASA JPL), Juxi Leitner (RF), Chuong Nguyen (RF), Sourav Garg (PhD), Sean McMahon (PhD), Medhani Menikdiwela (PhD), James Mount (PhD), James Sergeant (PhD)

Project aim: How we can use and innovate on existing computer vision and robotic vision techniques to make them perform well in the challenging conditions robots will encounter in real-world situations? To answer this question, the Robust robotic visual recognition for adverse conditions project is developing robust algorithms that solve the fundamental robotic tasks of place and object recognition under challenging environmental conditions, such as darkness, adverse atmospheric and weather conditions, seasonal change and translating them into real-world industrial applications.



▲ A kangaroo captured in darkness using a low light camera at the Centre’s Robotic Vision Summer School retreat in Kioloa, NSW.

RV2—NOVEL VISUAL SENSING AND HYBRID HARDWARE FOR ROBOTIC OPERATION IN ADVERSE CONDITIONS AND FOR DIFFICULT OBJECTS

Research leader: Chuong Nguyen (RF)


Research team: Peter Corke (CI), Robert Mahony (CI), Michael Milford (CI), Donald Dansereau (Research Affiliate), Adam Jacobson (PhD), Dan Richards (PhD), James Sergeant (PhD), Dorian Tsai (PhD)

Project aim: The key question being addressed in this project is “How can existing and new special vision hardware be exploited and developed in conjunction with new algorithm development to enable robotic vision systems to operate well under the wide range of challenging conditions that robots and applied computer vision technology encounter in the real world?”


To address this, the project advances the performance of robot vision algorithms using low-light cameras, rotational filters, hyperspectral cameras and thermal cameras to improve robot autonomy under all viewing conditions.

In developing new algorithms that exploit the particular advantages these innovative cameras offer, we enabled new performance benchmarks in tasks such as scene understanding, place recognition and object recognition. These hardware-software solutions can deal with challenges such as reflections, transparency and low light conditions and can also be applied to all Centre projects involving visual sensing, making their own systems’ performance under adverse conditions more robust.

Previous project RV3 – Robustifying place recognition and SLAM with semantic scene understanding was merged into SR3 – Scene understanding for robotic vision in 2017.



▲ The challenge of perception and navigation in underground and above ground environments.




▲ Visual surface-based positioning method for vehicles.

Key Results in 2017

The first stage of the robotic vision positioning system for underground mining vehicles in collaboration with Caterpillar, Mining3 and the Queensland Government was completed.

This project is part of an Advance Queensland Innovation Partnerships Grant and an associated publication “Enhancing Underground Visual Place Recognition with Shannon Entropy Saliency” was presented at the 2017 Australasian Conference on Robotics and Automation (ACRA) in December. Researchers CI Michael Milford and Associated RFs Adam Jacobson and Fan Zeng completed three sets of field trips across two different underground mining sites spread across Australia (mine site names withheld for confidentiality reasons).



▲ Some of the extensive media coverage of the project received around the world.



▲ (l to r) Associated RFs Fan Zeng and Adam Jacobson, CI Michael Milford, Research Affiliate Thierry Peynot.



▲ (l to r) Former Caterpillar Engineer Peter Beasley, CI Michael Milford and PhD Researcher Adam Jacobson hundreds of metres underground at a Queensland mine site.

Detecting hazards in challenging conditions

Driving in adverse conditions—like heavy rain, fog, or poor light can put even the most experienced driver to the test. As the idea of autonomous vehicles becomes more of a reality, we need to know the on-board navigation systems are up for these challenges, which can present themselves at any time and with little warning.

Water hazard detection, in particular, is crucial for autonomous navigation of outdoor robots. Traditional methods for detecting water hazards include stereo matching, texture-colour classification and polarisation difference. While they work well with clean images, they nevertheless have a limited effective range. In Centre project RV2—*Novel visual sensing and hybrid hardware for robotic operation in adverse conditions and for difficult objects*, Centre Research Fellow Chuong Nguyen and his team, extended these methods to work with challenging wet and raining conditions where existing methods fail.

The team proposed a new stereo matching for wet-muddy ground, along with extending u- and v-disparities to better detect the ground plane. Finally, they incorporated a classification with temporal detection to improve the accuracy of detecting water puddles in both on- and off-road conditions.

The team validated their approaches against other detection methods, and released their video and image data sets for independent studies. The resulting system can detect water hazards up to 100m distant.



Daniel Richards was confirmed as a PhD researcher in early 2017 and continued to work on low light vision research, including research trips to the Defence Science and Technology (DST) Group and work on developing custom low light hardware systems.



▲ Combined semantic-traditional place recognition pipeline.

PhD researcher **Sourav Garg's** paper "Improving Condition- and Environment-Invariant Place Recognition with Semantic Place Categorisation" co-authored with **Associated RF Adam Jacobson, Swagat Kumar (Tata Consultancy)** and **CI Michael Milford** was accepted and presented at IROS2017.

Ordinary or specific place recognition involves reconciling current sensory input on a robot or mobile platform with previously gathered or stored maps, and recognising where you are within that map.

Place categorisation on the other hand involves recognising the type of place - kitchen, supermarket, ocean, forest and so on. Following on from previous work in this area, in this work led by Sourav Garg and in collaboration with Swagat Kumar from Tata Consultancy, we leverage the powerful complementary nature of place recognition and place categorisation processes to create a new hybrid place recognition system that uses place context to inform place recognition.

PhD Researcher James Mount also presented an accepted paper at ACRA 2017 titled “Image Rejection and Match Verification to Improve Surface-Based Localisation” co-authored with **CI Michael Milford**.

Positioning (or localisation) is a crucial capability for almost any vehicle or robot that moves. In this research, led by James, we developed techniques for localising with very high accuracy and low latency off the surface of the road using cameras, both during the day and at night. This sort of positioning system would complement (rather than completely replace) other positioning systems including GPS.

James also created and ran an all-day show titled “Self-Driving Cars – Making an Ethical Choice” at QUT’s Robotronica in August. The show ran hourly for a total of seven shows to over a thousand audience members. The premise of the show was simple, to show and discuss, with an actual miniature autonomous car, what the effects of various self-driving car ethical considerations are. To make the show interactive, the audience was able to vote live via their smartphone or tablet, on what decision they thought the robot should make. The day was a huge success and the researchers continue their efforts of engagement with industry, government and the public on the impacts and opportunities of self-driving cars and the artificial intelligence technology that drives them.



▲ PhD Researcher James Mount presenting to the crowd at QUT’s Robotronica



▲ Full house for the show at QUT’s Robotronica, totalling more than 1000 people over the day



▲ CI Michael Milford presenting at the “Closing the Technology Loop on Self-Driving Cars” forum.

CI Michael Milford hosted a major self-driving car forum at QUT titled “Closing the Technology Loop on Self-Driving Cars” in June. More than 100 representatives from state and local government, industry and start-ups discussed and debated autonomous vehicles.

Developing New Specialist Camera Equipment

Light field cameras are a new paradigm in imaging technology with the potential to greatly augment computer and robotic vision. Unlike conventional cameras that only capture an image from a single perspective, light field cameras instantaneously capture multiple images of the same scene from slightly different perspectives, which may help deal with occlusions, highly reflective surfaces, and refractive objects that often break most modern robotic vision techniques.

Unfortunately, there is a scarcity of commercially-available light field cameras appropriate for robotic applications—none deliver light fields at video frame rates within a sufficiently small size and mass. Creating a full camera array comes with synchronisation, bulk and bandwidth issues.

PhD Researcher Dorian Tsai, Research Engineer Steve Martin, Research Affiliate Donald Dansereau, Research Affiliate Thierry Peynot and Centre Director Peter Corke developed a light field camera using mirrors of different shapes and orientations to reflect the scene into an upwards-facing camera, to create an array of virtual cameras with overlapping field-of-view at specified depths. Our custom mirror-based light field camera adapter, called the MirrorCam, allows video frame rate light fields with a design that is cheap, simple in construction and accessible.

The MirrorCam was mounted onto a robotic arm and used in the first-ever visual servoing experiments with a light field camera. We showed that our method outperformed conventional monocular and stereo image-based visual servoing under field-of-view constraints and occlusions. Current research is focused on how to use these developments to servo towards, and ultimately pick up and manipulate transparent objects using robotic vision.



▲ MirrorCam mounted on the Kinova MICO robot manipulator. Nine mirrors in a 3D printed frame to create nine virtual cameras, which provides low-cost video frame-rate light fields.

Don't Trip! Automatic Safety Inspection for Trip Hazards on Construction Sites


Construction sites are dangerous places, with injuries and some fatalities each year. Trip hazards are an ever-present danger, and represent an interesting practical and scientific challenge.

It's not feasible for a robot observer to directly manipulate or test every trip hazard in the environment, nor is it safe to watch thousands of humans trip over objects as a training method. The most practical method is to learn from RGB or RGB-D data to detect trip hazards.

The problem is also interesting and challenging because it's not the object type itself that primarily determines how much of a trip hazard something is, but rather its current affordance. A ladder standing up against a wall is far less of a trip hazard than when lying on the floor - so any learning system must be capable of distinguishing affordance from object class. This research is now being expanded into other domains including mining.

We conducted a comprehensive investigation into the performance characteristics of 11 different deep learning colour and depth fusion approaches, including four fusion and one non fusion approach; using both colour and two types of depth images. The resulting best performing systems were able to achieve a high level of competency in automatically detecting trip hazards in environments they had never seen before.

PhD researcher Sean McMahon's paper "Multimodal Trip Hazard Affordance Detection on Construction Sites" co-authored with CI Niko Sünderhauf, CI Michael Milford and Research Affiliate Ben Upcroft was accepted to the IEEE Robotics and Automation Letters journal and also presented at the International Conference on Intelligent Robots and Systems (IROS) 2017.



▲ Recognizing trip hazards on construction sites.

Understanding: Semantic Representations

The Semantic Representations (SR) program develops models, representations and learning algorithms that will allow robots to reason about their visual percepts, describe what they see and plan actions accordingly. The representations explored in this program will enable the recognition of activities from observer and actor viewpoints, fine-grained understanding of human-object, object-object and robot-object interactions and large-scale semantic maps of environments. The program also investigates the mapping of visual inputs to sequence outputs for achieving a given task (e.g., describing a scene textually or providing a symbolic representation for a robotic task based on human demonstration).

SR projects

SR1—VIDEO REPRESENTATIONS

Research leader: Basura Fernando (RF)

Research team: Stephen Gould (CI), Gordon Wyeth (CI), Anoop Cherian (RF), Sareh Shirazi (RF), Fahimeh Rezazadegan (PhD), Rodrigo Santa Cruz (PhD), Bohan Zhuang (PhD)

Project aim: Part of robots' understanding and interacting with their surrounds and humans is being able to operate in a dynamic environment. This project considers ways in which videos or short video segments can be represented usefully for robots, plus ways they can monitor, understand and predict the actions and interactions of a human. It also explores ways robots can use video feed to predict the consequences of their actions. The **Video representations** project extends these elements further to investigate how robots can learn from observation and an understanding of dynamic scenes to work on collaborative or cooperative tasks.

Previous projects SR2 – Representing Human-Object interactions and SR5 – Perceptions & manipulation were merged with SR1 – Video representations.

SR3—SCENE UNDERSTANDING FOR ROBOTIC VISION

Research leader: Niko Sünderhauf (CI)

Research team: Tom Drummond (CI), Michael Milford (CI), Ian Reid (CI), Feras Dayoub (RF), Yasir Latif (RF), Trung Thanh Pham (RF), Ming Cai (PhD), Mehdi Hosseinzadeh (PhD), Kejie Li (PhD), William Hooper (PhD) Lachlan Nicholson (PhD), Saroj Weerasekera (PhD), Huangying Zhan (PhD)

Project aim: For many robotic applications, understanding a scene comes from models of the environment that will give a robot information about geometric and semantic concepts and possible actions it can take ("affordances").

The **Scene understanding for robotic vision** project is algorithms and representations for robots to acquire this contextual information, and make reasoned considerations about the uncertain and dynamic environments they are operating in.

As part of this work, the project has developed maps of the environment that are labelled semantically with a standard set of class labels, that are consistent and accurate within the scene's 3D structure. We are investigating how we can use information from millions of examples to improve scene structures without having to impose hard constraints, like Manhattan world models. In addition to our other objectives, we are also providing semantic representations for the ACRV-SLAM project (AA2).

With these maps and labels established, we will develop representations that allow for a scene to be deconstructed, use the component parts for planning robotic

navigation and acquisition or manipulation of objects. The representation of uncertainty is a key element of this project. While this is well understood in the context of geometry, our research question is how do we determine, represent and use uncertainty resulting from inference over semantic entities. The Scene understanding for robotic vision project draws on advances from the Fundamental Deep Learning project (VL1, finishing in 2017) to bridge this gap. As this project advances, we will move beyond simple labels to develop deeper, more powerful label sets, such as those describing affordances.



▲ Research Assistant Chad Odsey assembles an Ikea table (part of the Centre-produced Ikea dataset), and the rank-pooling action recognition systems correctly identifies that he is screwing in the leg.



▲ A combined point and object based map, computed in real-time by our system. Grey areas show the point cloud generated by ORB-SLAM while coloured areas show where known object classes have been recognised and reconstructed (e.g. cyan=monitor, red=chair).


SR4—JOINT VISION AND LANGUAGE REPRESENTATIONS FOR ROBOTS

Research leader: Anton Van Den Hengel (CI)


Research team: Stephen Gould (CI), Chunhua Shen (CI), Anthony Dick (AI), Basura Fernando (RF), Chao Ma (RF), Qi Wu (RF), Peter Anderson (PhD)

Project aim: Unlike many computer vision tasks, which can be precisely circumscribed (e.g. image segmentation) robots in an open world must be able to learn and respond to unforeseen questions about unforeseen images, or develop action strategies for previously unseen circumstances.


The aim of this project is to build joint vision and language representations to describe scenes (captioning), answer queries (VQA), and describe and define robotic actions. The joint vision and language representations for robots project moves beyond natural language to consider more general image-to-sequence models, where a sequence may be intended as robot control commands.



▲ Three different scenarios where the method can be applied. Top Row to Bottom Row: image-guided inpainting of an ORB-SLAM depth map; a Kinect depth map with structured holes; and a LIDAR depth map respectively. From left to right: the input image; single view depth prediction results; sparse/incomplete depth map obtained geometric sensor; fused dense depth maps.



▲ Dense scene reconstruction using CNN-based priors. The camera video feed (left) is processed by a CNN to produce a surface normal prediction (bottom right) and this is fused with geometric information from the image stream to produce a dense reconstruction (top right).



▲ The Adelaide/ANU winning entry in the official VQA challenge. Top: the task is to correctly answer free-form questions about images; Bottom: the architecture used for the task.

Key Results in 2017

A novel rank pooling-based action recognition method was developed and evaluated by researchers at our ANU node in collaboration with QUT researchers. This has led to a real-time implementation that was deployed on the IKEA furniture dataset created by the team.

The system also delivers state-of-the-art performance on a number of public benchmarks and a number of papers were presented at the Conference on Computer Vision and Pattern Recognition (CVPR) 2017 and the International Conference on Computer Vision (ICCV) 2017. For example “Generalised Rank Pooling for Activity Recognition” by RF Anoop Cherian, RF Basura Fernando, Data61-CSIRO and ANU Senior Researcher Mehrtash Harandi and CI Stephen Gould was presented at CVPR.

Adelaide and Monash researchers delivered a new method for filling in depth based on uncertainty modelling in CNNs. This runs in real-time and can fill in parts of the scene where a dense or sparse depth sensor has not delivered a measurement. For example, in LIDAR it can be used for smart interpolation. For Kinect data it can fill in where holes arise from non-reflective objects or stereo shadows. The real-time implementation was used in the early prototype of the Amazon Challenge perception pipeline. The paper “Just-In-Time Reconstruction: Inpainting Sparse Maps using Single View Depth Predictors as Priors” by PhD Researcher Saroj Weerasekera, PhD Researcher Thanuja Dharmasiri,


Associated Research Fellow Ravi Garg, CI Tom Drummond and CI Ian Reid has been accepted for the 2018 International Conference on Robotics and Automation (ICRA) to be held in Brisbane, Queensland.

An object-based semantic RGB-D SLAM was developed in a QUT and Adelaide collaboration. This real-time system maps the world (using ORB-SLAM) but recognises objects and reconstructs the objects in their own coordinate frame. The result is a map that comprises points and objects, with each object having a hi-fidelity geometric reconstruction. This work builds on other work from the team on unsupervised scene segmentation, also used in the Amazon Challenge perception pipeline. The paper “Meaningful maps with object-oriented semantic mapping” by CI Niko Sünderhauf, RF Trung Pham, RF Yasir Latif, CI Michael Milford and CI Ian Reid was presented at the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) in Vancouver, Canada.

We continued development of integration of deep learning with SLAM, delivering the world’s first system that fuses dense geometry from a DTAM-like system via a prior provided from a deep learning system. More specifically, a deep network predicts the scene surface normals and these are used to regularise the geometry computed using more standard photometric consistency. A real-time system has been developed and demonstrated, and a licence is under negotiation with an Oxford start-up, 6Degrees-xyz. The paper “Dense monocular

reconstruction using surface normal” by PhD Researcher Saroj Weerasekera, RF Yasir Latif, Associated RF Ravi Garg and CI Ian Reid was presented at the 2017 IEEE International Conference on Robotics and Automation (ICRA) in Singapore.

A system jointly developed between Adelaide and ANU was ranked 1st in the VQA Challenge 2017.



Visual Question Answering is an important stepping stone on the route to natural interaction with robots and this work is therefore germane to our efforts to integrate Vision and Language, for tasks such as Human-Robot interaction and for robotic navigation in a semantic space.

The paper “Tips and Tricks for Visual Question Answering: Learnings from the 2017 Challenge” by Adelaide PhD Researcher Damien Teney, PhD Researcher Peter Anderson, Xiaodong He and CI Anton van den Hengel was presented at the Computer Vision and Pattern Recognition (CVPR) conference workshop corresponding to the challenge in Hawaii, USA.

What is Deep Learning and how does it work?

Imagine a world where Facebook automatically finds and tags friends in your photos; where Skype translates spoken conversations in real time; or where Google Deepmind's AlphaGo computer program masters the ancient game of Go. It may not be as far off as you think.

Welcome to the world of machine learning and deep-neural networks, or more simply, 'deep learning'.

Deep learning is a subset of machine learning that examines computer algorithms to learn and improve independently. One such technique is called 'neural networks'. Inspired by the nerve cells (neurons) that make up the human brain, neural networks comprise layers (also termed 'neurons') that are connected adjacent to one another; the more layers, the 'deeper' the network. The algorithms in this technique are the foundation for deep learning and play a central part in image recognition—and robotic vision.

Riddled with seemingly insurmountable problems, Neural networks have been largely ignored by machine learning researchers in the past, until a confluence of events thrust neural networks back into the forefront of the field.

"What gave neural networks the biggest leg-up," explains CI Ian Reid, "was the advent of a mammoth amount of labelled data."

In 2007, a pair of computer scientists at Stanford University and Princeton University launched ImageNet, a database of millions of labelled images from the Internet. Now ImageNet provides neural networks with about 10 million images and 1000 different labels. This helped form the foundation for neural networks to become a central tool of robot vision, and gains have been made since.

"A big step has been the emergence of convolutional neural networks," Ian said.

"As with traditional neural networks, convolutional counterparts are made of layers of weighted neurons."

These neurons are not solely modelled on the workings of the brain but also from the visual system itself.

As networks get deeper and researchers unwrap the secrets of the human brain on which they are modelled, the networks will become ever more nuanced and sophisticated.

"As we learn more about the algorithms coded in the human brain and the tricks evolution has given us to help us understand images," says Director Peter Corke, "we will be reverse-engineering the brain and borrowing these tricks."

Acting: Vision & Action

The Vision & Action program is about robots manipulating real objects by pushing the boundaries of speed, coordination and complexity.

VA projects

The Vision & Action program is about robots manipulating real objects by pushing the boundaries of speed, coordination and complexity.


VA1—LEARNING ROBUST HAND-EYE COORDINATION FOR GRASPING IN NOVEL ENVIRONMENTS

Research leader: Juxi Leitner (CI)

Research team: Peter Corke (CI), Rob Mahony (CI), Chunhua Shen (CI), François Chaumette (PI), Chris Lehnert (Research Affiliate), Norton Kelly-Boxall (PhD), Sean McMahon (PhD), Douglas Morrison (PhD), Adam Tow (PhD), Fangyi Zhang (PhD), Zheyu Zhang (PhD)

Project aim: Replicating hand-eye coordination in complex visual environments involves developing robust motion control of robotic platforms based on vision data, being able to deal with the variations and complexities the robot will encounter as it operates in a real-world environment.

