Part of the conference series
Breakthrough science and technologies
Transforming our future

Robotics and autonomous systems – visions, challenges and actions

Conference report

THE
ROYAL
SOCIETY

Introduction

On 13 November 2015, the Royal Society hosted a high-level conference on the subject of robotics and autonomous systems (RAS).

The conference brought together national and international scientists, technologists and leaders from across academia, industry and government, to scope the long-term vision (20 to 30 years) and challenges in this exciting and rapidly developing field. Presentations and discussion sought to define actions required to ensure maximum impact for society from RAS.

This conference is the second in a new series organised by the Royal Society, entitled *Breakthrough science and technologies: transforming our future*, which will address the major scientific and technical challenges of the next decade.

Each conference focuses on one technology and covers key issues including the current state of the UK industry sector, future direction of research and the wider social and economic implications. *The Transforming our future* conference series is organised through the Royal Society's Science and Industry programme, which aims to reintegrate science and industry across the Society's activities and to promote the value and importance of science by connecting academia, industry and government.

This report is not a verbatim record, but summarises the discussions that took place during the day and the key points raised. Comments and recommendations reflect the views and opinions of the speakers and not necessarily that of the Royal Society. Full versions of the presentations can be found on our website.

royalsociety.org/topics-policy/industry-innovation/transforming-our-future/

Executive summary

This conference set out to uncover the barriers to successful innovation and translation in robotics and autonomous systems (RAS), and to identify the fundamental breakthroughs and skills required to strengthen and build links across industry and the science base.

A report by McKinsey estimated that the application of advanced robotics could generate a potential impact of $1.9 – 6.4 trillion per year by 2025.[1] Recognising its potential, in 2012 the UK Government identified RAS as one of the Eight Great Technologies supporting the UK industrial strategy, which have the potential to rebalance the nation's economy, create growth and jobs.

As RAS are one of the most important emergent key enabling technologies, it represents an important area for scientific and engineering endeavour, now and in the future. RAS will be of immense socioeconomic impact, pervading all areas of society including communication, agriculture, medicine and transport. There is great potential for industrial advances, including the creation of new start-up companies, providing economic opportunities for the UK and elsewhere. Realising the potential for these new technologies requires transformational science and engineering closely allied with the needs of industry and society.

While robotics has come a long way, developing the capabilities of robots that operate successfully in unpredictable environments is still a considerable challenge. There is a need for multi-sensory inputs, from machine-learning driven computer vision to haptic intelligence, which surpass current capabilities. Real innovation is expected to come from shared control between robot and human, with "human-in-the-loop" actuation and motion mapping enabling machines to learn how to predict and adapt safely.

Humans will benefit from RAS capable of carrying out difficult and dangerous tasks beyond our own capabilities, working at the depths of the ocean, in nuclear power plants, or carrying out delicate surgeries in remote locations. However, conventional training and education will not be sufficient for industries using RAS technology – curricula do not prepare students adequately enough to manage the issues created by complex systems.

As RAS are taken from the research base and into industry, the transition from low to high technology readiness levels requires better access to more tangible assets to provide realistic environments in which to demonstrate, test and de-risk. Whether it be underwater facilities to test marine robots or operating driverless cars on the road, such demonstrations help prove commercial value to industry and investors.

National Academies have an important role in the successful development and uptake of emerging technologies such as RAS by using the expertise from within their Fellowships to advise government and other stakeholder groups, including industry. By organising this conference, the Royal Society has strengthened its links to the UK's RAS community and will continue to engage with it in the future.

---

[1] McKinsey Global Institute: Disruptive technologies: Advances that will transform life, business, and the global economy (May 2013); http://www.mckinsey.com/insights/business_technology/disruptive_technologies (accessed February 2016)

The current state of the art in RAS

RAS are taking on a new role and purpose. They will be part of our response to national challenges: an ageing population, safer transport, efficient healthcare, productive manufacturing and secure energy.

The conference included expert representatives from several horizontal sectors, who provided updates on the current state of their industries, and described potentially disruptive future technological breakthroughs.

Transport

Mobile autonomy is a pervasive technology that is now relevant to ocean monitoring and exploration, waste disposal, entertainment, mining, haulage, farming and shopping. In fields such as transport, where casualty rates are a particular cause for concern, autonomy may play a role in improving safety statistics.