This project aims to go beyond engineered visual features and environments to develop demonstrator systems that allow manipulation of real world objects including cups, pens, tools, vegetables, toys and many other examples. A key aspect of this project is developing systems and architectures that are robust enough to deal with a variety of conditions and can adapt easily to new tasks, new objects and new environments.



▲ Centre PhD Researcher Fangyi Zhang pictured with a Baxter robot during his internship at the University of Maryland. Fangyi's research focuses on learning visuo-motor policies through reinforcement learning and transferring policies from simulation to the real world.

VA2—HIGH PERFORMANCE VISUAL CONTROL AND NAVIGATION IN COMPLEX CLUTTERED ENVIRONMENTS

Research leader: Rob Mahony (CI)

Research team: Peter Corke (CI), Jonathan Roberts (CI), Jochen Trunpf (AI), Alex Martin (RE), Juan Adarve (PhD), Peter Kujala (PhD), Sean O'Brien (PhD), Jean-Luc Stevens (PhD), Dorian Tsai (PhD)

Project aim: The project is exploring real-world control scenarios in which there are multiple ways to complete a set task, including motions that interact with the environment, for instance moving a glass out of the way to reach a bottle, or picking up a sequence of items from a cluttered workspace in a learned order, rather than a pre-programmed order. Achieving this successfully involves two capabilities, which this project is addressing:

- / moving quickly, while maintaining stability, through a cluttered environment by calling on integrated control and sensing systems, in this case, developing algorithms to visually detect obstacles quickly and independently formulate strategies to avoid them
- / coming up with, and executing an effective and efficient solution to a complex task by calling on integrated decision and planning capabilities, in this case, a vision-based solution that exploits semantic information and other relevant cues.

The project is exploring real-world control scenarios in which there are multiple ways to complete a set task, including motions that interact with the environment.



▲ Organisers of the “New Frontiers for Deep Learning in Robotics” workshop held at the 2017 Robotics: Science and Systems (RSS) conference. Aaron Courville (University of Montreal), Michael Milford (Centre CI, QUT), Niko Sünderhauf (Centre CI, QUT), Chelsea Finn (UC Berkeley), Piotr Mirowski (DeepMind), David Cox (Harvard), Yarin Gal (University of Cambridge), Juxi Leitner (Centre RF, QUT)

Key Results in 2017

Research Fellow Juxi Leitner, CI Niko Sünderhauf, CI Michael Milford and Centre Director Peter Corke collaborated with Associate Professor Pieter Abbeel from UC Berkeley to present the workshop “New Frontiers for Deep Learning in Robotics” at the 2017 Robotics: Science and Systems (RSS) conference held in Boston, USA. The workshop was sponsored by Google and OSARO and attracted over 200 people.

The workshop “Deep Learning for Robotic Vision” was held at the 2017 Conference on Computer Vision and Pattern Recognition (CVPR) in Hawaii involving researchers from our Adelaide and QUT nodes. The workshop was led by CI Gustavo Carneiro and involved CI Niko Sünderhauf, RF Juxi Leitner, RF Vijay Kumar, RF Trung Pham, CI Michael Milford, CI Ian Reid and Centre Director Peter Corke. It was an international collaboration with Google Brain, Google Research, the University of Oxford and the University of Freiburg. The workshop attracted around 200 attendees. The group are organising a special issue on “Deep Learning for Robotic Vision” for the International Journal of Computer Vision (IJCV) in 2018 following the success of this workshop.

By learning to predict the outcome of human and robot actions on an environment we enable a robot to physically perform a human demonstrated task without knowledge of the thought processes or actions of the human, only their visually observable state transitions. The approach was shown to work in proof-of-concept table-top, object manipulation tasks and demonstrate generalisation to previously unseen states, while reducing the priors required to implement a robot task learning system. This is one of the biggest short falls of existing approaches of Learning from Demonstration, Reinforcement Learning and Inverse Reinforcement Learning, the need to have a lot of trials or very strong priors. Adam is currently on leave and working for a Chinese start-up in the warehouse automation space.

Building on a visit in 2016, QUT and international partner INRIA have continued their joint research in combining novel deep learning techniques with classical visual servo control approaches. Quentin Bateau, a PhD researcher supervised by Eric Marchand and François Chaumette in Rennes was leading these efforts on the French side and RF Juxi Leitner from Brisbane. After attending the Robotic Vision Summer School, Quentin also visited QUT for a few weeks to implement, train and test various approaches.

Researchers started by exploring possibilities to combine convolutional neural networks with image based visual servoing, specifically

photometric approaches developed at INRIA. The naïve integration yielded promising results and over the course of the year better results were achieved by tweaking the network architecture and the data training method. A paper, presented at the Deep Learning workshop at RSS 2017, highlighted that this can be trained quickly if one finds a smart way of training. In our case the idea was that accuracy is more important when you are close to s^* , your desired feature state. Taking this into account during data augmentation yielded significant benefits.

Adam Tow's PhD research at the intersection of learning, vision and manipulation led to some interesting ideas on how to train a robot's actions. In the paper “What Would You Do? Acting by Learning to Predict” a novel way to learn tasks directly from visual demonstrations was proposed.

The efforts to extend this work, which required scene specific training, to scene-agnostic visual servoing networks has been submitted (and in the meantime accepted) to be presented at ICRA 2018. Quentin has finalised his thesis and is now involved in a French start-up company.

Internships and Deep Q networks

Fangyi Zhang is a third year PhD researcher in the centre. His research has been focusing on learning visuo-motor policies through reinforcement learning and transferring policies from simulation to the real world, which are two sub-projects of the VA program.

Fangyi's PhD research was started by evaluating the feasibility of learning vision-based planar reaching using a Deep Q Network (DQN), showing that DQN is able to learn reaching in simulation, while the learned policies do not transfer directly to real robots with real cameras observing real scenes. He then proposed modular deep Q networks to transfer policies from simulation to the real world in a low-cost manner with a small number of labelled real images. To further weaken the reliance on labelling real data which is expensive or even impractical in many robotic applications, Fangyi proposed an adversarial discriminative approach to transfer visuo-motor policies from simulated to real environments, which reduced the required amount of labelled real data by 50% for object reaching in clutter with a seven DoF robotic arm (Baxter). These results have been presented at robotics conferences such as ACRA and CVPR, among which the ACRA 2017 paper on modular deep Q networks appeared as a best paper finalist. Currently, Fangyi is investigating how to extend the adversarial discriminative transfer and modular approaches to more complicated robotic manipulation tasks and also further reduce the demand for labelling real data.

Apart from his PhD research, Fangyi has also been actively taking part in cross-node collaborations. In October 2017, he stayed in the ANU node for 10 days to collaborate with fellow PhD researcher Zheyu Zhuang and CI Robert Mahony for a project of using Lyapunov functions as bridges to link the state-of-the-art deep learning techniques and current robotic control frameworks for better robustness (Zheyu's PhD research). In addition, he also stayed in the Autonomy Robotics Cognition (ARC) lab at the University of Maryland in the USA for 3 months in 2016, participating in a project enabling a Baxter robot to do housework in a kitchen scenario.



On his way back to Australia, he visited the Berkeley Artificial Intelligence Research (BAIR) lab at UC Berkeley, presenting his work and also exchanging ideas on deep learning for robotics with top scholars in the field. All these experiences made his PhD a fruitful journey and also an indispensable part of our centre's culture.

Centre Team Wins Amazon Robotics Challenge with Low-Cost Robot

Last year, the Centre's team took out the top prize in the 2017 Amazon Robotics Challenge in Nagoya, Japan.



Competing as one of **16 teams** from **10 countries**, participants were tasked with building their own automated robot, including hardware and software, to successfully pick and stow items in a warehouse.

The Challenge combined object recognition, pose recognition, grasp planning, compliant manipulation, motion planning, task planning, task execution, error detection and error recovery.

The robots were scored by how many items they successfully picked and stowed in a fixed amount of time.

For the challenge, the Australian Centre for Robotic Vision developed **"Cartman", a Cartesian robot** that can move along three axes at right angles to each other, like a gantry crane, and featured a rotating gripper that allowed the robot to pick up items using either suction or a simple two-finger grip.

"Cartman provides us the right tools and is specialised for the task at hand, picking out of boxes", explains **Project Leader Juxi Leitner**.

"We are world leaders in robotic vision and we are pushing the boundaries of computer vision and machine learning to complete these tasks in an unstructured environment."


Cartman's vision system was the result of hours of training data time, according to Dr Anton Milan from The University of Adelaide. "We had to create a robust vision system to cope with objects that we only got to see during the competition."





One feature of our vision system was that it worked off a very small amount of hand annotated training data and we needed just seven images of each unseen item for us to be able to detect them, aided by the use of weighing scales.”

- ## Members of the winning Team
- ▶ Ming Cai, University of Adelaide
 - ▶ Peter Corke, QUT
 - ▶ Jordan Erskine, QUT*
 - ▶ Anthony Gillespie, QUT*
 - ▶ Riccardo Grinover, QUT*
 - ▶ Alec Gurman, QUT*
 - ▶ Tom Hunn, QUT*
 - ▶ Norton Kelly-Boxall, QUT
 - ▶ Vijay Kumar, University of Adelaide
 - ▶ Juxi Leitner, QUT
 - ▶ Darryl Lee, QUT*
 - ▶ Chris Lehnert, QUT
 - ▶ Steve Martin, QUT
 - ▶ Matt McTaggart, QUT
 - ▶ Anton Milan, University of Adelaide
 - ▶ Doug Morrison, QUT
 - ▶ Nathan Perkins, QUT*
 - ▶ Trung Pham, University of Adelaide
 - ▶ Gerard Rallos, QUT*
 - ▶ Andrew Razjigaev, QUT*
 - ▶ Ian Reid, University of Adelaide
 - ▶ Thomas Rowntree, University of Adelaide
 - ▶ Rohan Smith, QUT*
 - ▶ Adam Tow, QUT
 - ▶ Sean Wade-McCue, QUT
 - ▶ Saroj Weerasekera, University of Adelaide
 - ▶ Zheyu Zhuang, ANU
- *undergraduate students*



Special thanks to the Centre's team sponsors, Osaro (www.osaro.com) and Photoneo (www.photoneo.com).



Learning: Visual Learning

This program explores the enormous potential that still exists towards solving previously impossible problems in machine perception. The recent breakthroughs from the machine learning community have allowed researchers to address new visual learning problems, as well as solve old problems.

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

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

VL projects

In 2017, the *Fundamental deep learning* project (VL1) wound up, leaving the *Learning for robotic vision* project (VL2) as the sole project in this program.

VL1—FUNDAMENTAL DEEP LEARNING

Research leader: Vijay Kumar (RF)

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

Project aim: Fundamental deep learning is at the core of successfully creating robots that see and is therefore at the core of our Centre’s purpose. Therefore it is imperative that we are actively at the forefront of current machine learning techniques. This includes exploring developing and exploiting, novel network architectures; developing detection and instance- or pixel-level annotations for thousands of open sets of classes; developing weakly supervised, online-trained, zero-shot or unsupervised learning models; developing active learning with and from temporal data.

VL2—LEARNING FOR ROBOTIC VISION


Research leader: Chunhua Shen (CI)

Research team: Ian Reid (CI), Gustavo Carneiro (CI), Chao Ma (RF), Vijay Kumar (RF), Hui Li (PhD), Ben Meyer (PhD), Tong Shen (PhD), Bohan Zhuang (PhD), Yuchao Jiang (Masters by Research)

Project aim: This project is investigating learning that is specific to robotic vision tasks, particularly where there are resource constraints, such as in an embedded vision system. An example of this type of system is COTSbot, which has both storage and power constraints in its operation. To address this, we are exploring avenues such as video segmentation; “any time” algorithms; fast, approximate, asymmetrically computed inference; unsupervised learning; online and lifelong learning for robotic vision; and suitable deep learning techniques.



▲ Example results showing object parsing (left) and semantic segmentation (right)



▲ This proposed deep metric learning model is capable of quickly searching an entire training set to select effective training samples.

Key Results in 2017

A unified framework was proposed to tackle the problem of robots reading text in the wild.

In contrast to existing approaches that consider text detection and recognition as two distinct tasks and tackle them one by one, the proposed framework settles these two tasks concurrently. The whole framework can be trained end-to-end, requiring only images, ground-truth bounding boxes and text labels. Published in International Conference on Computer Vision (ICCV2017).

We presented RefineNet, a generic multi-path refinement network that explicitly exploits all the information available along the down-sampling process in a fully convolutional network, enabling high-resolution prediction using long-range residual connections. Published in IEEE Conf. Computer Vision and Pattern Recognition (CVPR2017)

RefineNet was the perception system developed by Adelaide researchers and used in the 2017 Amazon Robotics Challenge.



Joint organisation of the workshop “Deep Learning for Robotic Vision” that happened in conjunction with Computer Vision and Pattern Recognition (CVPR 2017) - a fantastic team effort involving lots of people from Adelaide and QUT.

Large-scale datasets have driven the rapid development of deep neural networks for visual recognition.

However, annotating a massive dataset is expensive and time-consuming. Web images and their labels are, in comparison, much easier to obtain, but direct training on such automatically harvested images can lead to unsatisfactory performance, because the noisy labels of Web images adversely affect the learned recognition models. To address this drawback we propose an end-to-end weakly-supervised deep learning framework which is robust to the label noise in Web images. Published in IEEE Conf. Computer Vision and Pattern Recognition (CVPR2017)

The Bayesian data augmentation paper that explains the theoretical aspects of the use of adversarial training for data augmentation: **Toan Minh Tran, Trung Thanh Pham, Gustavo Carneiro, Lyle Palmer, Ian Reid**. A Bayesian Data Augmentation Approach for Learning Deep Models. Neural Information Processing Systems (NIPS) 2017.

Joint publication between Adelaide and Monash on a new method for deep metric learning that is potentially useful for deep learning optimisation methods used by all nodes of the centre: **Ben Harwood, Vijay Kumar, Gustavo Carneiro, Ian Reid, Tom Drummond**. Smart Mining for Deep Metric Learning. International Conference on Computer Vision (ICCV 2017).

Technology: Algorithms & Architecture

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

AA1—VOS-DISTRIBUTED ROBOTIC VISION

Research leader: Tom Drummond (CI)

Research team: Peter Corke (CI), Vincent Lui (RF), Steven Martin (RE), Will Chamberlain (PhD)

Project aim: This project is creating a Vision Operating System that provides a framework to bring multiple sensing and computational resources together to solve complex robotic vision problems. The system will enable robots to use external sensing resources, for instance CCTV cameras in the environment, or sensors mounted on other robots, as well as computation resource sensors, or large computing resource within the network. This kind of framework enables novel solutions to complex problems, where various resources can be combined in different ways to solve complex localisation, navigation, understanding and planning problems.

AA2—ACRV SLAM FRAMEWORK

Research leader: Viorela Ila (CI)

Research team: Matt Dunbabin (CI), Tom Drummond (CI), Richard Hartley (CI), Hongdong Li (CI), Ian Reid (CI), Laurent Kneip (AI), Feras Dayoub (RF), Yasir Latif (RF), Vincent Lui (RF), Mina Henein (PhD), Andrew Spek (PhD), Sean McMahon (PhD), Jun Zhang (PhD)

Project aim: This project is developing novel simultaneous localisation and mapping (SLAM) algorithms that can perform in challenging large-scale, dynamic, dense and non-rigid environments. In particular, it focuses on developing and integrating robot vision algorithms from robust vision, real-time vision and semantic vision areas, into a single SLAM-centred robot navigation framework. The framework will be demonstrated in real-world robot applications including autonomous underwater vehicles (AUVs), unmanned aerial vehicles (UAVs) and ground-based autonomous vehicles.


AA3—COMPUTER GRAPHICS SIMULATION FOR ROBOTIC VISION

Research leader: Peter Corke (CI)

Research team: Niko Sünderhauf (CI), Trung Thanh Pham (RF), Steven Martin (Research Engineer), John Skinner (PhD)


Project aim: A robotic vision system's performance depends on a number of highly variable factors, including the robot's initial state, the world it perceives, lighting conditions, unforeseen distractors (like moving objects) and unrepeatable sensor noise. A consequence of these factors is that we cannot faithfully repeat a robotic vision experiment, nor can we rigorously and quantitatively monitor different algorithms' performance. A critical bottleneck for machine learning applications is the limited amount of real-world image data that can be captured and labelled for both training and testing purposes.

To address these challenges the Computer graphics simulation for robotic vision project is investigating the potential of photo-realistic graphic simulation, based on state-of-the-art game-engine technology.



◀ Final sparse 3D reconstruction using multi-camera underwater image acquisition with color-coded magnitude of the uncertainty of the estimation (violet-high uncertainty, red-low uncertainty)

▼ Two orders of magnitude faster covariance recovery in bundle adjustment algorithm.



Key Results in 2017

ANU researchers developed a new method for globally-optimal image-based camera re-localisation.

This method can be applied to solving the vehicle localisation problem for autonomous driving, and is also useful for robot localisation, camera tracking, as well as spatial computation in virtual and augmented reality (VR/AR) applications. **Dylan Campbell** (ANU PhD Researcher), **Lars Petersson**, **Laurent Kneip** and **Cl Hongdong Li** received the MARR PRIZE Honorable Mention award for their paper “Globally-Optimal Inlier Set Maximisation for Simultaneous Camera Pose and Feature Correspondence” at the 2017 IEEE International Conference on Computer Vision (ICCV) held in Venice, Italy. The Marr Prize is considered one of the top honours for a computer vision researcher.

Additionally, the team achieved first place and a “Best Algorithm Award” in the Non-Rigid Structure from Motion (NRSFM) Challenge held at the Conference on Computer Vision and Pattern Recognition (CVPR).

They also developed a successful computer vision algorithm for monocular non-rigid dynamic scene 3D reconstruction and presented the paper “Monocular Dense 3D Reconstruction of a Complex Dynamic Scene from Two Perspective Frames” at the IEEE International Conference on Computer Vision (ICCV)



Klement Istenic, a visiting PhD researcher from the University of Girona, won first prize in the Student Poster Competition at the OCEANS conference in Aberdeen, UK. The paper “Mission-time 3D Reconstruction with Quality Estimation” was a collaboration between RF Viorela Ila at ANU and colleagues at the University of Girona.

The paper “Fast Incremental Bundle Adjustment with Covariance Recovery” by **RF Viorela Ila**, **Lukas Polok** (Apple Inc), **Marek Solony** (Brno University of Technology) and **Klemen Istenic** (University of Girona) was presented at the International Conference on 3D Vision (3DV), Qingdao, China, 2017. The paper won the Best Paper Honorable Mention at the conference.

The article “SLAM++. A Highly Efficient and Temporally Scalable Incremental SLAM Framework” by **RF Viorela Ila**, **Lukas Polok** (Apple Inc), **Marek Solony** (Brno University of Technology) and **Pavel Svoboda** (Brno University of Technology) was accepted into the International Journal of Robotic Research (IJRR).

Robust Visual Odometry in Underwater Environments

Autonomous underwater vehicles (AUVs) are increasingly being used for underwater surveys and to investigate archaeological finds.

The accurate estimation of pose and velocity of an AUV is critical to ensure the repeatability and validity of scientific data that is captured using sensors on board the AUV. A low-cost and effective way is by using stereo camera sensors to perform visual odometry (VO). However, this is a difficult problem underwater due to poor imaging condition and inconsistent motion caused by water flow. This project created a robust and effective stereo underwater VO system that can overcome these difficulties and accurately localise the AUV. Experimental results demonstrate that the proposed pipeline outperforms existing VO systems in underwater environments, as well as obtaining a comparative performance on the KITTI benchmark dataset.




▲ Fig. 1 Selected typical sample images of the Underwater Coral Dataset.

As is shown in Fig.1, the imaging conditions in water are poor due to light attenuation, poor/artificial illumination, haze and scattering. When the AUV operates in shallow waters, scattering of the sun light is highly problematic. Motion blur can also be present and is due to motion of the robot while the camera shutter is open. The vehicle/camera's motion is inconsistent with oscillation, especially in shallow underwater areas, due to the water waves. All these problems greatly increase the difficulty in estimating the robot's location.

We extensively tested and evaluated possible solutions and proposed a stereo underwater VO system that is able to robustly and accurately localise the AUV. Our system is built upon a popular visual-based localisation system in robotics called ORB-SLAM2. We carefully modified this system to accommodate challenging conditions and demonstrated that the proposed pipeline outperforms existing VO systems in underwater environments (see Fig.2).

Participants:

- | | |
|---|--|
| / RF Viorela Ila
(project leader) | / RF Feras Dayoub |
| / PhD researcher Jun Zhang | / CI Matthew Dunbabin |
| / Associate Investigator
Laurent Kneip | / Klemen Istenic (visiting PhD
researcher from the University
of Girona) |



▲ Fig. 2: (a) Camera trajectory (red) and the 3D structure (green) produced by our proposed method. (b) Comparison of trajectories generated by our proposed method (red), LIBVIS02 (yellow) and ORB-SLAM2 (cyan). ORB-SLAM2 can only survive the first 60 frames shown in the orange rectangle. Underwater Coral Dataset.

Pepper: a peek into the future of social robotics

Pepper is the world's first humanoid robot that can recognise emotions. Created by SoftBank Robotics, it also mimics human behaviours such as following the conversation around it by looking at whoever is speaking. We are using Pepper as a social robotics research platform to explore the potential for humanoid robots.

"The degree to which social robots could change society is the basis of our new research program," said project leader Belinda Ward.

"Pepper is probably the most 'personable' robot on the market in terms of its perceived emotional intelligence, making it a fantastic platform to investigate the suitability of social robots, which is still a very new field.

The \$1.5 million social robotics research program developed by Centre COO Sue Keay, is funded by the Advance Queensland initiative and will involve the Centre's collaboration with SoftBank Corp's subsidiary, ST Solutions Australia (STSA).

Belinda said the funding would allow her team to explore the different applications of these robots across a range of settings and conditions and their effectiveness in each.

"While a social robot in every home is probably a long way off for society, I see a place for social robots supporting human staff in every hospital, aged care facility and classroom, as companions and helpers, with many more settings to be considered in the future.

"Would a shy child in a classroom be more willing to ask a question of the robot rather than the class? Would a nursing home resident enjoy reminiscing about their past to a robot? Would a hospital patient feel comfortable chatting with a robot as it tidied their room?"

Belinda's team will also work with STSA to improve Pepper's vision and navigation systems, making the robot even more versatile in the future.

"What we learn from human-computer interactions with Pepper will inform the next generation of service robots, building an effective social component into their task-oriented programming."

Humanoid Robotics project leader Belinda Ward with Softbank's Pepper Robot. *Image courtesy of QUT Media.* ▶



Revolutionising the health sector

The Australian Centre for Robotic Vision is leading the way in robotic research globally with a number of new applications and technologies that will revolutionise the health sector.


From the use of robotic systems in magnetic resonance imaging (MRIs) and mammograms, detecting melanomas and cancer, to the introduction of a futuristic scan that detects morbidity, our Centre is undertaking ground-breaking research in the medical sphere.

Gustavo Carneiro, Chief Investigator at the Centre's University of Adelaide node, specialises in robotic machine learning. In recent years, Gustavo has been focused on developing a deep learning algorithm with more than 500 images to help detect and classify masses in mammograms.

"Breast cancer is one of the major diseases affecting the lives of women across the globe. The analysis of breast masses from mammograms represents an important task in diagnoses, which is predominantly a manual process at present and subject to the assessment of a clinical expert. We are pleased to be undertaking world-leading research into how robotics can help improve the accuracy, accessibility and cost of this process," said Gustavo.

Gustavo and his team have also recently published a deep learning algorithm to detect morbidity. The system looks at elderly patients' chest CT scans that don't reveal anything outwardly unhealthy. Using the algorithm, the scan can detect (with almost 70% accuracy) if the patient will survive in the next five years.

"By 2050, a quarter of Australia's population will be over 65, with significant economic costs and productivity challenges. With reliable, low-cost robotics, we can dramatically reduce the costs of healthcare, while increasing its efficiency and accessibility."



Chief Investigator Gustavo Carneiro

Research Performance (KPIs)

ARC Centre of Excellence Key Performance Indicators (KPIs)	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017
Conference publications	25	42	50	53	50	66	50	47
Journal publications	12	15	25	22	25	32	25	24
Paper prizes or awards	1	6	3	1	5	4	5	7
Disclosures/ Patents	0	0	0	0	0	1	2	5
Quality of research outputs (% of papers published in peer-reviewed outlets)	100	100	100	100	100	100	100	100
Percentage of publications relevant to SRPs	80	80	80	100	80	100	100	100
Number of invited talks/papers/ keynote lectures given at major international meetings (incl. international conferences held in Australia)	5	6	10	41	10	38	10	34
Number and nature of commentaries about the Centre's achievements (in professional magazines and the popular press, e.g newspapers, TV, etc.)	5	49	10	94	20	84	20	181

SECTION 3 | Research Impact



Research Impact

At the Australian Centre for Robotic Vision, we are conducting world-leading research that will help bring transformational change to the world. Specifically, we are focusing our work towards making an impact on three critical challenges facing the world today, including:

Our environment



Our health care capabilities



Our resources and their management, focusing on food production



Where we live

Our environment is a complex network of both natural and built environments and it is critical to our wellbeing.

Our natural environment sustains us physically with fundamental necessities such as drinkable water, breathable air and nutrient-rich food. But it also provides beauty and wonder that nourishes the human spirit. Managing this environment in a sustainable way is vital, assuring our own survival and quality of life well into the future. However, our world is vast and as humans, we can only ever inspect a tiny fraction of it at a time. We rely heavily on extrapolating this data in an attempt to understand the full condition of our natural environment, but this means local variations can be overlooked, in both space and time.

Robotic vision has the potential to increase the areas covered and the frequency of inspection, while at the same time reducing the cost. Ultimately, robots with intelligent vision capabilities would be equipped to not only inspect but also remediate our natural environment.

Just as important is our built environment, our cities, roads, dams, tunnels, pipelines and electrical distribution networks. As with the natural environment, vision-capable robots offer the potential to reduce the costs of inspection as well as reduce risks to humans who currently perform this often hazardous work.

How we live

A key concern for any society is the health of its people. In Australia, we are faced with the challenge of an ageing population. Current estimates suggest that by 2050, a quarter of us will be over 65 years of age, and ageing has significant economic costs.

Firstly, the dependency ratio increases, meaning fewer people in the workforce to carry the burdens of production and caring for those who are no longer working.

Secondly, the cost of healthcare increases. "Healthy, independent living" describes how people can continue to live at home, even if they need some assistance. The key to independent living is access to affordable assistive technology including devices or systems that provide practical solutions to everyday activities. Robotics is already having an impact in this area, developing technologies such as exoskeletons, care assistants, communications devices and devices to improve mobility. The effects of these and other robotic healthcare solutions are that they reduce costs, allowing more people to remain independent; to care for themselves, and to preserve their dignity and their freedom.

Introducing various automation systems to hospitals will help reduce healthcare costs by bringing all the advantages that automation gives to other sectors, such as increasing productivity, ensuring higher quality work, removing menial duties from highly trained nursing staff and removing workers from harmful or high-risk tasks. For hospitals, automation extends beyond surgery to areas such as internal logistics, patient handling, rehabilitation and care.



Harvey the Capsicum Picking Robot

Means to live

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

Food production is reliant on human labour, the least labour intensive industries are extensive crop agriculture and dairy production, which are highly mechanised and are becoming more automated.

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

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

SECTION 4 | National Benefit



National Benefit—meeting Australia’s innovation, science and research priorities

Robots are already in wide use on factory floors in Australia’s food manufacturing industries and in the field as autonomous mining and port container vehicles. However, robots are noticeably absent from a number of other sectors.

So where are they all?

Without vision capabilities, there are a wide array of potential applications that are closed to robots. Complex manual assembly, packing, navigation, machine operation, fruit picking, crop spraying, remote assistance, smart homes and appliances, autonomous driving, together with environmental surveying and monitoring.

At the Australian Centre for Robotic Vision, we are developing vision-enabled robotic systems that can understand and respond to their environment, operate reliably over long periods in complex, unstructured surroundings and interact safely and effectively with people as well as objects.

By creating robots that see, we will be able to help solve Australia’s future environmental, economic and social challenges. These challenges include climate change, growing populations with increasingly ageing populations who will need healthcare. Some of the issues identified through the nation’s Science and Research Priorities that our Centre’s research can address include:

- / labour shortages and low productivity growth in key industries
- / rising Occupational Health & Safety (OH&S) compliance costs
- / the need for increased productivity due to an ageing population
- / rising healthcare costs and demands for customised healthcare solutions
- / ageing infrastructure
- / growing demands for minerals, energy and food
- / the need to preserve the environment.

Current trends suggest robots will be fundamental to improving productivity in our resources, services and manufacturing sectors, supporting the growth of new industries while fostering an entrepreneurial and innovative knowledge economy that will benefit Australia.

We are specifically addressing these key priorities to create national benefit by:

- / undertaking world-class research, grounded by national challenges with a focus on 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 by transferring trained people into important future industries and new enterprises
- / building businesses’ awareness of the technologies we are creating, through planned communication and a variety of engagement models
- / actively collaborating with others in the national innovation system, including non-partner universities, organisations and industry, through collaborative and contract research
- / fostering strong international engagement with top international researchers in our field, relationships that can in turn be leveraged by Australian industry.

The Australian Government Science and Research Priorities

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



Food

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

Soil and water

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

Transport

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



Cyber security

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

Energy

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



Resources

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

Advanced manufacturing

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



Environmental change

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

Health

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



National Benefit (KPIs)

ARC Centre of Excellence Key Performance Indicators (KPIs)	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017
Contribution to Australia's Strategic Research Priorities The Centre will contribute to: <ul style="list-style-type: none"> / Living in a changing environment / Promoting population health and wellbeing / Managing our food and water assets / Securing Australia's place in a changing world / Lifting productivity and economic growth Percentage of publications relevant to SRPs	Annually	80%	80%	80%	100%	100%	100%	100%	100%
Measure of expansion of Australia's capability in the priority area(s)	At review		24 postdoctoral fellows trained, 60 PhD students trained, 390 honours students involved, 10 additional staff attracted to partner universities.						
Measure of international reputation and competitiveness	At review		The number of publications in high impact archival journals and at the best scientific conferences is very impressive. In particular in the robotics field there is a clear recognition of the brand/identity provided by the Centre." Independent Scientific Review.						
/ Number of awards		1	13	2	17	3	10	3	20
/ Number of fellowships		2	3	3	4*	3	2	3	2
/ International grants		0	1	1	0	2	1	2	6

Weren't we promised The Jetsons?

Remember The Jetsons—the animated sitcom that reflected popular 1960s imaginings of technology as having all the answers? Life, it envisaged, would continue much the way it had, but with our every whim served by platoons of sentient robots.

So, where are all our subservient robots?

“Robots in factories are 60-year-old technology—highly effective, yet simple,” Peter Corke, the Centre’s Director, explains. “They have no idea what they are doing and they don’t work outside the factory.

“Say you take them to an apple orchard—the robot doesn’t know where it is; it doesn’t even know what an apple is or where to find one.”

Essentially, for awareness of environment, the robots need to ‘see’.

“The contextual understanding of the environment is impossible without vision,” says Ian Reid, our Deputy Director.

“(It’s) the missing link that will let robots operate in an unstructured environment.”

This is where the Australian Centre for Robotic Vision comes to the rescue. Using cameras to guide robots to carry out tasks requires a mind-bogglingly complex set of algorithms and advanced 3D geometry. When we take into account a dynamic environment, where random objects can appear on the scene and move at speed across it, the challenges rise exponentially.

The biggest problem is image recognition, where a robot can understand what it “sees” and find an answer to a question like “What is an apple?” But

this can quickly spin off into something not-so-simple at all. What happens when the robot is confronted, not with an apple, but a pear?

Recent advances in the field of machine learning over the past four to five years “has a huge role to play in robotic vision,” says Ian.

“Computer vision has become much better at seeing the world through a camera and converting the images it sees into geometric information.”

The breakthrough, he explains, came in about 2012 with “deep neural networking”, algorithms that learn and improve on their own.

“In this way, robots will help each other learn about their worlds and how they can be interpreted,” Peter says.

“Only one robot needs to get the knowledge, which it can then share,” he explains. “In some ways they will be like the Borg (the cybernetic aliens from science-fiction series, Star Trek). If the knowledge that one acquires is imperfect, then it can be shared with and improved by other robots.”

Maybe our Jetson’s dream will become a reality sooner than we think!



Harvey, the capsicum-picking robot, in action at a protected cropping facility

SECTION 5 | Engagement



End-User Engagement

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

Professor Peter Corke
Director, Australian Centre for Robotic Vision

The Centre's End User Engagement Strategy guides our leadership role in developing Australia's fledgling robotic vision industry. In 2014 when we began our work in the Centre, we took stock of the current industry in Australia. We found it fragmented and still very much in its infancy. The Centre's role is not to compete with existing businesses in the Australian marketplace, but to identify the companies we can partner with to build and deliver technological solutions that will serve across a range of industries.

Our key focus is on developing robotic vision technologies in those areas where their application is identified as having the highest potential economic benefit to Australia.

To fulfil our leadership role in the new field of robotic vision while maximising our research impact, we are continually working to build and maintain sustainable partnerships across the research sector, together with public and private enterprises. An important part of this work is generating additional income and sources of funding. By increasing our network of end users, we are given the opportunity to

conduct research that attracts investment outside our ARC funding. In 2017, the Australian Centre for Robotic Vision attracted over \$9 million in additional external funding including industry, government, other ARC grant funding and university contributions.

Robotic vision has the potential to impact on all industries and as a result the Centre's End User Engagement Strategy must incorporate a diversity of approaches. The Centre achieves this through:

- / Creating Australia's first Robotics and Computer Vision roadmap
- / Annual Symposium and local demonstrations where existing and potential end users can see our technologies in operation
- / Sector-specific industry workshops engaging directly with businesses to understand what challenges they are facing and which of these challenges can be addressed through robotic vision solutions
- / Public media and events, including news stories, social media, trade journals, website and printed collateral
- / One-on-one consultations where our members host meetings with representatives from various industry segments, related to our targeted impact areas
- / Bringing high profile international conferences to Australia
- / VIP tours, the Centre regularly invites high-profile visitors to our various nodes to showcase our evolving work


- / Publication of research papers and articles in respected trade and peer-reviewed journals (see Outputs p. 95).

We track our engagement activities using a CRM system, which records contacts, meetings and other relevant information. Ultimately, we aim to lobby governments to work with our national robotic vision community to build a strong robotic vision industry in Australia.

CENTRE STAKEHOLDERS

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

Centre Stakeholder Map



End-User Links KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017
Number of government, industry and business community briefings	Annually	10	79	10	54	10	64	10	196
Number and nature of public awareness/outreach programs	Annually	4	40	4	12	12	17	12	64
Number and nature of commentaries about the Centre's achievements (in professional magazines and the popular press, e.g newspapers, TV, etc.)	Annually	5	49	10	94	20	84	20	181
Currency of information of the Centre's website <i>*Number of website updates annually</i>	Annually	12 updates	60 updates	12	116	12	36	12	61
Number of website hits	Annually	50k	34k	200k	98k	300k	110k	500k	504k
Number of talks given by Centre staff open to the public	Annually	5	10	10	12	10	19	10	38

Asparagus Harvesting Partnership kicks off at ANU

In April, the Centre's ANU node partnered with Mulyan Pty Ltd, an asparagus farm in Cowra NSW, to develop an autonomous harvesting robot for green asparagus.

"Harvesting asparagus is difficult to achieve using traditional agricultural machinery due to the unstructured and sparse nature of the crop and the fact that only a part of the crop is harvested each day. We have a team of farm hands that go back over every acre of asparagus in the farm every day during the harvest period, making asparagus one of the most expensive crops in the portfolio in terms of manpower," said Ed Fagan, of Mulyan Farms.

The Centre's ANU node is carrying out an 18-month project to develop a prototype vision-based harvesting system that identifies asparagus spears for harvest in real time using dense geometric SLAM techniques and image-based visual servoing to cut, capture and collect the spears, storing them in cold water to seal in their maximum taste as quickly as possible.



Centre researchers from ANU are working with Mulyan Pty Ltd, an asparagus farm in Cowra NSW to develop an autonomous harvesting robot for green asparagus. *Image courtesy of Alex Martin*

Communication and media engagement

 **500,000** Social media hits on Centre research

Our communication and media engagement activities target a broad audience, from our closest stakeholders, to the general public.

Our activities in this area include:

- / promoting the Centre through traditional media channels, targeting local, national and international news, current affairs programs and other media outlets
- / maintaining several social media channels with targeted messaging and content
- / revamping our website (roboticvision.org) as our core repository of information about the centre, by promoting our activities, achievements and services
- / commissioning STEMmatters to write and publicise articles promoting the Centre's research
- / holding public forums and presentations to primary and secondary schools, as well as professional associations and interest groups, to showcase our work and give people the opportunity to interact with robots first hand.

TRADITIONAL MEDIA

In 2017, the Centre ran a targeted media campaign with the objective of promoting the Centre and our researchers, through stories in popular publications. We engaged specialist science communication consultancy, STEMmatters to help us achieve our goals. This year-long campaign showcased our significant wins and achievements. We published and promoted more than 60 media releases and news stories throughout 2017. These pieces led to 181 stories about the Centre in print, online and television channels. All stories are available in the News and Media section of our website.

SOCIAL MEDIA

We maintain a range of social media channels targeting different users and audiences. Our social media feeds combine a range of images, video presentations and demonstrations, as well as links to research updates, upcoming events, researcher profiles, together with news and related articles. A number of our social media feeds are cross-linked with one another, allowing followers to explore our Centre more deeply

based on their specific interests. Our social media channels include: Twitter, Instagram, Facebook, Google+, LinkedIn and our images are all viewable on Flickr. We now have more than 1,500 followers on Twitter.

WEBSITE

roboticvision.org is at the core of our online presence, providing comprehensive information about the Centre including our purpose, detailed profiles of our research programs, people, field testing, news and video, access to our publications, opportunities for industry collaboration and media linkage, workshops and events.

The website is also a key point of contact for our researchers around the world, giving them online access, through a secure portal, to Centre operational information, collaborative research forums, and policies and procedures including the Centre's treatment and protection of intellectual property.

COMMUNITY OUTREACH

In addition to supporting industry and professional associations to keep up with the technology and opportunities robotic vision





Image courtesy of the Queensland Museum

presents, the Centre is actively involved in building capacity at the high school, undergraduate and graduate levels, as well as making robotic vision more accessible to the general public. In 2017, these community outreach activities included:

- / Showcasing the Centre's research and robots at major public events including QUT's Robotronica, the World Science Festival (Brisbane) and The University of Adelaide's showcase, Ingenuity
- / Sponsoring the Unmanned Aerial Vehicles (UAV) Outback Challenge, held in Dalby, Queensland. High school students from across the globe travel to Australia to compete in the event, one of the world's biggest drone competitions.
- / Guest and keynote speaker appearances at industry and professional association events
- / Hosting public forums and tours of the Centre for future leaders and business groups
- / Hosting school group visits ranging from preschool to high school at the Centre's headquarters
- / Collaborating with existing science extension events such as Coding for Kids.

DISTINGUISHED VISITORS AND EXCHANGE OPPORTUNITIES

In 2017, the Centre hosted 40 international visitors from 30 organisations including Raia Hadsell, Senior Research Scientist at Google Deepmind; Professor Henrik Christensen, Director of the Institute for Contextual Robotics at UC San Diego; Dr Rodney Brooks, Founder, Chairman and CTO of ReThink Robotics; and Professor Ronald Arkin, Georgia Tech.

PUBLICATIONS AND PRESENTATIONS

Our researchers published a total of 71 papers, including 47 conference papers and 24 journal articles. We have included a full list of the Centre's publications from 2017 under Outputs (p. 95).

We also led 196 government, industry and business community meetings and presentations, bringing the complex world of robotic vision to a wide audience.

NETWORKS

Building research networks at home, and overseas, is an important part of sharing and expanding our knowledge in the computer vision and robotics sectors, together with

growing our Centre's profile as a leader in the field of robotic vision. We actively encourage our chief investigators, research fellows and PhD researchers to visit each other's laboratories, including those of our international partners and to host visits in return. This allows us to work more effectively together and build a truly international community around robotic vision.