Mapping, planning and navigation must all come together in autonomous vehicles. Professor Paul Newman FREng (University of Oxford), explained that today's driverless cars use infrastructure-free ego-motion, the 3D motion estimation of a camera system in an environment with no reference to satellites, paired with 2D scanning lasers to model the environment. Environmental noise, such as air conditioning units and weather, are overcome by new temporal calibration algorithms that enable cars to build cheap, accurate local maps just by driving through a city. These local maps then act as a substrate for mobile autonomy.

FIGURE 1

Driverless cars require the combination of spatial memories with real-time information.

Image used with permission from Professor Paul Newman.

The world changes, so although local views provide spatial memory, they require real-time information to be useful (Figure 1). Driverless cars will operate in all weathers and, if driven by machine vision, mapping must become independent of all environmental factors. One solution is the accumulation of data using experience-based navigation. Mapping in darkness is solved by using streetlights for point of interest (POI) recognition. Experience diversity makes localisation robust.

Newman's driverless car system has localised over 1000 km and in 2017 he plans to run the system in 40 "pods" in Milton Keynes. He proposes that future advances will use hybrid range-image perception methods to find people and objects and predict their intent. Robots in isolation will get better with use but a fleet of cars will share and inherit all local data sets, creating global perception.

Marine robotics

Professor David Lane CBE FREng (Heriot-Watt University) believes progress in marine robotics over the past 30 years demonstrates that scientific advances in autonomy, control, sensor processing and human-operator interaction have successfully translated into profitable businesses.

Early applications of advanced marine robotics were in the inspection, repair and maintenance of subsea infrastructures for the oil and gas industry. Traditionally, divers carried out inspections, but the depths they could reach were limited. Now, tele-operated robots with human operators connected to automated inspection units by cables are able to dive beyond 6000 metres. However, these systems are both expensive and subject to the cable wrapping around itself or other objects. This presents a commercial opportunity for the development of an autonomous vehicle.

FIGURE 2

Testing the valve-turning capability of the Girona 500 AUV in a realistic underwater environment.

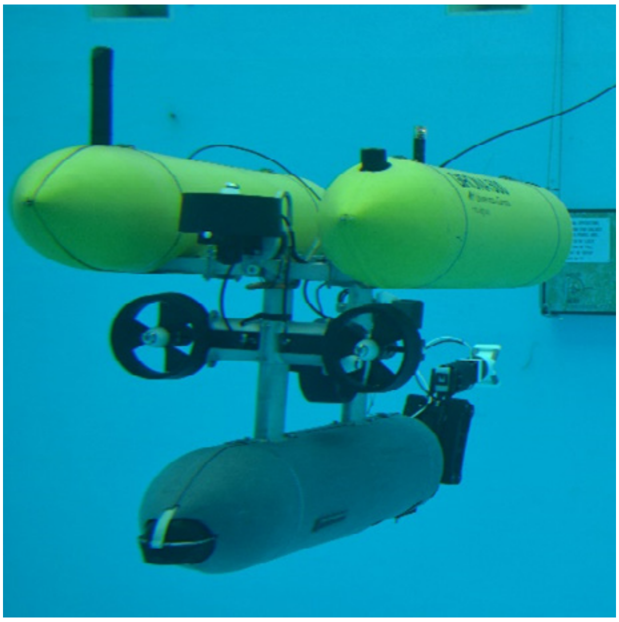

Image used with permission from Professor David Lane; taken from www.persistentautonomy.com.

The journey to exploit this commercial opportunity began in 2003 with the ALIVE programme, a multi-national, multi-disciplinary collaboration to develop an autonomous underwater vehicle (AUV) that can be deployed from a surface vessel, autonomously navigate to a subsea structure that has a docking panel on it, and turn a valve on that docking panel. While this system was able to complete this mission successfully, it was not robust.

By 2013, the Subsea 7 team had developed a completely autonomous inspection vehicle that is already undertaking commercial work for Shell. It uses a map of the subsea structure to navigate and has some decision-making capabilities to enable it to focus on POI. It is a simple, robust system capable of bringing back inspection data for analysis by a human operator.

The Subsea 7 Seabyte project developed a stable remote operating vehicle autopilot capable of maintaining its position and carrying out a clean 90-degree turn in a moving current. Difficult even for expert pilots, this simple skill represented immediate commercial benefit in the real world. Critically, this point may have been missed by developers without the benefit of *in situ* underwater demonstrations in a realistic environment and in collaboration with end users (see Figure 2).