CONFERENCES AND WORKSHOPS

As well as giving us the opportunity to showcase our work, international conferences play a key role in growing our networks both at home and abroad. The Centre supports our researchers' attendance at these events to share our findings, learn about important work other researchers are doing, and to recruit new researchers and students. These activities lift our international reputation as a research centre, which helps us attract world-leading researchers to Australia, further enhancing our international linkages and collaborative partnerships. Centre researchers presented 56 papers at 26 conferences in nine countries, including here in Australia. We also ran several very successful workshops as part of the programs of international conferences. At the



Images courtesy of QUT Media

Robotics: Science and Systems Conference (RSS) in Boston, our “New Frontiers for Deep Learning in Robotics” workshop featured luminary speakers, including convolutional neural network founder Yann LeCun of Facebook AI, who spoke about challenges and opportunities in deep learning in robotics. With approximately 200 people in attendance, it was one of the most well attended workshops at the conference. As well as invited keynotes, the workshop was accompanied by a series of spotlight presentations and a poster session, with a best paper awards ceremony (sponsored by the Centre, Osaro and Google) and a workshop dinner.

Another workshop, “Brave new ideas for motion representations in videos II” was held at the Computer Vision and Pattern Recognition conference (CVPR), held in Honolulu. The workshop was enthusiastically received by almost 100 attendees.

At the same conference, the Centre also ran a “Deep Learning for Robotic Vision” workshop. This addressed recent advances in deep learning techniques, along with

the many unique challenges presented by robotics and the need for new approaches. This workshop focused on work related to deep learning techniques in computer vision as they applied to a broad range of robotic devices, from self-driving cars, drones and bipedal robots.

We are looking forward to the International Conference on Robotics and Automation (ICRA) being held in Brisbane in May 2018. The Chair of the Centre’s Advisory Committee, Alex Zelinsky, and our Director, Peter Corke, led the successful 2014 bid for Australia to host this high-profile event. It will be the first time the IEEE Robotics and Automation Society’s flagship conference will be held in the southern hemisphere. We have been working closely with the Society and the event organisers, to ensure the resounding success of this event. ICRA is a premier international forum for those working in robotics research, and will create an opportunity for us to comprehensively promote and profile our work to an audience of over 2,000 world-leading researchers.

MANUFACTURING MEETS ROBOTIC VISION

With the support of PwC, a valued member of our End-User Advisory Board, the Centre ran a series of innovation challenges. Aimed at identifying people and companies in the Australian robotics and vision communities that we could collaborate with to solve challenges, leverage our expertise and work with on future commercialisation opportunities.

Each challenge in the series is focused on a different industry sector. Throughout the year the Centre hosted two events, the first was held in Brisbane in March, focusing on manufacturing; the second focused on food production both were held in Brisbane in June.

These events created huge interest from the companies that participated.

Commercialisation

A core objective of the Centre is developing unique intellectual property that we can license or sell in commercial markets. We have established guidelines and policies on how to manage our intellectual property, along with approaches for open-sourcing some of our technologies. The Centre maintains a register of all intellectual property that we develop during the course of our activities, along with any dependencies on background intellectual property. Individual researchers assign the intellectual property they have developed to their host university, as collaborating

partners in the Centre Agreement. The Agreement sets out the terms of ownership among project parties, who consult with the Centre Executive on protecting and commercialising intellectual property as it is developed. In addition to commercialising our intellectual property, we can achieve further impact with our research through our open source policy. In cases where intellectual property we create has low potential for commercialisation, the Centre recommends the “BSD-3-Clause License” created by the Open Source Initiative (opensource.org). Our monthly Centre News bulletin promotes the

different code and their repositories, that our researchers open source.

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



Cartman picking and packing at the Amazon Robotics Challenge

AlphaOne and Nuarda: Australian Centre for Robotic Vision's newest start-ups

Greater than 80% of start-up founders in Australia are university graduates. Universities Australia released a report in March 2017, highlighting the important role universities play in the country's largest job creating sector, projecting they will create more than half a million jobs in future decades. Currently the sector is contributing over \$160 billion to the Australian economy each year.

With our strong culture of supporting entrepreneurship, Centre researchers launched two new enterprises: AlphaOne.ai and Nuarda.

AlphaOne.ai

Centre Chief Investigator Tom Drummond from our Monash node, teamed up with two former students to create AlphaOne.ai. The start-up provides consulting services that help businesses learn how to deal with issues like handling big data for their customers, logistics and processes, or dealing with a lot of visual inspection. These are expensive problems for companies to solve.

"I think a lot of business leaders are looking at these new technologies and thinking they should find out more," Tom said.

"The expertise doesn't reside in these businesses, so they need an external party to help develop, then support that technology into the future, so it can be integrated into their business processes."

That's where AlphaOne.ai comes in.

The start-up also serves another need, according to Tom: Keeping talented people from leaving Melbourne.

"There are lots of opportunities for them to move to Silicon Valley and take jobs there, but Melbourne is a great city," Tom said. "So having an opportunity to prevent the 'brain drain' provides a way to build up the entrepreneurial activity within the city."

In addition to their work in Melbourne, the group is already starting to handle clients outside of Australia.

Nuarda

Nuarda has evolved out of the Centre's CloudVis project, which allows computationally intensive processing tasks to be done in the cloud. People and robots can access the latest algorithms without needing specialised hardware, or software.

Research engineer, Steve Martin and software engineer, Gavin Suddrey launched the start-up as a means of bringing computer vision technologies to the mass consumer market.

The pair identified the task of taking state-of-the-art technologies from the lab to the real-world as a significant challenge, including the amount of domain expertise needed to set up many of these systems. At the same time, QUT bluebox announced its robotics accelerator program and the team saw an opportunity to gain some seed funding. The program also gave them invaluable guidance on defining their market and they further refined their product: Nuarda.

As part of the team's involvement in QUT bluebox, they entered Nuarda into the 'QUT bluebox Innovation Challenge', making their way into the finals from a field of 300 applicants.

Nuarda won an industry-sponsored prize from the people at ShapeLabs.

Steve and Gavin are now working on commercialising their product.

International and National Links & Networks KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017
Number of international visitors and visiting fellows	Annually	2	8	8	11	8	22	8	40
Number of national and international workshops held/organised by the Centre	Annually	1	1	2	15	3	10	3	13
Number of visits to overseas laboratories and facilities	Annually	2	11	8	25	8	56	8	91
Examples of relevant interdisciplinary research supported by the Centre <i>*Number of interdisciplinary projects within the Centre</i>	Annually	0	0	1	3	2	5	3	3
Number of talks given by Centre staff open to the public	Annually	5	10	10	12	10	19	10	38

SECTION 6 | People





Ninety researchers across the Centre travelled to Tangalooma Resort Moreton Island for the Centre's fourth annual symposium RoboVis 2017.

Research Training and Professional Development

Research Training

We strive to provide a rich training experience for early career researchers and over the Centre's seven-year life we will train 24 postdoctoral fellows and 80 PhD students. Eighteen new PhD researchers were welcomed to the Centre this year, along with three new postdoctoral researchers.

The Centre offers a range of professional development opportunities for its researchers and staff, such as:

- / Travel to other Centre nodes or our overseas partner universities, to experience different research environments
- / The opportunity to attend top international conferences and regular seminars at Centre nodes
- / Meeting and interacting with eminent visiting researchers from our national and international partners

- / Research training for early career researchers, delivered in both face-to-face and online formats and supplemented by a specialist toolbox that we host on the Centre's intranet.

ROBOVIS 2017

Our annual symposium, *RoboVis*, gives researchers from Centre nodes around the country the opportunity to come together as a group and share their research in a relaxed environment. This year, 90 researchers and guests, gathered at the beautiful Tangalooma Resort on Queensland's Moreton Island to focus on areas of knowledge leadership and career development, share their research and connect with their interstate colleagues. We were joined by the Chair of our Centre Advisory Committee, Dr Alex Zelinsky, and members of our End-User Advisory Board, Russel Rankin, Alan Davie, Trent Lund, Peter Katsos, Rob Wood and Andrew Harris.

RoboVis 2017 included guest speakers, a technical program, three-minute thesis spotlights, technical demonstration sessions and a research poster session. Highlights of the event included a "Women in STEM" breakfast forum with QUT's Mathilde Desselle and an industry presentation from PwC's Trent Lund, who is also a member of the Centre's End-User Advisory Board. We enjoyed some socialising with a trivia night, a bonfire night, wild dolphin feeding, along with the eagerly anticipated presentation of the 2017 Centre Awards.

We closed the symposium with a knowledge leadership workshop for our research fellows and PhD researchers. This was the last in a 2017 series of customised knowledge leadership training modules to help support our early career researchers' professional development.



▲ The Centre ran 8 knowledge leadership workshops in 2017, ending with a full-day workshop with all Centre PhD researchers and research fellows at RoboVis

PhD Researcher **Peter Anderson**, pictured here with Centre Director, Peter Corke, won the 2017 Award for Best Three-Minute PhD Spotlight, plus the award for Best Poster Session



Research Fellow **Trung Pham**, pictured here accepting his award for Best Technical Demonstration Session from Peter Corke.



Centre Research Fellow (now Chief Investigator), **Niko Sünderhauf**, won the 2017 Award for Best Centre Citizen. This award is presented to the researcher who best exhibits the spirit of collaborative enquiry needed to advance research in robotic vision and make the Centre a success.



The Centre's Amazon Robotics Challenge Team received more recognition for their accomplishments, taking out the Award for Best Team Project. Members present at RoboVis 2017 included (L to R): Juxi Leitner, Douglas Morrison, Norton Kelly-Boxall, Sean Wade-McCue, Zheyu Zhuang, Vijay Kumar, Trung Pham, Thomas Rowntree and Peter Corke. (Fellow team members, Adam Tow, Anton Milan, Steve Martin, Ming Cai, Chris Lehnert, Saroj Weerasekera, Riccardo Grinover, Tom Hunn, Matt McTaggart, Nathan Perkins, Gerard Rallos, Rohan Smith, Jordan Erskine, Anthony Gillespie, Darryl Lee, Andrew Razjigaev, Alec Gurman and Ian Reid weren't present for this photo, but do feature in our story on p. 35)



Ben Harwood, Tom Drummond and Vijay Kumar accepted the 2017 Award for Best Collaborative paper from Centre Director, Peter Corke. They also accepted this award on behalf of Gustavo Carneiro and Ian Reid for their work on “Smart Mining for Deep Metric Learning”.



The 2017 Award for Contribution to Public Debate was awarded to **Michael Milford** (not present) for his outstanding thought leadership on the topic of robotic vision technology, in particular its impact on society.



Niko Sünderhauf, Juxi Leitner, Vijay Kumar, Trung Pham

(fellow award-winners Gustavo Carneiro and Michael Milford were not present) accepted their award from Peter Corke for Best Profile-Raising Event In Robotics And Computer Vision Communities. This was for their workshop “Deep Learning for Robotic Vision” that was held at the 2017 Conference on Computer Vision and Pattern Recognition (CVPR) in Honolulu, Hawaii.

LEADERSHIP DEVELOPMENT

The Centre strongly focuses on developing our early career researchers. We are developing our PhD researchers and research fellows into knowledge leaders who will transfer their robotic vision knowledge into either industry or academia. We envisage a time when we will see our graduates and alumni founding new companies in Australia's fledgling robotic vision industry and we want to ensure we give them the skills and support they need to take robotic vision to the world.

We have developed a custom program of knowledge leadership workshops with organisational development specialists, Evexia, designed to give early career researchers the skills they need to lead others in any sphere of influence. We held two workshops at each Centre node and another whole-of-Centre workshop as part of our RoboVis 2017 activities. We also delivered custom in-house training at each node on the subjects of Careers and Start-ups.

The result of our focus on developing research leadership, as well as leading researchers, is clearly demonstrated by the level of recognition our Centre staff gain for their talent, both at home and abroad. The Australian Centre for Robotic Vision continues to be an attractive prospect to other researchers to join us bringing our total cohort to 75 early career researchers

including 18 research fellows and 57 PhD researchers. Our international quota remains high with all our research fellows and 65% of our PhD researchers hailing from more than 15 countries.

Gender diversity

We have to acknowledge that all our Centre investigators are male, reflecting significant gender imbalance. These imbalances are acute in computer science and only slightly better in engineering. We are actively pursuing ways to redress this, by employing a number of tactics, such as addressing unconscious bias in advertising when recruiting for research fellows and PhD researchers.

We also engaged external consultants, Gender Matters, to help us investigate other ways to build gender diversity in our field. In October 2017, at our annual RoboVis symposium, the Centre launched our Robotic Vision Gender Equity Plan. This plan incorporated feedback from across the centre, in addition to the recommendations from the Gender Matters' report. A copy of our Robotic Vision Gender Equity Plan can be downloaded from the Centre's website. The Chair of our Centre Advisory Committee, Alex Zelinsky is a recognised Male Champion of Change.

In addition to these initiatives, Centre COO Sue Keay has been named one of

Australia's inaugural 30 Superstars of STEM, encouraging female role models in robotics and computer vision. Sue has also been working with the AnitaB.org Institute to bring the internationally recognised Grace Hopper conference to Australia in 2019 to celebrate the achievements of women technologists. These female-focussed conferences in Australia have been credited with improvements in gender equity in fields such as Physics (Women in Astronomy) and Mathematics (Women in Mathematics).

We also have a number of high profile female researchers associated with the Centre including former Centre Advisory Committee, member Professor Michelle Simmons, who was named Australian of the Year in January 2018. Centre Research Affiliate Dr Anjali Jaiprakash who has been named one of 25 Women in Robotics to know about and was recently named one of MIT Tech Review's Top innovative researchers under 30 in the Asia-Pacific region. And new Centre advisory committee member Kylie Ahern, who founded Australia's Science magazine, COSMOS.

Our number of female PhD researchers is improving each year, currently we have 54 PhD researchers, five of whom are female with three new female PhD researchers joining the Centre in early 2018.

Diversity of Staff

Position	#	FTE	Gender Ratio (female:male)
Chief Investigator	15	3.6	0:15
Partner Investigator	6	0.3	0:6
Associate Investigator	14	—	0:14
Operations Team - Professional Staff*	8	4.3	7:1
Research Engineers/Software Engineer	3	1.8	0:3
Project Managers & Officers	3	1.2	3:0
Research Fellows	14	14	2:12
PhD researchers	54	54	5:49

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

Dr Sue Keay: Superstar of STEM

The Centre's own Chief Operating Officer, Dr Sue Keay, was recognised by Science & Technology Australia as one of their inaugural "Superstars of STEM". Sue was one of only 30 women, from around the country selected from more than 300 applicants. The program's goal is to help women learn how to speak about their science and inspiring others to consider careers in STEM (science, technology, engineering and mathematics).

Sue has been the Centre's COO since its inception in 2014, and is passionate about raising the profile of technology among women, in an effort to make sure Australia is poised to take advantage of new technologies. Speaking about "Superstars of STEM", Sue said, "The program is vital to ensure women participate, and stay engaged, in STEM."

"There are very few women in technology. There are even fewer in robotics and computer vision. Yet the opportunities to change the world through these fields is incredible," she said.

Science & Technology Australia's President, Professor Emma Johnston said "We want Australian girls to realise that there are some amazing, capable and impressive women working as scientists and technologists too, and that they work in and out of the lab in places you might not expect,"

As a mother of two daughters, Sue has always felt a responsibility to make sure there are more STEM opportunities for young women.

"There were significant improvements in the workplace from my mother's generation to my own, but I don't see a similar amount of progress benefitting my daughters' generation. This has to change and I'm keen to be an active participant in changing things," she said.

Dr Sue Keay, Centre Chief Operating Officer.



Mentoring

Our goal for all Centre researchers is for them to have a mentor they can learn from, while building a solid and trusting relationship. A mentoring relationship gives researchers an informed second opinion, and helps them gain insights and awareness of their own performance through the eyes of a 'critical friend'. Through this relationship, researchers can identify their personal development needs, as well as learning from the experience and advice of their mentor.

Researchers select their own mentors, with support from the Centre's Administrative Coordinator. We have a number of mechanisms to support our mentoring program, ensuring mentoring relationships are working effectively for all concerned. We intentionally allow a level of flexibility in the mentoring relationship, with a view of it being a support role, rather than a replacement of existing supervisors.

The role of a mentor may include:

- / sharing their expertise and experience to help individual researchers develop their talents

- / listening, clarifying, giving feedback and challenging researchers to consider issues from different points of view
- / opening doors, assisting researchers to network and further develop their careers
- / providing an avenue for researchers to raise issues they may have and discussing them in a safe space.

The role of a researcher working with a mentor may include:


- / taking responsibility for identifying and achieving development towards career goals
- / initiating meetings with their mentor, managing meeting dates and times and setting the agenda
- / being open to and appreciative of different perspectives, as well as constructive and honest feedback
- / being considerate of the demands placed on their mentor's time.

Education

During May, the Centre's QUT Node launched the Robot Academy. Providing university-level, short video lessons, as

well as full online courses, the academy assists undergraduates from around the world to understand and prepare for, robotic and robotic vision technologies. Based on the Centre's earlier "Massive Open Online Course" (MOOC) efforts, these resources are available for use, free of charge, at any time. QUT also hosts the RoboHub podcast as part of the Robot Academy.

Since its launch, our Robot Academy has had more than 125,000 views from 30,000+ unique users, in over 100 countries. More than half of these visitors return regularly. A number of universities are adopting lessons from QUT's Robot Academy for use in their own teaching.



2017 was also the year that Centre Director Peter Corke released the Second Edition of his book **Robotics, Vision and Control: Fundamental Algorithms in MATLAB**, published by Springer.

"Robotic vision is the combination of robotics and computer vision and involves the application of computer algorithms to data acquired from sensors. This second edition is completely revised, updated and extended with coverage of Lie groups, matrix exponentials and twists, inertial navigation, differential drive robots, lattice planners, pose-graph SLAM and map-making, restructured material on arm-robot kinematics and dynamics, series-elastic actuators and operational-space control, Lab colour spaces, light field cameras, structured light, bundle adjustment and visual odometry and photometric visual serving."

Research Training and Professional Development KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017
Number of professional training courses for early career researchers, staff and PhD students	Annually	0	0	4	7	4	5	4	15
Number of Centre attendees at all professional training/development courses offered by the Centre (include courses offered for external stakeholders and clients)	Annually	0	0	30	120	50	83	50	211
Number of new postgraduate students working on core Centre research and supervised by Centre staff (include PhD, Masters by research and Masters by coursework)	Annually	10	6	20	23	10	17	20	18
Number of new postdoctoral researchers recruited to the Centre working on core Centre research	Annually	7	6	9	11	1	3	8	3
Number of new Honours students working on core Centre research and supervised by Centre staff	Annually	4	10	8	7	8	12	8	40
Number of postgraduate completions and completion times, by students working on core Centre research and supervised by Centre staff	Annually	0	0	0	0	0	6	20	5
Number of Early Career Researchers (within five years of completing PhD) working on core Centre research	Annually	7	6	16	14	16	16	16	14
Number of students mentored	Annually	45	6	90	29	90	157	90	157
Number of mentoring programs offered by the Centre (include programs for students, new staff, external stakeholders and clients)	Annually	0	1	2	2	2	2	2	2

Our People

The Centre's largest expenditure is on people, supporting research fellows and PhD researchers. We welcomed 18 new postgraduate students to the Centre this year, in addition to three new postdoctoral research fellows. We also welcomed 40 new Honours students, exceeding our target of eight students for the year.

The Centre brings together a critical mass of outstanding researchers from around the world, from our accomplished PhD and early career researchers, right through to veterans of industry and academia. Their expertise spans key robot vision areas such as machine, mapping and navigation, visual servo control, 3D reconstruction, low-level and high-speed vision and distributed systems. Our expertise and accomplishments, are being recognised on the international stage, as well as at home. Our extensive media engagement campaign that ran throughout the year resulted in excellent coverage of our transformative work, showcasing the skills and talents of our researchers.

Our strong reputation for quality research and researchers means our people are in demand. In 2017 we said farewell to four of our Research Fellows. We would like to acknowledge and thank:

- ✓ Dr Chuong Nguyen, from our ANU node who joined CSIRO as a Research Engineer and Experimental Scientist, while continuing his association with the Centre as a Research Affiliate.
- ✓ Dr Guosheng Lin was a Centre Research Fellow at our University of Adelaide node, leaving Australia for Singapore and taking up the position of Assistant Professor at Nanyang Technological University
- ✓ Dr Markus Eich, who was based at our QUT node, returned to Germany to join the Continental Automotive Group
- ✓ Dr Anoop Cherian was based at our ANU node. He is now a Research Scientist at the Mitsubishi Electric Research Laboratories (MERL) in Cambridge, USA

Attracting and retaining talent in the Centre is a major challenge in this fiercely competitive market. To address this and mitigate any risk,

we have a succession plan to ensure we are always growing and supporting leadership talent within the Centre. We actively work to promote and maintain the Centre as a rewarding place to work, while continuing to make it an appealing prospect for potential researchers. We also keep in close contact with our alumni, who are often a good source of talent, with key industry experience, together with contacts that can potentially help the Centre.

NEW CHIEF INVESTIGATORS

Congratulations are in order for four of the Centre's Associate Investigators, who in 2017 were promoted to the position of Chief Investigator. They are Jonathan Roberts, Matthew Dunbabin, and Niko Sünderhauf from QUT, and Tat-Jun Chin from the University of Adelaide.

Jonathan joined the Centre in 2014 and is a Professor in Robotics at QUT. His research interests are in field robotics together with design and medical and healthcare robots.

Matthew joined the Centre in 2014. His wide-ranging research interests include vision-based navigation, image-based habitat classification, adaptive sampling and path planning, as well as robot and sensor network interactions.

Niko joined the Centre as a Research Fellow in 2014. In 2017, he was appointed as a lecturer at QUT, became an Associate Investigator with the Centre before becoming a Chief Investigator in 2017. His research interests include robust place recognition for changing environments.

Tat-Jun joined the Centre in 2014 and is an Associate Professor at the University of Adelaide. His award-winning research explores the areas of robot estimation and geometric optimisation.

CENTRE CULTURE

As we build our Centre we are also all developing a Centre culture, with a set of pervasive values, beliefs and attitudes that characterise who we are. As people join the Centre, we introduce the Centre's Vision and Values, our history, and our plans for the future via an extensive onboarding process. We are a nimble enterprise, with the flexibility and capacity to handle diversity and a range of new ideas.

With Nodes across Australia and partners spanning the globe, the internet is a vital tool to ensure Centre cohesion and a collaborative culture. The Centre maintains an active intranet shared workspace using Atlassian's Confluence. Our staff and partners can also connect, interact and share information through our public social media channels.

We produce two regular online newsletters for the Centre. Our monthly internal newsletter gives all Centre personnel the opportunity to share and celebrate each other's successes. Our quarterly public newsletter keeps industry, government, collaborators, visitors and friends up to date on the Centre's activities. We share information on our upcoming public forums and other events, news and triumphs in the public research space and accomplishments on our public website. Our external newsletter has a readership of 1200 people and growing.

Virtual team best practice

One of the central aims of ARC Centres of Excellence is to encourage collaboration between organisations. The glue that supports this collaboration is a hub of professional staff, located at our Centre's nodes.

Our professional staff are a vital part of our Centre, breaking down geographic barriers and helping us achieve our objectives.

As a team, Centre professional staff have developed their own Vision and Team Charter that supports their collaborative work and ensures high-quality service to the Centre's researchers. Some of the best practice activities our professional team use are:

- / Developing team norms and communication protocols
- / Abandoning email in favour of instant messaging, videoconferencing and phone calls
- / Prioritising the Centre over local thinking
- / Committing to regular meetings despite the challenges of conflicting work schedules
- / Gathering and sharing, information with the right people, at the right time.

As well as looking after the ARC reporting requirements at each node, our operations team also

manage training opportunities for Centre researchers, coordinate all Centre events, and organise travel to our overseas partner organisations. Our operations team encourages everyone to collaborate across the nodes, organise local events, and encourage researchers to contribute to Centre activities by providing pieces for our newsletters, and sharing success stories across the Centre and beyond.


“The Centre professional staff are a relied-upon network characterised by their innovative and proactive approach to solving challenging problems. Assisting with the many events and reporting requirements that are part of running a successful national research endeavour,” says Monash node leader Professor Tom Drummond.



Operations team members at the team retreat during RoboVis 2017 (l to r) Carol Taylor, Sandra Pedersen, Tracy Kelly, Kate Aldridge, Sarah Allen, Thuy Mai. Sue Keay and Tim Macuga not pictured.

Centre staff profiles

Chief Investigators (CIs)



Gustavo Carneiro
University of Adelaide



Tat-Jun Chin
University of Adelaide



Peter Corke
QUT



Tom Drummond
Monash



Matt Dunbabin,
QUT



Richard Hartley
ANU



Hongdong Li
ANU



Robert Mahony
ANU



Michael Milford
QUT



Ian Reid
University of Adelaide



Jonathan Roberts
QUT



Chunhua Shen
University of Adelaide



Niko Sünderhauf
QUT



Anton van den Hengel
University of Adelaide



Gordon Wyeth
QUT

Associate Investigators (AIs)



Luis Mejias Alvarez
QUT



Nick Barnes
ANU/Data61



Ross Crawford
QUT



Anthony Dick
University of Adelaide



Anders Eriksson
QUT



Clinton Fookes
QUT



Jason Ford
QUT



Felipe Gonzalez
QUT



Jonghyuk Kim
ANU



Tristan Perez
QUT



Fatih Porikli
ANU/DATA61



Qinfeng 'Javen' Sh
University of Adelaide



David Suter
University of Adelaide



Jochen Trumpf
ANU

Partner Investigators (PIs)



Paul Newman
University of Oxford, UK



Philip Torr
University of Oxford, UK



Marc Pollefeys
ETH Zurich, Switzerland



Frank Dellaert
Georgia Tech, USA



François Chaumette
INRIA, France



Andrew Davison
Imperial College
London, UK

Research Fellows



Suman Bista
QUT



Feras Dayoub
QUT



Thanh-Toan Do
University of Adelaide



Basura Fernando
ANU



Viorela Ila
ANU



Vijay Kumar
University of Adelaide



Yasir Latif
University of Adelaide



Juxi Leitner
QUT




Vincent Lui
Monash



Chao Ma
University of Adelaide



Valerio Ortenzi
QUT



Trung Pham
University of Adelaide



Sareh Shiraz
QUT




Qi Wu
University of Adelaide

Associated Research Fellows




Ravi Garg
University of Adelaide




Alireza Khosravian
University of Adelaide



Hamid Reza Tofighi
University of Adelaide



Fan Zeng
QUT



Adam Jacobson
QUT

Research Affiliates



Cesar Cadena
ETH Zurich



Donald Dansereau
Stanford University



Zongyuan Ge
IBM Research, Australia



Stephen Gould
Amazon



Anjali Jaiprakash
QUT



Pan Ji
University of Adelaide



Laurent Kneip
Shanghai Tech



Chris Lehnert
QUT



Chris McCool
QUT



Mark McDonnell
University of
South Australia



Anton Milan
University of Adelaide



Chuong Nguyen
CSIRO



Ajay Pandey
QUT



Thierry Peynot
QUT



Ahmet Sekercioglu
Monash University



Ben Upcroft
Oxbotica



Liao 'Leo' Wu
QUT

PhD Researchers



Peter Anderson
ANU



Gil Avraham
Monash University



Ming Cai
University of Adelaide



William Chamberlain
QUT



Shin Fang Ch'ng
University of Adelaide



Arif Chowdhury
ANU



Thanuja Dharmasiri
Monash University



Rafael Felix
University of Adelaide




Sourav Garg
QUT



Luis Guerra
Monash



Jian "Edison" Guo
ANU



Ben Harwood
Monash University

PhD Researchers (continued)



Mina Henein
ANU



William Hooper
QUT




Mehdi Hosseinzadeh
University of Adelaide



Jasmin James
QUT



Chris Jeffrey
QUT



Norton Kelly-Boxall
QUT



Peter Kujala
QUT



Hui Li
University of Adelaide



Kejie Li
University of Adelaide



Yu Liu
University of Adelaide



Sean McMahon
QUT



Andres Felipe
Marmol Velez
QUT



Medhani Menikdiwela
ANU



Ben Meyer
Monash



Douglas Morrison
QUT



James Mount
QUT




Vladimir Nekrasov
University of Adelaide



Lachlan Nicholson
QUT



Sean O'Brien
ANU



Fahimeh Rezazadegan
QUT



Dan Richards
QUT



Rodrigo Santa Cruz
ANU



James Sergeant
QUT



Tong Shen
University of Adelaide



John Skinner
QUT




Andrew Spek
Monash University



Jean-Luc Stevens
ANU



Timo Stoffregen
Monash



Brendan Tidd
QUT



Adam Tow
QUT

PhD Researchers (continued)




Dorian Tsai
QUT



Sean Wade-McCue
QUT



Xian Wang
University of Adelaide



Josh Weberruss
Monash University



Saroj Weerasekera
University of Adelaide



Huangying Zhan
University of Adelaide



Fangyi Zhang
QUT



Jun Zhang
ANU



Yi "Joey" Zhou
ANU



Bohan Zhuang
(completed)
University of Adelaide



Zheyu Zhuang
ANU




Yan Zuo
Monash University

Associate PhD Researchers



Sadegh Aliakbarian
ANU




Adrian Johnston
University of Adelaide



Cristian Rodriguez
ANU



Ehab Salahat
ANU



Cedric Scheerlinck
ANU

Graduates // Congratulations to all our 2017 graduates!




Juan Adarve
PhD, ANU




Zongyuan Ge
PhD, QUT



Yuchao Jiang
Masters by Research,
University of Adelaide



Zhibin Liao
PhD, University
of Adelaide



Bohan Zhuang
PhD, University
of Adelaide

Centre Operations team



Sue Keay
Chief Operating Officer



Kate Aldridge
Centre Coordinator



Tracy Kelly
Centre Finance & Administration Officer (Part-Time)



Tim Macuga
Communications & Media Officer (Part-Time)



Sarah Allen
QUT Node Administration Officer and PA to Centre Director



Carol Taylor
ANU Node Administration Officer (Part-Time)



Thuy Mai
University of Adelaide Node Administration Officer (Part-Time)



Sandra Pedersen
Monash University Node Administration Officer (Part-Time)


Engineers



Steve Martin
Research Engineer



Alex Martin
Research Engineer



Gavin Suddrey
Software Engineer

Project Officers



Belinda Ward
Project Leader, Humanoid Robotics



Tabetha Bozin
Project Lead, Robotic & Computer Vision Roadmap (August to November 2017)



Sandy Holmes
Project Officer, Robotic & Computer Vision Roadmap (from December 2017)



SECTION 7 | Governance



Governance

Good governance is the responsibility of our Centre Executive Committee, with oversight provided by the Centre's Advisory Committee. The Centre Executive are accountable to the Australian Research Council, a statutory agency responsible for Australia's National Competitive Grants Program, and also represent our four domestic research partners. The Executive are assisted by a Research Committee as described on p. 14. The Centre also has an End-User Advisory Board (EUAB) but will restructure in 2018 to combine the Centre Advisory Committee and EUAB to form a new Centre Advisory Board. The Centre's Organisational Structure is outlined on p. 81.

Centre Advisory Committee

Our **Centre Advisory Committee** oversees the Centre's overall strategic direction. Its membership comprises an independent Chair and up to six committee members representing academia and industry all with expertise in the relevant science along with a proven track record in technology transfer.



Dr Alex Zelinsky AO (Chair)

Alex is Australia's Chief Defence Scientist and head of the Defence Science and Technology Group (formerly the Defence Science and Technology Organisation). Before joining Defence, Alex was Group Executive for Information Sciences at the CSIRO and Director of the organisation's ICT Centre.

Alex was CEO and co-founder of Seeing Machines, a high-technology company that develops computer vision systems. The company was a start-up from the Australian National University in Canberra, where Alex was Professor of Systems Engineering. Seeing Machines is now listed on the London Stock Exchange.

Previously, Alex was a robotics and computer vision researcher at the AIST Electrotechnical Laboratory in Japan. He also taught and conducted research, in computer science at the University of Wollongong. Alex has extensive experience in advising Australian federal and state governments, 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. In 2013, he was awarded the Trevor Pearcey Medal, the ICT industry's premier prize for lifetime achievement in the profession. In 2015 Alex received Engineers Australia's MA Sargent Medal, the most prestigious award made by the College of Electrical Engineers. Alex was appointed an **Officer of the Order of Australia (AO)** in the 2017 Queen's Birthday Honours "for distinguished service to defence science and technology, to systems engineering, and to education as an academic and researcher."



Scientia Professor Michelle Simmons

Michelle is an ARC Laureate Fellow and Director of the highly successful Centre of Excellence for Quantum Computation and Communication Technology and **2018 Australian of the Year**. She has pioneered unique technologies internationally to build electronic devices in silicon at the atomic scale, including the world's smallest transistor, the narrowest conducting wires, and the first transistor where a single atom controls its operation. Her work opens up the prospect of developing a silicon-based quantum computer, a powerful new form of computing with the potential to transform information processing.

Michelle is one of a handful of researchers in Australia to have twice received a Federation Fellowship and a Laureate Fellowship, the Australian Research Council's most prestigious award of this kind. She has won both the Pawsey Medal (2006) and Lyle Medal (2015) from the Australian Academy of Science, for outstanding research in physics and upon her appointment, was one of the youngest Fellows of this Academy. She was named Scientist of the Year by the New South Wales Government in 2012, and in 2014 became one of only a few Australians inducted into the American Academy of Arts and Sciences. In 2015, she was awarded the CSIRO Eureka Prize for Leadership in Science and, in 2016, the Foresight Institute Feynman Prize in Nanotechnology for her work in "the new field of atomic-electronics, which she created". In 2017, Michelle was named L'Oréal-UNESCO Asia-Pacific Laureate in the Physical Sciences.

She is Editor-in-Chief of the online journal, *Nature Quantum Information*.



Professor Mandyam Srinivasan

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

Srini holds an undergraduate degree in Electrical Engineering from Bangalore University, a Master's Degree in Electronics from the Indian Institute of Science, a PhD in Engineering and Applied Science from Yale University, a DSc 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, the Royal Society of London, the Royal Institute of Navigation, and the Academy of Sciences for the Developing World, and Inaugural Federation Fellowship. He won Australian Prime Minister's Science Prize in 2006, the UK Rank Prize for Optoelectronics in 2008, the Distinguished Alumni Award of the Indian Institute of Science in 2009. He was also awarded Membership of the Order of Australia (AM) in 2012, in 2014 he won the Queensland Science Championship plus the Harold Spencer-Jones Gold Medal of the Royal Institute of Navigation.



Professor Hugh Durrant-Whyte

Hugh is the Chief Scientific Advisor, Ministry of Defence, United Kingdom. He is on leave from his role as a Professor, and ARC Federation Fellow, at The University of Sydney. His research work focuses on robotics and distributed sensor networks, having 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 (2006-2010). He is a Fellow of the Academy of Technological Sciences and Engineering (FTSE), a Fellow of the Institute of Electrical and Electronic Engineers (FIEEE), a Fellow of the Australian Academy of Science (FAA), a Fellow of the Royal Society (FRS). He served as the Chief Executive Officer of National ICT Australia Limited (NICTA) from December 2010 to November 2014. He was awarded the MA Sargent Medal in 2017.



Sir Michael Brady

Mike is Emeritus 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. 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 is a Fellow of the Royal Society, the Royal Academy of Engineering, the Academy of Medical Sciences, the Institute of Physics and the British Computer Society. He is also an Honorary Fellow of the Institution of Engineering and Technology. In 2000, he was awarded the Institution of Engineering and Technology Faraday Medal and the IEEE Third Millennium Medal for the UK, the Henry Dale Prize by the Royal Institution in 2005 and the Whittle Medal by the Royal Academy of Engineering in 2010. He was knighted in the 2004 New Year's honours list for his services to engineering.

Centre Executive Committee

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

The Committee includes representatives from the Centre's four domestic partners—QUT, the University of Adelaide, the Australian National University and Monash University—and provides the leadership and direction that is critical to the successful operation of our Centre and achievement of our ambitious research program. The members of our Executive Committee are:



Centre Director, Distinguished Professor Peter Corke (QUT)



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



Chief Operating Officer, Dr Sue Keay



Professor Robert Mahony (Australian National University)



Professor Tom Drummond (Monash University)

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

End-User Advisory Board

Our End-User Advisory Board includes up to six representatives from industry, particularly in the Centre's main areas of application. Members serve a two-year term.

The End User Advisory Board's role with the Centre is to share their experience of industry needs and give advice on how to best align the Centre's work with those needs. Once a year, the End User Advisory Board delivers a brief report to the Centre Director and Executive Committee, which provides comments and advice on the Centre's End-User Engagement activities.

In 2017, the End User Advisory Board met twice including in June at the PwC headquarters in Sydney and again at the Centre's annual *RoboVis* symposium in October.



Russel Rankin (Chair)

Russel Rankin has more than thirty years' experience in the food and beverage industry in various senior commercial and research positions. He has an inherent ability to identify and drive innovation opportunities that deliver a competitive advantage, and he has extensive experience in commercialising innovations from within business and research organisations.

Russel is Director and Founder of Food Innovation Partners Pty Ltd, a company that makes connections between commercial enterprises, research organisations, government agencies, financiers, and marketing and industry bodies. His company delivers business, innovation and commercialisation services to the food industry, as well as business development services for companies and research organisations. Through his ability to translate research outcomes into commercial competitive advantage, and translate commercial opportunities back into viable research strategies, he is able to match food companies with research providers, develop pre-competitive, syndicate projects with multiple commercial partners, help companies' access government support programs, assess equity and make acquisitions, and help food businesses develop and commercialise new ideas.

Prior to starting Food Innovation Partners, Russel was General Manager of Innovation with the National Food Industry Strategy, a federal-government initiative established to provide leadership to Australia's food industry. Before venturing into the commercial arena, Russell worked in food research with the CSIRO for more than 25 years.

Russel is an entrepreneurial thinker. As well as Food Innovation Partners, has established a number of new businesses designed to take a range of new innovative food and beverage products to market.

In addition to chairing our End User Advisory Board, Russel is Chair of the Advisory Board to Queensland Department of Agriculture and QUT's Agricultural Robotics program, a member of the Advisory and Assessment Board for the South Australian Government's Advanced Food Manufacturing program, member of the Advisory Board to KFSU Pty Ltd, a company that makes dietary fibre from sugar cane, and Director of The Food Market Company, Freshly & Co and Beauty Drink Pty Ltd.

Russell has previously been a Board member on the Clean Technology Food & Foundries Investment Program, an initiative of AusIndustry under the Commonwealth Department of Industry, Innovation, Climate Change, Science, Research and Tertiary Education and has been an Advisory Board Member for Oski Drinks Pty Ltd.



Alan Davie

Alan has worked across a range of industries in Australia, the US, the UK, South East Asia and the Pacific. He has over forty years' experience in planning, assessment, engineering and construction plus management of major resources, urban development and infrastructure redevelopment projects, tourism feasibility and planning studies. He has qualifications in project management, town planning, environmental engineering and civil engineering, and is experienced in people management, company management and stakeholder communications.

Alan has held long-term positions on company and advisory boards for companies such as Sinclair Knight Merz Pty Ltd, Project Dynamics Pty Ltd, ANU Enterprise Pty Ltd, Australian Scientific Instruments Pty Ltd, Very Small Particle Company Pty Ltd and Griffith University's School of Environmental Engineering.

He is now the Managing Director of his own consultancy, which carries out assessment and gives advice on strategic infrastructure projects.

Alan spent a large part of his career with a global engineering company, during which time he held a number of management positions, including Queensland and Pacific Business Development Manager, Queensland Operations Manager, General Manager Water and Environment and General Manager International Development Assistance.



Andrew Harris

Andrew Harris is a Professor of Chemical and Biomolecular Engineering at The University of Sydney, and the Australian director of Laing O'Rourke's future engineering and innovation consultancy, the Engineering Excellence Group. Laing O'Rourke is Australia's largest private engineering and construction business, with a local turnover in the order of \$AU3 billion per annum.

Andrew received his PhD from the University of Cambridge in 2002, and has spent the bulk of his career at the interface between industry and academia. He is a Chartered Engineer, and a Fellow of both the Institution of Chemical Engineers (IChemE) and Engineers Australia (IEAust).

Andrew is also a non-executive Director of Hazer Group (ASX:HZR), a listed clean tech organisation. In 2016, he was recognised as one of Australia's 50 most innovative engineers by Engineers Australia.



Trent Lund

Trent is the Lead Partner for Innovation & Digital Ventures at PwC Australia. He helps organisations leverage emerging technologies to transform ideas into customer-centred, commercial outcomes. With two decades of industry knowledge, Trent has worked around the globe—throughout the Asia-Pacific region, the United Kingdom and the Middle East.