> “Requirements are where the treasure is!”
>
> **Professor David Lane**

Lane’s latest project, Pandora (Persistently Autonomous Robots), uses probabilistic machine learning to teach robots. Rather than programming, robots learn to mimic human actions. The benefit of this probabilistic approach is an ability to offset disturbance and respond to failure. Using these techniques, the latest AUV can now manipulate a valve while maintaining stability underwater, inspect anchor chains and carry out other, relatively sophisticated commercial applications.

Robots in agriculture

Sustainable food production for the planet’s growing population is a significant challenge, and one where robotics, automation and vision technology can make a significant contribution. Professor Peter Corke (Queensland University of Technology) described recent work in the application of automation to broad-acre agricultural and horticultural production.

The growth in global food demand, in particular for animal proteins, has come at a time when farming is facing other challenges, including competition for agricultural land with other uses, e.g. housing, an ageing farming population, climate variation and a rise in cost of inputs. In Australia, for example, productivity started to drop in 2000, bringing lower yields from the same acreage.

New technologies, such as herbicides and mechanisation, have historically increased productivity. Ever larger machines offer maximum efficiency relative to the cost of human operation (i.e. one driver in a much larger tractor), but also create problems such as soil compaction. These challenges led to the introduction of controlled traffic farming methods, but this approach is unable to provide the step growth in productivity that is now required.

Corke and his team decided to remove the need for large, human-operated machines altogether, creating an automated system employing many small modular robots. Small, light vehicles do not crush the ground, can work non-stop and automatically resupply themselves.

However, to cover the same amount of ground requires a larger number of small autonomous robots. To deliver a value proposition that beats the current single large vehicle approach, the small, modular robots must be low-cost and ideally avoid using expensive real-time kinematic GPS guidance technology.

Corke came at this challenge from a new perspective. Farmers use their eyes, so why not employ robotic vision to agricultural tasks? Although robotic vision components are mass produced and cheap, Corke still had to develop a system with two capabilities – crop row following and collision detection and avoidance. To carry out either within a real, changing environment is challenging.

Using stabilised, low quality imaging inputs from a camera combined with inertial sensing, Corke's AGBOT II system warps this input into an overhead view, which it then uses to classify each crop row, during the day or night. To solve collision avoidance, it was necessary to combine height data using stereovision with algorithms to estimate novelty or surprise.

The AGBOT II is width adjustable to suit any crop types and can host multiple tools, such as robot arms to pick weeds or herbicide applicators. This technology has the potential to be highly economical, reducing the cost of weed removal by up to 90% by offering precise herbicide application.

The next generation AGBOT will incorporate weed classification to enable specific herbicide delivery or even use new technology, e.g. solar-powered microwaves, to replace herbicides altogether. Eventually, Corke envisions a future where agricultural robots give individual attention to each plant.

Thinking globally, Corke believes that robotics may have a place in increasing opportunities for farmers in poorer parts of the world. In countries such as China and India, the trend towards urbanisation will create a labour deficit on farms making new farming technologies even more essential.

Healthcare

Physics makes scaling forces difficult but inspiration can come from biomimetics. By applying information about how micro-organisms such as bacteria swim through the body, Professor Bradley Nelson (ETH-Zurich) has developed magnetic actuation methods capable of generating force and torque at the micro- and nanoscale, *in vitro* and in three dimensions.

Nelson's first prototype, the OctoMag, used eight actuators to provide five degrees of freedom and was capable of carrying out retinal surgery in animal trials, placing drug delivery inserts directly into the eye. This technology also offers more accurate steering of catheter tips used in cardiac ablation.

FIGURE 3

Aeon Phocus during its first use in humans (21 August 2015).

Image used with permission from Professor Bradley Nelson.

OctoMag has now evolved into the Aeon Phocus, which is used in human surgery. Its ability to accurately steer catheter tips allows surgeons to carry out tele-automated catheter controlled heart surgery in conscious patients (Figure 3).

New 3D technologies allow the nano-printing of more compatible, non-cytotoxic and biodegradable microstructures made using polymers. Swarms of microstructures have been demonstrated to move as fast as microns per second in mice. This affords a new design space, where nanoscale, functional end effectors can be fabricated at the size of red blood cells; *in vivo* microbots enabling cell transfections and gene expression. Other applications include water mediation, manufacturing and environmental monitoring. The potential is vast, but investment is essential as the journey from research to clinical trials in humans and certification is complex and expensive.