Rob Wood

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



Peter Katsos

Peter is the General Manager for ABB Robotics Australia. ABB is a global leader in power and automation technologies. He has worked in industrial automation and robotics for over 25 years and has gained a wide range of experience ranging from service, design installation and project management of turnkey systems through to sales, business development and management. He holds a Bachelor's Degree in Electrical Engineering and Computing.

Research Committee

At the end of 2017 we broadened membership of the Centre's Research Committee to include our Executive Committee, all Chief Investigators and our project leaders. The Committee meets monthly via videoconference, face-to-face each quarter, to review research progress and make decisions on the Centre's research direction. Each year, the Research Committee conducts a review of the Centre's strategic research direction, making adjustments as required to ensure it stays on track to meeting its many objectives.

Governance KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017
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 Board	0	1	1	1	2	1	1
Vision and usefulness of the Centre Strategic Plan	At review	Annual Review by Board							
The adequacy of the Centre's performance measure targets	At review	Annual Review by Board							
Effectiveness of the Centre in bringing researchers together to form an interactive and effective research team	At review								
Weeks spent by Centre researchers at other nodes		10	1	45	51	45	31	50	102
Number of new joint research projects		0	0	1	1	2	2	3	3
Capacity building of the Centre through scale and outcomes	At review	<p>The Centre will develop our early career staff through supervisor mentoring, a Centre-run mentoring program, and opportunities for supervising PhD and undergraduate students, teaching and grant writing.</p> <p>We will develop PhD students through supervisor mentoring, a Centre-run mentoring program and opportunities for supervising undergraduate students.</p>							

SECTION 8

Finance and Operations




Image courtesy of QUT's Science and Engineering Faculty

The Australian Centre for Robotic Vision receives funding from two main sources, the Australian Research Council (ARC) and the Centre's collaborating organisations. In addition to cash, our collaborating and partner organisations provide significant resources through "in-kind" contributions, mainly in the form of our researchers' time.

We also source additional funding via industry engagement and relevant industry projects, raising over \$23 million over the life of the Centre so far.

Financial Performance

Our 2017 Income and Expenditure Report, see p. 93, summarises the Centre's financial performance for the calendar year.

The majority of our expenditure is on salaries and staff expenses, with smaller amounts allocated to equipment purchases, travel and professional development and operating expenses.

The Centre accumulated a significant carry-forward of funds in 2014 due to delays in establishing the legal agreements governing the operation of the Centre which prevented us from recruiting. Carry-forward continued into 2015 and 2016 but we were actively spending both income and carry forward in 2017. The carry-forward is fully allocated

in our budget and we project it to be fully expended with annual expenditure forecast to exceed "new" annual income for the remainder of the Centre's life to ensure funds will be fully expended. New income is approximately \$3.7million, in cash, that we are contracted to receive each year until 2021.

The Centre Executive determines budget allocations based on the Centre's original bid submission, with all nodes contributing to operating expenses. We have centralised some expenditure at QUT, with budgets across the life of the Centre developed by the mutual consent of all members of the Executive.

We have created a pool of untied cash to support strategic initiatives by using the (not guaranteed) indexation funds sent to the Administering Organisation by the ARC. These funds will be used in 2018 to support the Centre's new Gender Equity Plan by funding 50% of the costs of female postdoctoral research fellows. The initiative is designed to address some of the structural barriers preventing female participation in robotics and computer vision.

SUMMARY CONTRIBUTIONS FROM ALL CENTRE PARTNERS

Each year, our administrative and collaborative partner organisations contribute \$AU980,000 as cash, a total of more than \$AU6.86 million over the life of the Centre. Almost \$AU997,000 is contributed as in-kind, or \$AU6.98 million over the life of the Centre.

Together, our international partner organisations contribute \$AU139,000 per annum of in-kind, totalling \$AU973,000 over the Centre's seven-year lifespan.

The Centre's collaborating Partners, where most of our researchers are based, also provide access to a broad range of robotic vision equipment, which is conservatively valued at over \$AU1 million per annum (\$7 million in total, over the life of the Centre). A summary of the total cash and in-kind contributions from our Partners over the Centre's seven-year term is shown on p. 94. All figures quoted are in Australian dollars.

Each year, our administrative and collaborative partner organisations contribute



\$AU980,000

as cash,

...a total of more than

\$AU6.86 million

over the life of the Centre

Statement of Operating Income and Expenditure for year end 31 December 2017

Income	2014	2015	2016	2017
ARC Centre Grant distributed as follows	2,714,290	2,714,284	2,714,285	2,714,285
Monash University	116,929	116,929	116,929	116,929
University of Adelaide	244,400	244,400	244,400	244,400
Australian National University	230,000	230,000	230,000	230,000
Queensland University of Technology	360,000	350,000	350,000	350,000
ARC Indexation	82,027	132,110	180,499	223,919
Interest from QUT (commenced 2016)			6,359	64,175
Other Income to Administering Organisation (QUT)				220,471
Total Income	3,747,646	3,787,723	3,842,472	4,164,179

Expenditure	2014	2015	2016	2017
Purchased Equipment	42,256	51,829	57,325	1638
Shared Equipment/Facilities	-	815	-	
Travel and Professional Development	24,276	253,180	321,076	609,953
Maintenance (IT and lab)	8,350	19,639	71,769	109,516
Salaries/Personnel exp	373,061	2,337,787	2,813,795	3,138,116
Other	38,002	180,876	192,309	248,253
Total Expenditure	485,946	2,844,126	3,456,274	4,107,476
Surplus/Deficit	3,261,700	943,597	386,198	56,703
Previous year carry forward	-	3,261,700	4,205,297	4,591,495
Total carry forward surplus*	3,261,700	4,205,297	4,591,495	4,648,198

Finance KPIs

Performance Measure	Reporting Frequency	Target 2014	Outcome 2014	Target 2015	Outcome 2015	Target 2016	Outcome 2016	Target 2017	Outcome 2017
Annual cash contributions from administering and collaborating organisations	Annually								
ANU		\$230,000	\$230,000	\$230,000	\$230,000	\$230,000	\$230,000	\$230,000	\$230,000
Monash		\$116,400	\$116,929	\$116,400	\$116,929	\$116,400	\$116,929	\$116,400	\$116,929
QUT		\$350,000	\$360,000	\$350,000	\$350,000	\$350,000	\$350,000	\$350,000	\$350,000
Adelaide		\$244,400	\$244,400	\$244,400	\$244,400	\$244,400	\$244,400	\$244,400	\$244,400
Annual in-kind contributions from Administering collaborating organisations	Annually								
ANU		\$332,000	\$114,000	\$332,000	\$150,000	\$316,000	\$411,926	\$322,000	\$544,445
Monash		\$151,000	\$151,000	\$151,000	\$229,000	\$159,000	\$180,275	\$163,000	\$193,090
QUT		\$309,000	\$309,000	\$309,000	\$250,000	\$309,000	\$490,759	\$309,000	\$645,186
Adelaide		\$153,000	\$200,000	\$153,000	\$205,000	\$184,000	\$211,172	\$184,000	\$267,411
Annual cash contributions from Partner organisations	Annually	Nil					Nil		-
Annual in-kind contributions from Partner organisations	Annually								
Georgia Tech		\$18,000	\$3,600	\$18,000	\$18,000	\$18,000	\$18,000	\$18,000	\$0
Imperial College		\$18,100	\$3,600	\$18,100	\$18,100	\$18,100	\$18,100	\$18,100	\$18,100
INRIA		\$14,500	\$0	\$14,500	\$14,500	\$14,500	\$14,500	\$14,500	\$14,500
Data61/CSIRO		\$30,000	\$5,800	\$30,000	\$30,000	\$30,000	\$30,000	\$30,000	\$30,000
Oxford		\$32,000	\$3,200	\$32,000	\$32,000	\$32,000	\$32,000	\$32,000	\$16,000
Swiss Federal Institute		\$26,400	\$0	\$26,400	\$26,400	\$26,400	\$26,400	\$26,400	\$0
Other research income sourced by the Centre—End User (industry, public sector, ARC Linkage and Discovery in non-core areas, CRC)	Annually	\$0	\$1m	\$1m	\$4.1m	\$1.5m	\$8m	\$1.5m	\$9.5m
Number of new organisations collaborating with, or involved in, the Centre	Annually	2	2	2	2	2	4	2	13
Level and quality of infrastructure provided to the Centre	At review								

Outputs

2017 Publications

*denotes Core Centre Research Output

BOOK (1)

*Corke, P. I. (2017). *Robotics, Vision and Control: Fundamental Algorithms in MATLAB* (2nd ed.). Springer International Publishing.

JOURNAL ARTICLES (56)

*Adarve, J. D., & Mahony, R. (2017). Spherpix: A Data Structure for Spherical Image Processing. *IEEE Robotics and Automation Letters*, 2(2), 483–490. <http://doi.org/10.1109/LRA.2016.2645119>

Anderson, B. D. O., Shi, G., & Trumppf, J. (2017). Convergence and State Reconstruction of Time-Varying Multi-Agent Systems From Complete Observability Theory. *IEEE Transactions on Automatic Control*, 62(5), 2519–2523. <http://doi.org/10.1109/TAC.2016.2599274>

Andert, F., Ammann, N., Krause, S., Lorenz, S., Bratanov, D., & Mejias, L. (2017). Optical-Aided Aircraft Navigation using Decoupled Visual SLAM with Range Sensor Augmentation. *Journal of Intelligent & Robotic Systems*, 1–19. <http://doi.org/10.1007/s10846-016-0457-6>

Armin, M. A., Barnes, N., Alvarez, J., Li, H., Grimpen, F., & Salvado, O. (2017, September 14). Learning Camera Pose from Optical Colonoscopy Frames Through Deep Convolutional Neural Network (CNN). *Springer*. https://doi.org/10.1007/978-3-319-67543-5_5

Ball, D., Ross, P., English, A., Milani, P., Richards, D., Bate, A., Upcroft, B., Wyeth, G., Corke, P. (2017). Farm Workers of the Future: Vision-Based Robotics for Broad-Acre Agriculture. *IEEE Robotics & Automation Magazine*, 1–1. <http://doi.org/10.1109/MRA.2016.2616541>

*Bangura, M., & Mahony, R. (2017). Thrust Control for Multirotor Aerial Vehicles. *IEEE Transactions on Robotics*, 33(2), 390–405. <http://doi.org/10.1109/TR0.2016.2633562>

Bawden, O., Kulk, J., Russell, R., McCool, C., English, A., Dayoub, F., Lehnert, C., and Perez, T. (2017). Robot for weed species plant-specific management. *Journal of Field Robotics*, 34(6), 1179–1199. <http://doi.org/10.1002/rob.21727>

Bewley, A., & Upcroft, B. (2017). Background Appearance Modeling with Applications to Visual Object Detection in an Open-Pit Mine. *Journal of Field Robotics*, 34(1), 53–73. <http://doi.org/10.1002/rob.21667>

Bongiorno, D. L., Bryson, M., Bridge, T. C. L., Dansereau, D. G., & Williams, S. B. (2017). Coregistered Hyperspectral and Stereo Image Seafloor Mapping from an Autonomous Underwater Vehicle. *Journal of Field Robotics*. <http://doi.org/10.1002/rob.21713>

Bruce, J., Jacobson, A., & Milford, M. (2017). Look No Further: Adapting the Localization Sensory Window to the Temporal Characteristics of the Environment. *IEEE Robotics and Automation Letters*, 2(4), 2209–2216. <http://doi.org/10.1109/LRA.2017.2724146>

*Drory, A., Li, H., & Hartley, R. (2017). A learning-based markerless approach for full-body kinematics estimation in-natura from a single image. *Journal of Biomechanics*, 55, 1–10. <http://doi.org/10.1016/j.jbiomech.2017.01.028>

*Drory, A., Li, H., & Hartley, R. (2017). Estimating the projected frontal surface area of cyclists from images using a variational framework and statistical shape and appearance models. *Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology*. <https://doi.org/10.1177/1754337117705489>

*Drory, A., Zhu, G., Li, H., & Hartley, R. (2017). Automated detection and tracking of slalom paddlers from broadcast image sequences using cascade classifiers and discriminative

correlation filters. *Computer Vision and Image Understanding*, 159(June 2017), 116–127. <http://doi.org/10.1016/J.CVIU.2016.12.002>

Fan, C., Chen, Z., Jacobson, A., Hu, X., & Milford, M. (2017). Biologically-inspired visual place recognition with adaptive multiple scales. *Robotics and Autonomous Systems*, 96, 224–237. <http://doi.org/10.1016/J.ROBOT.2017.07.015>

*Fernando, B., & Gould, S. (2017). Discriminatively Learned Hierarchical Rank Pooling Networks. *International Journal of Computer Vision*, 124(3), 335–355. <http://doi.org/10.1007/s11263-017-1030-x>

*Fernando, B., Gavves, E., Oramas M., J. O., Ghodrati, A., & Tuytelaars, T. (2017). Rank Pooling for Action Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4), 773–787. <http://doi.org/10.1109/TPAMI.2016.2558148>

Hinas, A., Roberts, J., & Gonzalez, F. (2017). Vision-Based Target Finding and Inspection of a Ground Target Using a Multirotor UAV System. *Sensors*, 17(12), 2929. <http://doi.org/10.3390/s17122929>

Hoang, T., Do, T.-T., Tan, D.-K. Le, & Cheung, N.-M. (2017). Enhance Feature Discrimination for Unsupervised Hashing. Retrieved from <http://arxiv.org/abs/1704.01754>

*Ila, V., Polok, L., Solony, M., & Svoboda, P. (2017). SLAM++ - A highly efficient and temporally scalable incremental SLAM framework. *The International Journal of Robotics Research*, 36(2), 210–230. <http://doi.org/10.1177/0278364917691110>

Irons, J. L., Gradden, T., Zhang, A., He, X., Barnes, N., Scott, A. F., & McKone, E. (2017). Face identity recognition in simulated prosthetic vision is poorer than previously reported and can be improved by caricaturing. *Vision Research*, 137, 61–79. <https://doi.org/10.1016/j.visres.2017.06.002>

Jaiprakash, A., O'Callaghan, W. B., Whitehouse, S. L., Pandey, A., Wu, L., Roberts, J., & Crawford, R. W. (2017). Orthopaedic surgeon attitudes towards current limitations and the potential for robotic and technological innovation in arthroscopic surgery. *Journal of Orthopaedic Surgery*, 25(1), 230949901668499. <http://doi.org/10.1177/2309499016684993>

*James, J., Ford, J. J., & Molloy, T. L. (2017). Change Detection for Undermodelled Processes Using Mismatched Hidden Markov Model Test Filters. *IEEE Control Systems Letters*, 1(2), 238–243. <http://doi.org/10.1109/LCSYS.2017.2713825>

Kim, J., Cheng, J., Guivant, J., & Nieto, J. (2017). Compressed fusion of GNSS and inertial navigation with simultaneous localization and mapping. *IEEE Aerospace and Electronic Systems Magazine*, 32(8), 22–36. <https://doi.org/10.1109/MAES.2017.8071552>

*Kumar, S., Dai, Y., & Li, H. (2017). Spatio-temporal union of subspaces for multi-body non-rigid structure-from-motion. *Pattern Recognition*, 71, 428–443. <http://doi.org/10.1016/J.PATCOG.2017.05.014>

Lai, T., Wang, H., Yan, Y., Xiao, G., & Suter, D. (2017). Efficient guided hypothesis generation for multi-structure epipolar geometry estimation. *Computer Vision and Image Understanding*, 154, 152–165. <http://doi.org/10.1016/J.CVIU.2016.10.003>

Latif, Y., Huang, G., Leonard, J., & Neira, J. (2017). Sparse optimization for robust and efficient loop closing. *Robotics and Autonomous Systems*, 93, 13–26. <http://doi.org/10.1016/J.ROBOT.2017.03.016>

*Lehnert, C., English, A., McCool, C., Tow, A. W., & Perez, T. (2017). Autonomous Sweet Pepper Harvesting for Protected Cropping Systems. *IEEE Robotics and Automation Letters*, 2(2), 872–879. <http://doi.org/10.1109/LRA.2017.2655622>

Li, Z., Wu, L., Ren, H., & Yu, H. (2017). Kinematic comparison of surgical tendon-driven manipulators and concentric tube manipulators. *Mechanism and Machine Theory*, 107, 148–165. <http://doi.org/10.1016/j.mechmachtheory.2016.09.018>

*Lin, G., Shen, C., van den Hengel, A., & Reid, I. (2017). Exploring Context with Deep Structured models for Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1. <http://doi.org/10.1109/TPAMI.2017.2708714> *In Press

*Liu, F., Lin, G., Qiao, R., & Shen, C. (2017). Structured Learning of Tree Potentials in CRF for Image Segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, PP(99), 1–7. <http://doi.org/10.1109/TNNLS.2017.2690453> *In Press

*Liu, L., Li, H., Dai, Y., & Pan, Q. (2017). Robust and Efficient Relative Pose With a Multi-Camera System for Autonomous Driving in Highly Dynamic Environments. *IEEE Transactions on Intelligent Transportation Systems*, 1–13. <https://doi.org/10.1109/ITS.2017.2749409> *In Press

*Liu, L., Shen, C., & Hengel, A. van den. (2017). Cross-Convolutional-Layer Pooling for Image Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(11), 2305–2313. <http://doi.org/10.1109/TPAMI.2016.2637921>

Lu, H., Cao, Z., Xiao, Y., Zhuang, B., & Shen, C. (2017). TasselNet: counting maize tassels in the wild via local counts regression network. *Plant Methods*, 13(1), 79. <http://doi.org/10.1186/s13007-017-0224-0>

Ma, C., Yang, C.-Y., Yang, X., & Yang, M.-H. (2017). Learning a no-reference quality metric for single-image super-resolution. *Computer Vision and Image Understanding*, 158, 1–16. <http://doi.org/10.1016/J.CVIU.2016.12.009>

*Marmol, A., Peynot, T., Eriksson, A., Jaiprakash, A., Roberts, J., & Crawford, R. (2017). Evaluation of Keypoint Detectors and Descriptors in Arthroscopic Images for Feature-Based Matching Applications. *IEEE Robotics and Automation Letters*, 2(4), 2135–2142. <http://doi.org/10.1109/LRA.2017.2714150>

*McCool, C., Perez, T., & Upcroft, B. (2017). Mixtures of Lightweight Deep Convolutional Neural Networks: Applied to Agricultural Robotics. *IEEE Robotics and Automation Letters*, 2(3), 1344–1351. <http://doi.org/10.1109/LRA.2017.2667039>

*McFadyen, A., Jabeur, M., & Corke, P. (2017). Image-Based Visual Servoing With Unknown Point Feature Correspondence. *IEEE Robotics and Automation Letters*, 2(2), 601–607. <http://doi.org/10.1109/LRA.2016.2645886>

Molloy, Timothy L., Ford, Jason J., & Mejias, L. (2017). Detection of Aircraft Below The Horizon for Vision-Based Detect And Avoid in Unmanned Aircraft Systems. *Journal of Field Robotics*. <http://doi.org/10.1002/rob.21719>

*Nascimento, J. C., & Carneiro, G. (2017). Deep Learning on Sparse Manifolds for Faster Object Segmentation. *IEEE Transactions on Image Processing*, 26(10), 4978–4990. <http://doi.org/10.1109/TIP.2017.2725582>

Nguyen, K., Fookes, C., Jillela, R., Sridharan, S., & Ross, A. (2017). Long Range Iris Recognition: A Survey. *Pattern Recognition*. <http://doi.org/10.1016/j.patcog.2017.05.021>

Oakden-Rayner, L., Carneiro, G., Bessen, T., Nascimento, J. C., Bradley, A. P., & Palmer, L. J. (2017). Precision Radiology: Predicting longevity using feature engineering and deep learning methods in a radiomics framework. *Scientific Reports*, 7(1), 1648. <http://doi.org/10.1038/s41598-017-01931-w>

Parra Bustos, A., & Chin, T.-J. (2017). Guaranteed Outlier Removal for Point Cloud Registration with Correspondences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. <http://doi.org/10.1109/TPAMI.2017.273482> *In Press

Petoe, M. A., McCarthy, C. D., Shivasani, M. N., Sinclair, N. C., Scott, A. F., Ayton, L. N., ... Blamey, P. J. (2017). Determining the Contribution of Retinotopic Discrimination to Localization Performance With a Suprachoroidal Retinal Prosthesis. *Investigative Ophthalmology & Visual Science*, 58(7), 3231. <https://doi.org/10.1167/iovs.16-21041>