The future direction of research

> “Full autonomy is a 20th century rhetoric. The real frontier is collaboration, which includes autonomy but different levels of autonomy at different moments under the control of a human operator.”
>
> **Professor Sethu Vijayakumar quoting from Professor David Mindell (MIT)**

Shared autonomy

Professor Sethu Vijayakumar (University of Edinburgh) believes real innovation will come from what he calls shared autonomy or punctuated autonomy, which is essentially shared control between robot and human.

Robots that interact must be able to sense the world around them, interact safely, and learn to predict and adapt according to real-world circumstances, while still allowing significant human input where appropriate.

Highlighting cutting edge research from his group, Vijayakumar demonstrated how anthropomorphic robots realise servoing with moving cameras to achieve robust, real-time feedback to map the local environment, create a unified percept and plan ahead to achieve closed loop behaviour. Bridging representation models are needed to understand the difference between what a human sees and the information required by the robot to carry out tasks. Techniques pioneered by Vijayakumar’s group based on interaction mesh and writhe representations pave the way for efficient, real time adaptation of motion plans in robots that interact closely with humans in co-working scenarios.

FIGURE 4

Shared autonomy prosthetic hand enables controlled grasp pressure without visual input from the user, which was demonstrated by Professor Vijayakumar at the conference.

© Touch Bionics

Safe interactive systems must also apply compliant actuation to deal with uncertainty, using algorithms able to optimise control policies with safe, real-time human-in-the-loop reactive capabilities. Efficient implementation that incorporated multiple representations was achieved by the use of probabilistic inference based reformulation of the classical optimal control problem. These techniques are crucial for safe, precise interactions with wearable exoskeletons such as the HASy (DLR), BLUE and miniBLUE compliant bipeds from the University of Edinburgh as well as active pelvis and ankle orthosis being developed in Vijayakumar's group.

Human-in-the-loop actuation and motion mapping is a way for machines to learn how to predict and adapt, learning internal dynamics associated with tasks from example. Using this data-driven approach, machines can predict the value of every possible action, without pre-programming all options, delivering on the fly adaptation at multiple scales.

This concept of shared autonomy was exemplified in a live demonstration of a prosthetic hand which Vijayakumar has been involved with the development of, where the EMG sensors decoded the user's volition while all the control of grasp pressure and adaptation to object shape was devolved to the smarts embedded in the hand electronics – reducing end user fatigue (Figure 4).

Vijayakumar ended by expressing the need and benefits of tighter collaboration between academia, industry, government research and funding bodies. An exciting new project with NASA served as an excellent exemplar. The University of Edinburgh is hosting the NASA Valkyrie Humanoid platform, a multi-million dollar testbed for developing advanced whole body dexterous manipulation and navigation capabilities – which is a follow up of the highly publicised DARPA Robotics Challenge. This is driven by the futuristic vision of lowering cost of space exploration by pre-deployment of assets, primarily through the use of robotic unmanned missions.

Experiencing the virtual twins

The move to 3DEXPERIENCE twins, which will replace the entire physical being of large constructions, provides an entirely holistic approach to science through software.

FIGURE 5

3DEXPERIENCE twin of Singapore created using Dassault Systèmes' 3DEXPERIENCity.

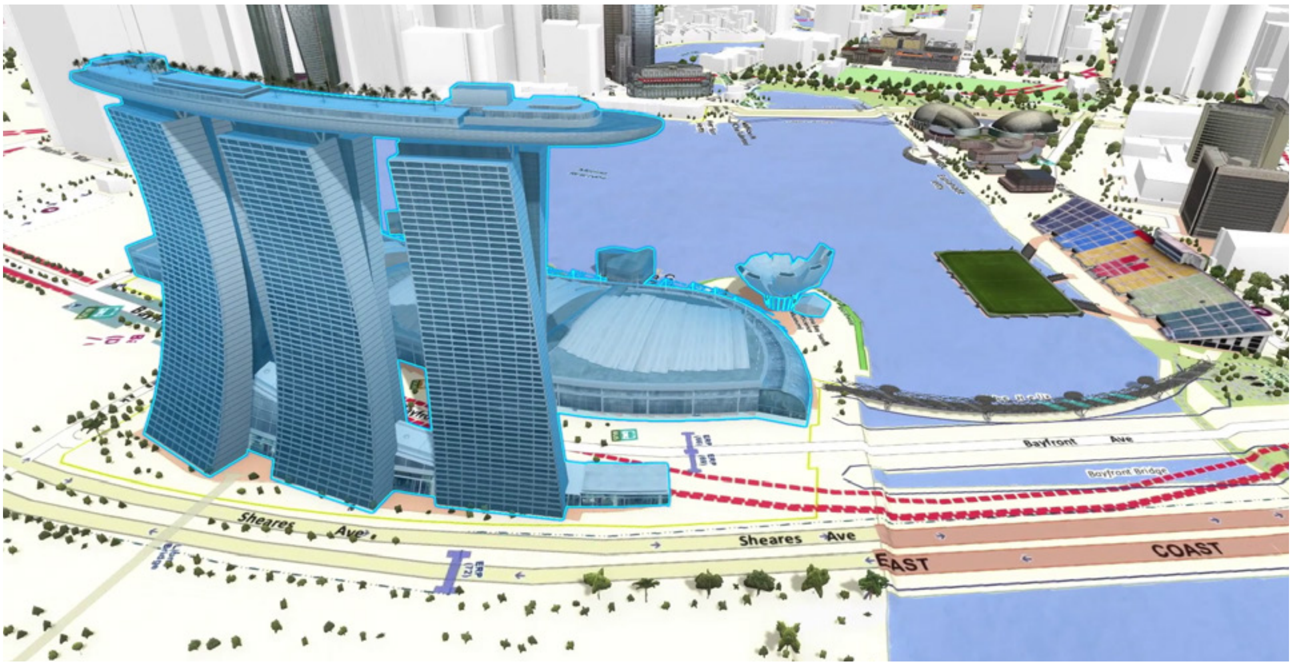

© National Research Foundation Singapore

According to Bernard Charlès, Dassault Systèmes, experiencing the virtual twins and innovation platforms will foster development of open, modular, system-driven and interconnected environments creating a true “Internet of Experiences.” This will not only entail revolution in design and a transformation in manufacturing and engineering, but will begin a transformation in society through the emergence of new ways of interacting, working and living together. In June 2015, Singapore announced it would use the Dassault Systèmes simulation platform to create a twin of the city. This digital model will provide a reference for automation systems to apply a multi-scale system approach to any enquiry, e.g. how far is an individual from a facility that might offer appropriate medical care or support (Figure 5)?

The nuclear disaster in Fukushima, where humans had to intervene to help manage the recovery efforts, brought to light a lag in technology for disaster recovery in many large nuclear facilities. Despite Japan being at the forefront of technology, the assumption that human intervention would be the simplest, cheapest and most effective solution in a nuclear disaster was a huge problem as in reality, no one understood all the complex drawings and documentation available. Both Russia and China have since come to recognise the benefits of experiencing the virtual twin of such facilities, providing safe and rapid evaluation of the situation prior to and during disaster and recovery.

Machine learning and computer vision

To be useful in the real world, robots must be able to sense the environment around them. Professor Andrew Blake FREng FRS (The Alan Turing Institute) discussed the two emerging paradigms of machine learning in computer vision: empirical detectors used to manage big data by learning at scale, and analysis by synthesis, based on building an “explanation” of the visual observations.

Big data is used to train closed ‘black-box’ empirical detectors that provide one-shot detection, with no prediction required, for efficient and accurate interpretation. The alternative is to construct data models using step-by-step explanation, allowing the software to learn from examples, and building a probabilistic distribution model from which to infer the possible solutions. Examples of the big data approach include the use of deep learning with the ImageNet large-scale hierarchical image database and the NIST Switchboard test platform for speech classification.

In contrast, analysis by synthesis turns the empirical approach on its head, synthesising an image, comparing it with the observed image, and driving refinements of the scene description from the differences. Application of Bayesian inference on structured probability distributions facilitates vision models that can deal with the complexity of natural images, using an approach inspired by ideas about processing in the human brain. Errors drive refinements and even in naïve visual models that ignore some of the relationships between pixels, a great deal of explanatory power is nonetheless possible.