- Ren, C. Y., Prisacariu, V. A., Kähler, O., Reid, I. D., Murray, D. W. (2017). Real-Time Tracking of Single and Multiple Objects from Depth-Colour Imagery Using 3D Signed Distance Functions. *International Journal of Computer Vision*, 124, 80–95. <https://doi.org/10.1007/s11263-016-0978-2>
- *Sa, I., Lehnert, C., English, A., McCool, C., Dayoub, F., Upcroft, B., & Perez, T. (2017). Peduncle Detection of Sweet Pepper for Autonomous Crop Harvesting—Combined Color and 3-D Information. *IEEE Robotics and Automation Letters*, 2(2), 765–772. <http://doi.org/10.1109/LRA.2017.2651952>
- *Saleh, F., Aliakbarian, M. S., Salzmman, M., Petersson, L., Alvarez, J. M., & Gould, S. (2017). Incorporating Network Built-in Priors in Weakly-supervised Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP(0), 1–1. <http://doi.org/10.1109/TPAMI.2017.2713785> *In Press
- Sandino, J., Wooler, A., & Gonzalez, F. (2017). Towards the Automatic Detection of Pre-Existing Termite Mounds through UAS and Hyperspectral Imagery. *Sensors*, 17(10), 2196. <http://doi.org/10.3390/s17102196>
- Stacey, G., & Mahony, R. (2017). The Role of Symmetry in Rigidity Analysis: A Tool for Network Localisation and Formation Control. *IEEE Transactions on Automatic Control*. Retrieved from <http://ieeexplore.ieee.org/document/8023875/> *In Press
- Stronks, H. C., Walker, J., Parker, D. J., & Barnes, N. (2017). Training Improves Vibrotactile Spatial Acuity and Intensity Discrimination on the Lower Back Using Coin Motors. *Artificial Organs*. <https://doi.org/10.1111/aor.12882>
- *Teney, D., Wu, Q., & van den Hengel, A. (2017). Visual Question Answering: A Tutorial. *IEEE Signal Processing Magazine*, 34(6), 63–75. <https://doi.org/10.1109/MSP.2017.2739826>
- Trumpf, J., & Trentelman, H. L. (2017). A converse to the deterministic separation principle. *Systems & Control Letters*, 101, 2–9. <http://doi.org/10.1016/j.sysconle.2016.02.021>
- *Tsai, D., Dansereau, D. G., Peynot, T., & Corke, P. (2017). Image-Based Visual Servoing With Light Field Cameras. *IEEE Robotics and Automation Letters*, 2(2), 912–919. <http://doi.org/10.1109/LRA.2017.2654544>
- Villa, T. F., Jayaratne, E. R., Gonzalez, L. F., & Morawska, L. (2017). Determination of the vertical profile of particle number concentration adjacent to a motorway using an unmanned aerial vehicle. *Environmental Pollution*, 230, 134–142. <http://doi.org/10.1016/j.envpol.2017.06.033>
- Wu, L., & Ren, H. (2017). Finding the Kinematic Base Frame of a Robot by Hand-Eye Calibration Using 3D Position Data. *IEEE Transactions on Automation Science and Engineering*, 14(1), 314–324. <http://doi.org/10.1109/TASE.2016.2517674>
- *Zhang, Q., & Chin, T.-J. (2017). Coresets for Triangulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1. <http://doi.org/10.1109/TPAMI.2017.2750672> *In Press
- Zulqarnain Gilani, S., Mian, A., Shafait, F., & Reid, I. (2017). Dense 3D Face Correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1. <http://doi.org/10.1109/TPAMI.2017.2725279> *In Press
- ## CONFERENCE PAPERS (56)
- Abbasnejad, M. E., Dick, A., & Hengel, A. van den. (2017). Infinite Variational Autoencoder for Semi-Supervised Learning. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 781–790). Honolulu, USA: IEEE. <http://doi.org/10.1109/CVPR.2017.90>
- Bista, S. R., Giordano, P. R., & Chaumette, F. (2017). Combining Line Segments and Points for Appearance-based Indoor Navigation by Image Based Visual Servoing. In IROS 2017 - IEEE/RSJ International Conference on Intelligent Robots and Systems (pp. 2960–2967). Vancouver. Retrieved from <https://hal.inria.fr/hal-01572353/>
- Bratanov, D., Mejias, L., & Ford, J. J. (2017). A vision-based sense-and-avoid system tested on a ScanEagle UAV. *International Conference on Unmanned Aerial Systems (ICUAS) 2017*. Retrieved from <https://eprints.qut.edu.au/108459/>
- *Campbell, D., Petersson, L., Kneip, L., & Li, H. (2017). Globally-Optimal Inlier Set Maximisation for Simultaneous Camera Pose and Feature Correspondence. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 1–10). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.10>
- *Chen, Z., Jacobson, A., Sünderhauf, N., Upcroft, B., Liu, L., Shen, C., Reid, I., & Milford, M. (2017). Deep learning features at scale for visual place recognition. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 3223–3230). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989366>
- *Cherian, A., Koniusz, P., & Gould, S. (2017). Higher-Order Pooling of CNN Features via Kernel Linearization for Action Recognition. In 2017 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 130–138). Santa Rosa, CA: IEEE. <http://doi.org/10.1109/WACV.2017.22>
- *Cherian, A., Fernando, B., Harandi, M., & Gould, S. (2017). Generalized Rank Pooling for Activity Recognition. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 1581–1590). Honolulu, HI, USA: IEEE. <http://doi.org/10.1109/CVPR.2017.172>
- Chirikjian, G. S., Mahony, R., Ruan, S., & Trumpf, J. (2017). Pose Changes From a Different Point of View. In ASME 2017 41st Mechanisms and Robotics Conference. Cleveland, Ohio: ASME. <http://doi.org/10.1115/DETC2017-67725>
- *Dharmasiri, T., Spek, A., & Drummond, T. (2017). Joint prediction of depths, normals and surface curvature from RGB images using CNNs. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 1505–1512). Vancouver, Canada: IEEE. <http://doi.org/10.1109/IROS.2017.8205954>
- Fernando, T., Denman, S., Sridharan, S., & Fookes, C. (2017). Going deeper: Autonomous steering with neural memory networks. In IEEE Conference on Computer Vision and Pattern Recognition. Hawaii. Retrieved from <https://eprints.qut.edu.au/114117/>
- *Garg, S., Jacobson, A., Kumar, S., & Milford, M. (2017). Improving Condition- and Environment-Invariant Place Recognition with Semantic Place Categorization. Retrieved from <http://arxiv.org/abs/1706.07144>
- Gong, D., Tan, M., Zhang, Y., Hengel, A. van den, & Shi, Q. (2017). Self-Paced Kernel Estimation for Robust Blind Image Deblurring. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 1670–1679). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.184>
- Gonzalez, F., & Johnson, S. (2017). Standard operating procedures for UAV or drone based monitoring of wildlife. In Proceedings of Unmanned Aircraft Systems for Remote Sensing (UAS4RS 2017). Hobart, Tasmania. Retrieved from <https://eprints.qut.edu.au/108859/>
- Hall, D., Dayoub, F., Perez, T., & McCool, C. (2017). A Transplantable System for Weed Classification by Agricultural Robotics. Retrieved from http://www.ferasdayoub.com/wp-content/uploads/2014/12/IROS17_0494_FI.pdf
- *Harandi, M., Salzmman, M., & Hartley, R. (2017). Joint Dimensionality Reduction and Metric Learning: A Geometric Take. In Proceedings of the 34th International Conference on Machine Learning (ICML). Sydney, Australia.
- *Harwood, B., G. V. K. B., Carneiro, G., Reid, I., & Drummond, T. (2017). Smart Mining for Deep Metric Learning. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 2840–2848). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.307>
- *Henein, M., Abello, M., Ila, V., & Mahony, R. (2017). Exploring The Effect of Meta-Structural Information on the Global Consistency of SLAM. Retrieved from <http://viorelaila.net/wp-content/uploads/2016/06/Henein17iros.compressed.pdf>
- *Hua, M.-D., Hamel, T., Mahony, R., & Allibert, G. (2017). Explicit Complementary Observer Design on Special Linear Group SL(3) for Homography Estimation using Conic Correspondences. In 56th IEEE Conference on Decision and Control (CDC). Melbourne, Australia. Retrieved from <https://hal.archives-ouvertes.fr/hal-01628177/>
- *Hua, M.-D., Hamel, T., Mahony, R., & Allibert, G. (2017). SL(3) for Homography Estimation using Conic Correspondences. Retrieved from <https://hal.archives-ouvertes.fr/hal-01628177/document>
- *Hua, M.-D., Trumpf, J., Hamel, T., Mahony, R., & Morin, P. (2017). Point and line feature-based observer design on SL(3) for Homography estimation and its application to image stabilization. In International Conference on Robotics and Automation (ICRA). Singapore. Retrieved from <https://hal.archives-ouvertes.fr/hal-01628175/>
- *Ila, Viorela, Polok, Lukas, Solony, Marek, Istenic, K. (2017). Fast Incremental Bundle Adjustment with Covariance Recovery. In International Conference on 3D Vision (3DV).
- *Istenic, K., Ila, V., Polok, L., Gracias, N., & Garcia, R. (2017). Mission-time 3D reconstruction with quality estimation. In OCEANS 2017 - Aberdeen (pp. 1–9). Aberdeen, UK: IEEE. <http://doi.org/10.1109/OCEANSE.2017.8084708>
- *James, J., Ford, J. J., & Molloy, T. L. (2017). Quickest detection of intermittent signals with estimated anomaly times. In 2017 Asian Control Conference - ASCC 2017. Gold Coast. Retrieved from <https://eprints.qut.edu.au/112160/>
- *Ji, P., Li, H., Dai, Y., & Reid, I. (2017). “Maximizing Rigidity” Revisited: A Convex Programming Approach for Generic 3D Shape Reconstruction from Multiple Perspective Views. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 929–937). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.106>
- *Ji, P., Zhang, T., Li, H., & Salzmman EPFL -CVLab Ian Reid, M. (2017). Deep Subspace Clustering Networks. In 31st Conference on Neural Information Processing Systems (NIPS 2017). Long Beach, CA, USA. Retrieved from <http://papers.nips.cc/paper/6608-deep-subspace-clustering-networks.pdf>
- *Johnston, A., Garg, R., Carneiro, G., Reid, I., & van den Hengel, A. (2017). Scaling CNNs for High Resolution Volumetric Reconstruction From a Single Image. In IEEE International Conference of Computer Vision (ICCV) (pp. 939–948). Venice. Retrieved from http://openaccess.thecvf.com/content_ICCV_2017_workshops/w17/html/Johnston_Scaling_CNNs_for_ICCV_2017_paper.html
- *Khosravian, A., Chin, T.-J., Reid, I., & Mahony, R. (2017). A discrete-time attitude observer on SO(3) for vision and GPS fusion. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5688–5695). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989669>
- *Kiani, K. A., & Drummond, T. (2017). Solving Robust Regularization Problems Using Iteratively Re-weighted Least Squares. In 2017 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 483–492). Santa Rosa, CA: IEEE. <http://doi.org/10.1109/WACV.2017.60>
- *Kim, J.-H., Latif, Y., & Reid, I. (2017). RRD-SLAM: Radial-distorted rolling-shutter direct SLAM. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5148–5154). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989602>
- *Kumar, S., Dai, Y., & Li, H. (2017). Monocular Dense 3D Reconstruction of a Complex Dynamic Scene from Two Perspective Frames. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 4659–4667). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.498>
- *Le, H., Chin, T.-J., & Suter, D. (2017). An Exact Penalty Method for Locally Convergent Maximum

- Consensus. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 379–387). Honolulu, USA: IEEE. <http://doi.org/10.1109/CVPR.2017.48>
- *Le, H., Chin, T.-J., & Suter, D. (2017). RATSAC - Random Tree Sampling for Maximum Consensus Estimation. In 2017 International Conference on Digital Image Computing: Techniques and Applications (DICTA) (pp. 1–8). Sydney, Australia: IEEE. <http://doi.org/10.1109/DICTA.2017.8227480>
- *Leitner, J., Tow, A. W., Sünderhauf, N., Dean, J. E., Durham, J. W., Cooper, M., Eich, M., Lehnert, C., Mangels, R., McCool, C., Kujala, Peter T., Nicholson, L., Pham, T., Sergeant, J., Wu, L., Zhang, F., Uproft, B., and Corke, P. (2017). The ACRV picking benchmark: A robotic shelf picking benchmark to foster reproducible research. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 4705–4712). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989545>
- *Li, H., Wang, P., & Shen, C. (2017). Towards End-to-End Text Spotting with Convolutional Recurrent Neural Networks. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 5248–5256). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.560>
- *Liu, L., Li, H., & Dai, Y. (2017). Efficient Global 2D-3D Matching for Camera Localization in a Large-Scale 3D Map. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 2391–2400). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.260>
- Lu, H., Zhang, L., Cao, Z., Wei, W., Xian, K., Shen, C., & Hengel, A. van den. (2017). When Unsupervised Domain Adaptation Meets Tensor Representations. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 599–608). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.72>
- *Maicas, G., Carneiro, G., & Bradley, A. P. (2017). Globally optimal breast mass segmentation from DCE-MRI using deep semantic segmentation as shape prior. In 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017) (pp. 305–309). Melbourne, Australia: IEEE. <http://doi.org/10.1109/ISBI.2017.7950525>
- *Menikdiwela, M., Li, H., Nguyen, C., Shaw, S., (2017). CNN-based small object detection and visualization with feature activation mapping, IVCNZ2017, Canterbury, New Zealand
- *Meyer, B. J., & Drummond, T. (2017). Improved semantic segmentation for robotic applications with hierarchical conditional random fields. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5258–5265). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989617>
- *Nguyen, C. V., Milford, M., & Mahony, R. (2017). 3D tracking of water hazards with polarized stereo cameras. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5251–5257). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989616>
- *Rezatofighi, S. H., G. V. K. B., Milan, A., Abbasnejad, E., Dick, A., & Reid, I. (2017). DeepSetNet: Predicting Sets with Deep Neural Networks. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 5257–5266). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.561>
- *Rezadegan, F., Shirazi, S., Uproft, B., & Milford, M. (2017). Action recognition: From static datasets to moving robots. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 3185–3191). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989361>
- *Sadegh Aliakbarian, M., Sadat Saleh, F., Salzmann, M., Fernando, B., Petersson, L., & Andersson, L. (2017). Encouraging LSTMs to Anticipate Actions Very Early. In The IEEE International Conference on Computer Vision (ICCV), 2017 (pp. 280–289). Venice, Italy. Retrieved from http://openaccess.thecvf.com/content_iccv_2017/html/Aliakbarian_Encouraging_LSTMs_to_ICCV_2017_paper.html
- *Song, Y., Ma, C., Gong, L., Zhang, J., Lau, R. W. H., & Yang, M.-H. (2017). CREST: Convolutional Residual Learning for Visual Tracking. In 2017 IEEE International Conference on Computer Vision (ICCV) (pp. 2574–2583). Venice, Italy: IEEE. <http://doi.org/10.1109/ICCV.2017.279>
- *Spek, A., Drummond, T. (2017) A compact parametric solution to depth sensor calibration. In 28th British Machine Vision Conference (BMVC). London: <https://bmv2017.london/proceedings/>
- *Spek, A., & Drummond, T. (2017). Joint pose and principal curvature refinement using quadrics. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 3968–3975). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989456>
- *Sünderhauf, N., Pham, T. T., Latif, Y., Milford, M., & Reid, I. (2017). Meaningful maps with object-oriented semantic mapping. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 5079–5085). Vancouver, BC, Canada: IEEE. <http://doi.org/10.1109/IROS.2017.8206392>
- *Wang, J., Cherian, A., & Porikli, F. (2017). Ordered Pooling of Optical Flow Sequences for Action Recognition. In 2017 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 168–176). Santa Rosa, CA: IEEE. <http://doi.org/10.1109/WACV.2017.26>
- *Wang, P., Liu, L., Shen, C., Huang, Z., Hengel, A. van den, & Shen, H. T. (2017). Multi-attention Network for One Shot Learning. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 6212–6220). Honolulu, HI, USA: IEEE. <http://doi.org/10.1109/CVPR.2017.658>
- *Weberruss, J., Kleeman, L., Boland, D., & Drummond, T. (2017). FPGA acceleration of multilevel ORB feature extraction for computer vision. In 2017 27th International Conference on Field Programmable Logic and Applications (FPL) (pp. 1–8). Ghent, Belgium: IEEE. <http://doi.org/10.23919/FPL.2017.8056856>
- *Weerasekera, C. S., Latif, Y., Garg, R., & Reid, I. (2017). Dense monocular reconstruction using surface normals. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 2524–2531). IEEE. <https://doi.org/10.1109/ICRA.2017.7989293>
- *Xiao, Z., Li, H., Zhou, D., Dai, Y., & Dai, B. (2017). Accurate extrinsic calibration between monocular camera and sparse 3D Lidar points without markers. In 2017 IEEE Intelligent Vehicles Symposium (IV) (pp. 424–429). Los Angeles, CA, USA: IEEE. <http://doi.org/10.1109/IVS.2017.7995755>
- *Yang, J., Ren, P., Zhang, D., Chen, D., Wen, F., Li, H., & Hua, G. (2017). Neural Aggregation Network for Video Face Recognition. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 5216–5225). Honolulu, HI, USA: IEEE. <http://doi.org/10.1109/CVPR.2017.554>
- *Zhou, Y., Kneip, L., & Li, H. (2017). Semi-dense visual odometry for RGB-D cameras using approximate nearest neighbour fields. In 2017 IEEE International Conference on Robotics and Automation (ICRA) (pp. 6261–6268). Singapore: IEEE. <http://doi.org/10.1109/ICRA.2017.7989742>
- *Zhuang, B., Liu, L., Shen, C., & Reid, I. (2017). Towards Context-Aware Interaction Recognition for Visual Relationship Detection. In The IEEE International Conference on Computer Vision (ICCV) (pp. 589–598). Venice, Italy. Retrieved from http://openaccess.thecvf.com/content_iccv_2017/html/Zhuang_Towards_Context-Aware_Interaction_ICCV_2017_paper.html
- *Zuo, Y., & Drummond, T. (2017). Fast Residual Features: Rapid Ensemble Learning for Semantic Segmentation. In Proceedings of the 1st Annual Conference on Robot Learning, in PMLR 78 (pp. 27–36). Retrieved from <http://proceedings.mlr.press/v78/zuo17a.html>
- SUBMITTED (69)**
- *Abbasnejad, M. E., Shi, Q., Abbasnejad, I., Hengel, A. van den, & Dick, A. (2017). Bayesian Conditional Generative Adversarial Networks. Retrieved from <http://arxiv.org/abs/1706.05477>
- *Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., & Zhang, L. (2017). Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. Retrieved from <http://arxiv.org/abs/1707.07998>
- *Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I., Gould, S., and Hengel, A. van den. (2017). Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments. Retrieved from <http://arxiv.org/abs/1711.07280>
- *Bruce, J., Sünderhauf, N., Mirowski, P., Hadsell, R., & Milford, M. (2017). One-Shot Reinforcement Learning for Robot Navigation with Interactive Replay. Retrieved from <http://arxiv.org/abs/1711.10137>
- Chen, Y., Shen, C., Wei, X.-S., Liu, L., & Yang, J. (2017). Adversarial Learning of Structure-Aware Fully Convolutional Networks for Landmark Localization. Retrieved from <http://arxiv.org/abs/1711.00253>
- Chen, Y., Tai, Y., Liu, X., Shen, C., & Yang, J. (2017). FSRNet: End-to-End Learning Face Super-Resolution with Facial Priors. Retrieved from <http://arxiv.org/abs/1711.10703>
- *Chen, Z., Jacobson, A., Sünderhauf, N., Uproft, B., Liu, L., Shen, C., Reid, I., Milford, M. (2017). Deep Learning Features at Scale for Visual Place Recognition. Retrieved from <http://arxiv.org/abs/1701.05105>
- *Cherian, A., & Gould, S. (2017). Second-order Temporal Pooling for Action Recognition. Retrieved from <http://arxiv.org/abs/1704.06925>
- *Cherian, A., Sra, S., & Hartley, R. (2017). Sequence Summarization Using Order-constrained Kernelized Feature Subspaces. Retrieved from <https://arxiv.org/pdf/1705.08583.pdf>
- Cherian, A., Stanitsas, P., Harandi, M., Morellas, V., & Papanikolopoulos, N. (2017). Learning Discriminative -divergence for Positive Definite Matrices (Extended Version) *. Retrieved from <https://arxiv.org/pdf/1708.01741.pdf>
- *Cruz, R. S., Fernando, B., Cherian, A., & Gould, S. (2017). DeepPermNet: Visual Permutation Learning. Retrieved from <https://arxiv.org/abs/1704.02729>
- *Dayoub, F., Sünderhauf, N., & Corke, P. (2017). Episode-Based Active Learning with Bayesian Neural Networks. Retrieved from <http://arxiv.org/abs/1703.07473>
- *Deng, R., Zhao, T., Shen, C., & Liu, S. (2017). Relative Depth Order Estimation Using Multi-scale Densely Connected Convolutional Networks. Retrieved from <http://arxiv.org/abs/1707.08063>
- *Do, T.-T., Nguyen, A., Reid, I., Caldwell, D. G., & Tsagarakis, N. G. (2017). AffordanceNet: An End-to-End Deep Learning Approach for Object Affordance Detection. Retrieved from <http://arxiv.org/abs/1709.07326>
- Do, T.-T., Tan, D.-K. Le, Hoang, T., & Cheung, N.-M. (2017). Compact Hash Code Learning with Binary Deep Neural Network. Retrieved from <http://arxiv.org/abs/1712.02956>
- Do, T.-T., Tan, D.-K. Le, Pham, T. T., & Cheung, N.-M. (2017). Simultaneous Feature Aggregating and Hashing for Large-scale Image Search. Retrieved from <http://arxiv.org/abs/1704.00860>
- *Eriksson, A., Olsson, C., Kahl, F., & Chin, T.-J. (2017). Rotation Averaging and Strong Duality. Retrieved from <https://arxiv.org/abs/1705.01362>
- *Gale, W., Carneiro, G., Oakden-Rayner, L., Palmer, L. J., & Bradley, A. P. (2017). Detecting hip fractures with radiologist-level performance using deep neural networks. Retrieved from <https://arxiv.org/abs/1712.09048>
- Guo, G., Wang, H., Shen, C., Yan, Y., & Liao, H.-Y. M. (2017). Automatic Image Cropping for Visual Aesthetic Enhancement Using Deep Neural Networks and Cascaded Regression. Retrieved from <https://arxiv.org/abs/1712.09048>
- Hall, D., Dayoub, F., Kulk, J., & McCool, C. (2017). Towards Unsupervised Weed Scouting for Agricultural Robotics. Retrieved from <http://arxiv.org/abs/1702.01247>