Unfortunately, analysis by synthesis is subject to failures over long timescales and is not robust enough to be useful on its own for complex applications, whereas the black-box approach is very robust but cannot explain the data in detail. Blake believes the solution is to combine the two approaches to create detection models that use black-box type decision forests and then apply analysis by synthesis. For example, the Microsoft Kinect team took an object recognition approach, simplifying complex human poses by using per-pixel classification to produce a more efficient human pose estimation from single depth images (Figure 6). Microsoft also fused detection with analysis by synthesis to produce software to track the detailed motion of the hand, potentially of great value in virtual reality systems.

FIGURE 6

Real-time human pose recognition in parts from single depth image.

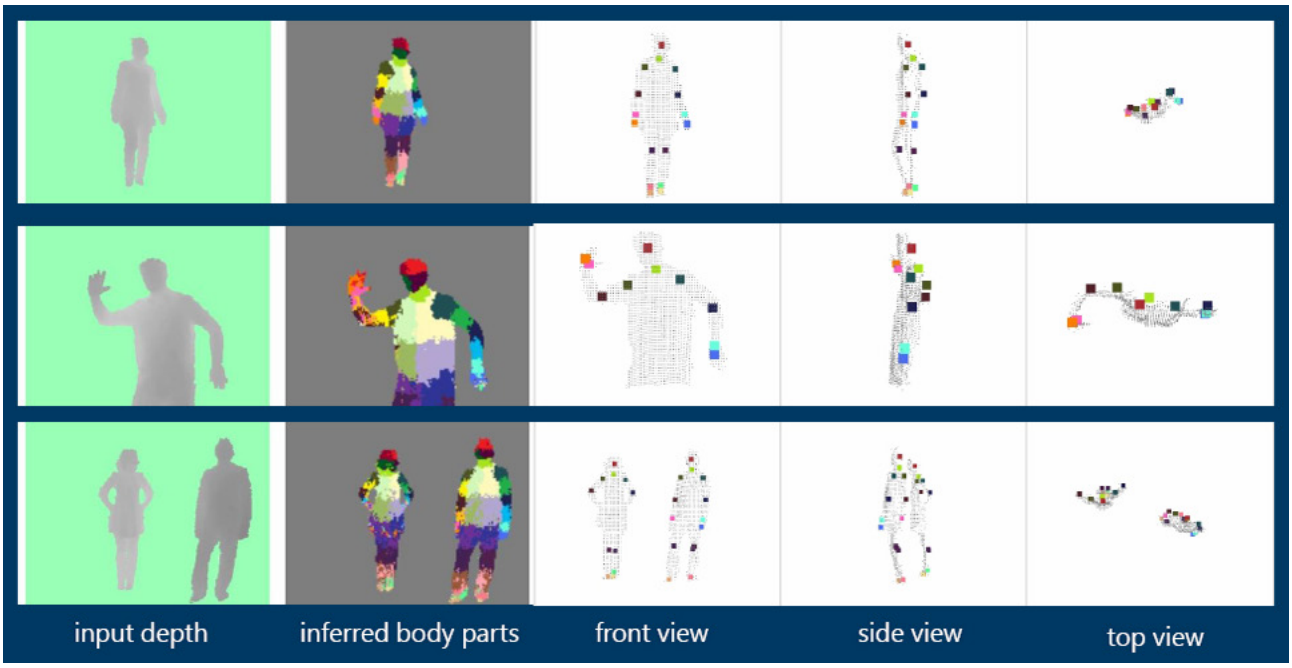

Image used with permission from Professor Andrew Blake.

Touch sensing and haptic feedback

Physical tasks rely heavily on visual and audio senses, but touch is also very important and less well understood. Professor Katherine J Kuchenbecker (University of Pennsylvania) described the necessity in going beyond vision and hearing to endow our robots with the sense of touch, thereby equipping them with haptic intelligence.

Kuchenbecker gave the example of robotic surgery systems, such as the da Vinci robot (Intuitive Surgical, Inc.), which provide 3D imagery and wristed dexterity but are not equipped with a sense of touch. Surgeons using these systems must learn to 'feel with their eyes', a difficult process resulting in potential dangers.

Current touch sensors commonly available are six-axis force sensors that can measure forces and torques, but these are delicate, expensive, and tend to drift in response to temperature changes. Developers must also mount them in series with the robotic tools, which is a big deterrent.

Kuchenbecker came up with a new approach to providing haptic feedback in robotic surgery, going back to the mechanoreceptors of the hand and evaluating which tactile cues provide the most information. This work suggested that a focus on vibration would be both technically viable and clinically beneficial.

FIGURE 7

The VerroTouch™ sensor and actuator kit for adding haptic feedback to surgical robots.

© Professor Katherine J Kuchenbecker

Over several iterations, Kuchenbecker and her team created VerroTouch™ (Figure 7), a set of cheap, simple retrofit components that can measure the instrument contact vibrations that occur during robotic surgery. Like a microphone and speaker, but using lower frequencies, the system enables clinicians to feel vibrations from the right and left instruments at their fingertips. Surgeons find this feedback to be natural and beneficial. Additional observations showed that the surgeon's skill level is often reflected in the vibration feedback – clumsiness is loud!

Kuchenbecker's findings about the utility of tactile vibrations in robotic teleoperation foreshadow a need to endow autonomous robots with a more capable sense of touch. However, such efforts are still in their infancy.

Tactile data is sensitive to environmental conditions, and improvements are needed in perception from touch. Small, low cost and reliable high bandwidth haptic sensors are essential to enable robots to handle delicate objects as capably as humans. Robots also need to learn to ignore the haptic effects of their own motion to better identify those sensations that are caused by contact with the world.

The challenges faced in turning research into commercial success

Throughout the conference, several big challenges to RAS became apparent in the next major steps towards developing ever more complex autonomy that can operate in real environments and interact with people.

> “The next major step towards everyday robots is all about the people.”
>
> **Professor Nick Roy**

Efficiency of learning

Blake believes there are shortfalls in the way computers currently learn compared to humans. This was illustrated by comparing ImageNet, which contains about 1.2 million training images, 50,000 validation images, and 100,000 test images used to train machine learning algorithms to make accurate decisions, to a child, which only requires a handful. Neural networks do not have the ability to apply feature inheritance or pattern logic to increase the efficiency of learning. There is a learning efficiency trade-off between tolerance to noise and the number of examples a system can read. Interestingly, it does not appear to be possible to move between the two. Blake believes that closing the gap between these two extremes will be important to faster machine learning in the future (Figure 8).

> “The bringing together of image processing and language is definitely the next big thing in vision technology,”
>
> **Professor Andrew Blake**

FIGURE 8

There appears to be a learning efficiency trade-off with no possibility of moving between the two.

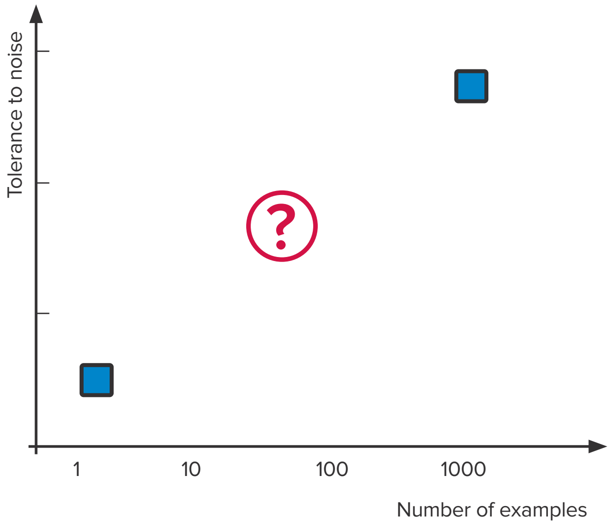

Image adapted with permission from Professor Andrew Blake.

Closing the semantic gap between human and robot operations

Professor Nick Roy (MIT) believes that closing the semantic gap between robot and human thinking is one of the most important challenges in RAS today.

Probability, miniaturisation of components and GPS have all helped drive RAS forward. They have enabled the development of robots that can get from point to point. Now, better on-board sensing and autonomy are crucial to build more complex systems capable of safe, reliable and trustworthy activity with greater autonomy. Uncrewed systems, such as drones, currently require an expensive team of pilots but in the future, systems will feature multiple vehicles controlled by one person.

Airborne drones use geometry to make decisions on how to get to a final destination. Current drone maps incorporate areas of free space, obstacles and unknowns (see Figure 9). When asking a human to describe the same journey, they are more likely to create a minimal sketch and describe the directions on how to get there using basic verbal instructions. Humans also know that if they get lost along the way, they can ask someone for further instructions. In contrast, much of the very precise, accurate and complex mapping carried out by today's drones is unnecessary. Ideally, drones should be able minimise the amount of data required to complete the same journey in the same way humans do.