- *Han, T., Wang, J., Cherian, A., & Gould, S. (2017). Human Action Forecasting by Learning Task Grammars. Retrieved from <https://arxiv.org/pdf/1709.06391.pdf>
- Jacobson, A., Scheirer, W., & Milford, M. (2017). Deja vu: Scalable Place Recognition Using Mutually Supportive Feature Frequencies. Retrieved from <http://arxiv.org/abs/1707.06393>
- *Ji, P., Reid, I., Garg, R., Li, H., & Salzmann, M. (2017). Non-Linear Subspace Clustering with Learned Low-Rank Kernels. Retrieved from <http://arxiv.org/abs/1707.04974>
- *Khosravian, A., Chin, T.-J., & Reid, I. (2017). A Branch-and-Bound Algorithm for Checkerboard Extraction in Camera-Laser Calibration. Retrieved from <http://arxiv.org/abs/1704.00887>
- *Latif, Y., Garg, R., Milford, M., & Reid, I. (2017). Addressing Challenging Place Recognition Tasks using Generative Adversarial Networks. Retrieved from <http://arxiv.org/abs/1709.08810>
- *Leal-Taixé, L., Milan, A., Schindler, K., Cremers, D., Reid, I., & Roth, S. (2017). Tracking the Trackers: An Analysis of the State of the Art in Multiple Object Tracking. Retrieved from <http://arxiv.org/abs/1704.02781>
- *Li, H., Wang, P., & Shen, C. (2017). Towards End-to-End Car License Plates Detection and Recognition with Deep Neural Networks. Retrieved from <http://arxiv.org/abs/1709.08828>
- Liu, W., Chen, X., Shen, C., Yu, J., Wu, Q., & Yang, J. (2017). Robust Guided Image Filtering. Retrieved from <https://arxiv.org/pdf/1703.09379.pdf>
- Ma, C., Huang, J.-B., Yang, X., & Yang, M.-H. (2017). Adaptive Correlation Filters with Long-Term and Short-Term Memory for Object Tracking. Retrieved from <http://arxiv.org/abs/1707.02309>
- Ma, C., Huang, J.-B., Yang, X., & Yang, M.-H. (2017). Robust Visual Tracking via Hierarchical Convolutional Features. Retrieved from <https://arxiv.org/pdf/1707.03816.pdf>
- *Ma, C., Shen, C., Dick, A., & Van Den Hengel, A. (2017). Visual Question Answering with Memory-Augmented Networks. Retrieved from <https://arxiv.org/pdf/1707.04968.pdf>
- *McTaggart, M., Morrison, D., Tow, A. W., Smith, R., Kelly-Boxall, N., Milan, A., Pham, T., Zhuang, Z., Leitner, J., Reid, I., Corke, P., and Lehnert, C. (2017). Cartman: Cartesian Manipulator for Warehouse Automation in Cluttered Environments. Retrieved from <http://arxiv.org/abs/1710.00967>
- *Meyer, B. J., Harwood, B., & Drummond, T. (2017). Nearest Neighbour Radial Basis Function Solvers for Deep Neural Networks. Retrieved from <http://arxiv.org/abs/1705.09780>
- *Milan, A., Pham, T., Vijay, K., Morrison, D., Tow, A. W., Liu, L., Erskine, J., Grinover, R., Gurman, A., Hunn, T., Kelly-Boxall, K., Lee, D., McTaggart, M., Rallos, G., Razjigaev, A., Rowntree, T., Shen, T., Smith, R., Wade-McCue, S., Zhuang, Z., Lehnert, C., Lin, G., Reid, I., Corke, P., and Leitner, J. (2017). Semantic Segmentation from Limited Training Data. Retrieved from <https://arxiv.org/abs/1709.07665>
- *Miller, D., Nicholson, L., Dayoub, F., & Sünderhauf, N. (2017). Dropout Sampling for Robust Object Detection in Open-Set Conditions. Retrieved from <http://arxiv.org/abs/1710.06677>
- *Morrison, D., Tow, A. W., McTaggart, M., Smith, R., Kelly-Boxall, N., Wade-McCue, S., Erskine, J., Grinover, R., Gurman, A., Hunn, T., Lee, D., Milan, A., Pham, T., Rallos, G., Razjigaev, A., Rowntree, T., Kumar, V., Zhuang, Z., Lehnert, C., Reid, I., Corke, P., and Leitner, J. (2017). Cartman: The low-cost Cartesian Manipulator that won the Amazon Robotics Challenge. Retrieved from <https://arxiv.org/abs/1709.06283>
- *Nguyen, A., Do, T.-T., Caldwell, D. G., & Tsagarakis, N. G. (2017). Real-Time Pose Estimation for Event Cameras with Stacked Spatial LSTM Networks. Retrieved from <https://arxiv.org/pdf/1708.09011.pdf>
- Nguyen, H. Van, Chesser, M., Rezatofghi, S. H., & Ranasinghe, D. C. (2017). Real-Time Localization and Tracking of Multiple Radio-Tagged Animals with an Autonomous Aerial Vehicle System. Retrieved from <https://arxiv.org/pdf/1712.01491.pdf>
- *Pan, L., Dai, Y., Liu, M., & Porikli, F. (2017). Depth Map Completion by Jointly Exploiting Blurry Color Images and Sparse Depth Maps. Retrieved from <http://arxiv.org/abs/1711.09501>
- *Pham, T. T., Do, T.-T., Sünderhauf, N., & Reid, I. (2017). SceneCut: Joint Geometric and Object Segmentation for Indoor Scenes. Retrieved from <https://arxiv.org/pdf/1709.07158.pdf>
- *Qiao, R., Liu, L., Shen, C., & Van Den Hengel, A. (2017). Visually Aligned Word Embeddings for Improving Zero-shot Learning. Retrieved from <https://arxiv.org/pdf/1707.05427.pdf>
- *Rezatofghi, S. H., Milan, A., Shi, Q., Dick, A., & Reid, I. (2017). Joint Learning of Set Cardinality and State Distribution. Retrieved from <https://arxiv.org/pdf/1709.04093.pdf>
- *Rezadegan, F., Shirazi, S., & Davis, L. S. (2017). A Real-time Action Prediction Framework by Encoding Temporal Evolution. Retrieved from <https://arxiv.org/abs/1709.07894>
- *Saha, S. K., Fernando, B., Cuadros, J., Xiao, D., & Kanagasigam, Y. (2017). Deep Learning for Automated Quality Assessment of Color Fundus Images in Diabetic Retinopathy Screening. Retrieved from <http://arxiv.org/abs/1703.02511>
- *Shen, T., Lin, G., Liu, L., Shen, C., & Reid, I. (2017). Weakly Supervised Semantic Segmentation Based on Co-segmentation. Retrieved from <http://arxiv.org/abs/1705.09052>
- *Shen, T., Lin, G., Shen, C., & Reid, I. (2017). Learning Multi-level Region Consistency with Dense Multi-label Networks for Semantic Segmentation. Retrieved from <http://arxiv.org/abs/1701.07122>
- *Shigematsu, R., Feng, D., You, S., & Barnes, N. (2017). Learning RGB-D Salient Object Detection using background enclosure, depth contrast, and top-down features. Retrieved from <https://arxiv.org/pdf/1705.03607.pdf>
- *Spek, A., Li, W.H., Drummond, T., (2017). A Fast Method for Computing Principal Curvatures from Range Images. Retrieved from arXiv preprint [arXiv:1707.00385](https://arxiv.org/abs/1707.00385)
- *Sünderhauf, N., & Milford, M. (2017). Dual Quadratics for Object Detection Bounding Boxes as Landmark Representations in SLAM. Retrieved from <http://arxiv.org/abs/1708.00965>
- Tan, D.-K. Le, Do, T.-T., & Cheung, N.-M. (2017). Supervised Hashing with End-to-End Binary Deep Neural Network. Retrieved from <http://arxiv.org/abs/1711.08901>
- *Teney, D., Anderson, P., He, X., & Hengel, A. van den. (2017). Tips and Tricks for Visual Question Answering: Learnings from the 2017 Challenge. Retrieved from <http://arxiv.org/abs/1708.02711>
- *Tow, A., Sünderhauf, N., Shirazi, S., Milford, M., & Leitner, J. (2017). What Would You Do? Acting by Learning to Predict. Retrieved from <http://arxiv.org/abs/1703.02658>
- *Toyer, S., Cherian, A., Han, T., & Gould, S. (2017). Human Pose Forecasting via Deep Markov Models. Retrieved from <https://arxiv.org/pdf/1707.09240.pdf>
- *Wade-McCue, S., Kelly-Boxall, N., McTaggart, M., Morrison, D., Tow, A. W., Erskine, J., Grinover, R., Gurman, A., Hunn, T., Lee, D., Milan, A., Pham, T., Rallos, G., Razjigaev, A., Rowntree, T., Smith, R., Kumar, Vijay, Zhuang, Z., Lehnert, C., Reid, I., Corke, P., and Leitner, J. (2017). Design of a Multi-Modal End-Effector and Grasping System: How Integrated Design helped win the Amazon Robotics Challenge. Retrieved from <http://arxiv.org/abs/1710.01439>
- *Wang, J., Cherian, A., Porikli, F., & Gould, S. (2017). Action Representation Using Classifier Decision Boundaries. Retrieved from <http://arxiv.org/abs/1704.01716>
- *Wang, X., Sekercioglu, Y.A., Drummond, T., Fremont, V., Natalizio, E., Fantoni, I. (2017). Relative Pose Based Redundancy Removal: Collaborative RGB-D Data Transmission in Mobile Visual Sensor Networks. Retrieved from arXiv preprint [arXiv:1707.05978](https://arxiv.org/abs/1707.05978)
- *Wang, X., You, M., & Shen, C. (2017). Adversarial Generation of Training Examples for Vehicle License Plate Recognition. Retrieved from <http://arxiv.org/abs/1707.03124>
- Wei, X.-S., Zhang, C.-L., Li, Y., Xie, C.-W., Wu, J., Shen, C., & Zhou, Z.-H. (2017). Deep Descriptor Transforming for Image Co-Localization. Retrieved from <http://arxiv.org/abs/1705.02758>
- Wei, X.-S., Zhang, C.-L., Wu, J., Shen, C., & Zhou, Z.-H. (2017). Unsupervised Object Discovery and Co-Localization by Deep Descriptor Transforming. Retrieved from <https://arxiv.org/pdf/1707.06397.pdf>
- *Wu, Q., Wang, P., Shen, C., Reid, I., & Van Den Hengel, A. (2017). Are You Talking to Me? Reasoned Visual Dialog Generation through Adversarial Learning. Retrieved from <https://arxiv.org/pdf/1711.07613.pdf>
- *Wu, Z., Shen, C., & Van Den Hengel, A. (2017). Real-time Semantic Image Segmentation via Spatial Sparsity. Retrieved from <https://arxiv.org/pdf/1712.00213.pdf>
- Yu, L., Jacobson, A., & Milford, M. (2017). Rhythmic Representations: Learning Periodic Patterns for Scalable Place Recognition at a Sub-Linear Storage Cost. Retrieved from <http://arxiv.org/abs/1712.07315>
- *Zhang, F., Leitner, J., Milford, M., & Corke, P. (2017). Sim-to-real Transfer of Visuo-motor Policies for Reaching in Clutter: Domain Randomization and Adaptation with Modular Networks. Retrieved from <https://arxiv.org/abs/1709.05746>
- *Zhang, F., Leitner, J., Milford, M., & Corke, P. I. (2017). Tuning Modular Networks with Weighted Losses for Hand-Eye Coordination. Retrieved from <http://arxiv.org/abs/1705.05116>
- *Zhang, J., Dai, Y., Porikli, F., He, M., Multi-Scale Salient Object Detection with Pyramid Spatial Pooling, APSIPA ASC 2017
- *Zhang, J., Wu, Q., Shen, C., Zhang, J., Lu, J., & Van Den Hengel, A. (2017). Asking the Difficult Questions: Goal-Oriented Visual Question Generation via Intermediate Rewards. Retrieved from <https://arxiv.org/pdf/1711.07614.pdf>
- Zhang, L., Wei, W., Shi, Q., Shen, C., Hengel, A. van den, & Zhang, Y. (2017). Beyond Low Rank: A Data-Adaptive Tensor Completion Method. Retrieved from <http://arxiv.org/abs/1708.01008>
- *Zhong, Y., Dai, Y., Li, H., Self-Supervised Learning for Stereo Matching with Self-Improving Ability, September 2017, <https://arxiv.org/pdf/1709.00930.pdf>
- *Zhuang, B., Wu, Q., Shen, C., Reid, I., & Hengel, A. van den. (2017). Care about you: towards large-scale human-centric visual relationship detection. Retrieved from <http://arxiv.org/abs/1705.09899>

Glossary

3DV	International Conference on 3D Vision
ACRA	Australasian Conference on Robotics and Automation (run by Australian Robotics and Automation Association)
AI	Associate Investigator
ANU	Australian National University
APRS	Australian Pattern Recognition Society
ARAA	Australian Robotics and Automation Association
ARC	Australian Research Council
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)
DPhil	Doctor of Philosophy (Oxford abbreviation)
DST	Defence Science and Technology
EOI	Expression of Interest
EUAB	End-User Advisory Board
FPGA	Field-Programmable Gate Array
IBVS	Image-Based Visual Servo
ICIP	IEEE International Conference on Image Processing
ICRA	IEEE International Conference on Robotics and Automation
ICT	Information and Communications Technology
IEEE	Institute of Electrical and Electronics Engineering
IET	Institution of Engineering and Technology
IROS	International Conference on Intelligent Robots and Systems
ISWC	International Semantic Web Conference
KPIs	Key Performance Indicators
MOOC	Massive Open Online Course
MVG	Multi View Geometry
NRP	National Research Priority
PhD	Doctor of Philosophy (international abbreviation)
PI	Partner Investigator
QUT	Queensland University of Technology
RHD	Research Higher Degree
RF	Research Fellow
SLAM	Simultaneous Localisation and Mapping
SRPs	Science and Research Priorities
VOS	Vision Operating System

Key Terms

algorithm // 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 // Intelligent behaviour demonstrated in machines.

autonomous // Without human intervention.

Bayesian (Bayes) nets (networks) // 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.

haptic controls // Incorporate tactile sensors that measure the forces exerted by the user and mimic human actions to control a machine.

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

machine learning // A type of artificial intelligence that gives computers the ability to learn based on large amounts of training data and without needing to be explicitly programmed.

neural network // A computer system very loosely modelled on neurons and synaptic connections found in biological brains.

semantics // Automatically applying meaningful human terms like 'kitchen' or 'coffee cup' to places or objects in the robotic vision system's environment. Semantics are important to help robots understand their environment by recognising different features and labelling or classifying them.



servo // A system that uses negative feedback to automatically correct its error.

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 surrounds.

support vector machine (SVM) // Classifies data by finding the best hyperplane that separates all data points of one class from those of another class.

Meet our Robot Family

In 2017, the Centre created a set of 'trading cards' each featuring one of the robots that we have created. As well as an image of our robots in action, the cards include vital statistics and fun facts about each one.



The Story of our Logo

Our logo represents the reunification of robotics and computer vision. It symbolises how robots might see in the future and recognises the importance of vision in the evolution of life on Earth.



540 million years ago, during the critical time period known as the Cambrian, the sense of vision, with its advanced and complex neurological network, was at the center of the Darwinian struggle for survival. Vision was a principal driver of evolution, providing animals with a map of their external world and concurrently invoking self-awareness - the recognition that the "self" viewing the world was also separate from it.

Vision also allowed animals to recognise similar forms and to associate with them, producing the inherent survival advantages involved in being part of a group.

Eventually, after 540 million years, humans and the human eye evolved. Humans then developed the technology to capture images using cameras, which mimic the human eye.

As the purpose of the Australian Centre for Robotic Vision is to give robots the gift of sight, our logo incorporates the most important elements of the eye.

Our Centre sits at the aperture (or opening) that allows light into the eye.

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

The blue circles represent the iris and the pupil, which control the amount of light entering the eye's natural crystalline lens. This clear, flexible structure works like the lens in a camera, shortening and lengthening its width in order to focus light rays.

The red shape represents a cross-section through an eye and symbolises the retina, where light rays come to a focusing point.

Embedded in the retina are millions of light sensitive cells, responsible for capturing light rays and processing them into impulses that are sent to the optic nerve. In a robot's eye these are digital sensors.

Just as vision played a major role in the evolution of life on Earth it can also spark the intelligence required for robots to be able to understand their environment, to make decisions and to perform useful tasks in the complex, unstructured and dynamically changing environments in which we live and work. Just as the minds of animals developed around the need to support a sense of vision, so to will the capabilities of robots.

Robotic Vision Partner Organisations



Robotic Vision Industry Partners



Australian Centre for Robotic Vision Head Office Contact Details

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

Design by Studio 55

CARTMAN:
THE PICK AND PACK ROBOT

WINNER: 2017 AMAZON
ROBOTICS CHALLENGE

CARTMAN STANDS FOR
CARTESIAN MANIPULATOR

3D PRINTER INSPIRED
AND PARTLY 3D PRINTED,
CARTMAN USES WEIGHING
SCALES AND TWO CAMERAS
TO IDENTIFY OBJECTS

THE AMAZON ROBOTICS CHALLENGE
COMBINED OBJECT RECOGNITION,
POSE RECOGNITION, GRASP
PLANNING, COMPLIANT MANIPULATION,
MOTION PLANNING, TASK PLANNING,
TASK EXECUTION, AND ERROR
DETECTION AND RECOVERY

CARTMAN CAN MOVE ALONG
THREE AXES AT RIGHT ANGLES
TO EACH OTHER, AND FEATURES
A ROTATING GRIPPER THAT
ALLOWS THE ROBOT TO PICK UP
ITEMS USING EITHER SUCTION
OR A SIMPLE TWO-FINGER GRIP

COST \$30K, WEIGHT 250KG.
CARTMAN WAS PACKED
AND CARRIED AS CHECKED
LUGGAGE TO JAPAN

CARTMAN BROKE A WRIST
DURING THE INTERNATIONAL
AMAZON ROBOTICS CHALLENGE
IN 2017 BUT RECOVERED AND
WENT ON TO WIN THE FINALS