Robots must resolve complex representational problems. For example, a drone required to deliver a parcel to an address cannot just arrive at GPS coordinates, which might be on top of the roof or on the side of the street, but must work out where to deliver the package. Roy believes the solution to this dilemma is through enabling communication between robot and human, bringing language to geometry.

Language has syntax, which controls how it is broken up, and can be expressed in corpora (samples) of "real world" text. Roy has demonstrated that robots can analyse syntax input from data-rich demonstration corpora to understand instructions delivered by a human operator and combine them with environmental parameters to compute the likelihood of correct pathways. Human and robot, sharing autonomy, can be highly efficient and instead of the human operator learning how to operate the robot, he or she simply uses normal instructions as they would when communicating with another human.

Abstract spatial concepts are pervasive in natural language. This means that corpora data can also teach robots to have more indirect percepts, enabling them to carry out more abstract instructions such as picking up the third object in a row, or arranging objects in a circle. Roy also believes that this approach might eventually enable robots to bridge the semantic gap using symbolic representations of the environment instead of language input.

FIGURE 9

Current drone mapping capabilities. (White – free space, Black – obstacles, Blue – unknown)

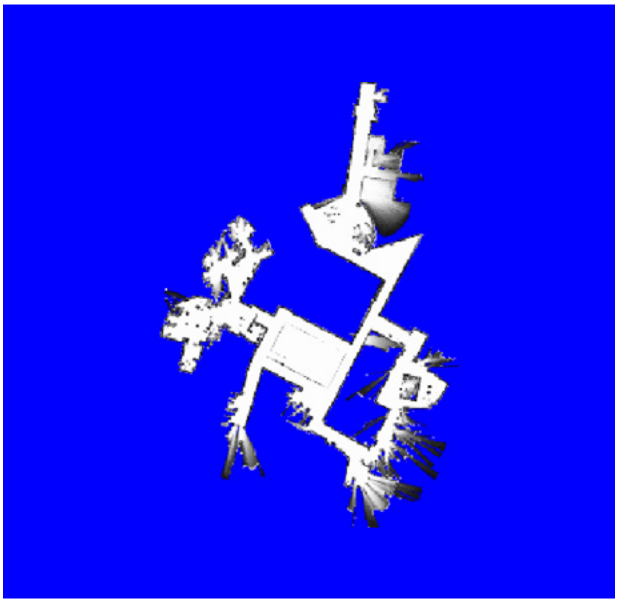

Image used with permission from Professor Nick Roy.

Legislation and regulation

> "Regulation works best when we understand what we're dealing with and we don't yet understand what we're dealing with."
>
> **Professor Nick Roy**

Legal and regulatory frameworks can struggle to keep up with technological advances. While it seems logical that accidents, such as crashed drones, might bear a shared responsibility between manufacturer and operator, the current lack of clarity in these matters raises an important engineering question regarding a developer's ability to demonstrate system capabilities. Roy observed the contradictions between certification and verification that are fundamentally about having knowledge set in stone, and autonomy, which inherently deals with the unpredictable.

Until recently, it was illegal for anyone other than government agencies to operate drones. While laws are changing slowly, Roy believes that it is important to have legal frameworks and regulatory guidelines in place that encourage development, without creating regulatory hurdles that are so high they might stifle innovation.

The skills base needed to deliver major scientific advances

There is a general expert consensus that it is important to open the doors to robotics for future generations. Studies and observations make it clear that young people growing up with advanced systems are able to understand and approach their use in a much more natural way. Young people recognise and understand virtual connection without physical connection – a child playing with an application-based remote controlled car will naturally demonstrate it to others by linking its virtual controls to the car's movements.

> "Computers were driven to today's excellence by lonely kids playing computer games. Perhaps robotics will be the same."
>
> **Professor Paul Newman**

Charlès argued that current education programmes continue to underestimate the value of systems approaches across all engineering disciplines. Application development involves a systems approach rather than a design approach; it is about connecting things together.

Holistic thinking is necessary for PhD students to define a research problem and formulate a question, but is still not well taught in schools.
Towards this end, Dassault Systèmes launched a STEM academic programme with DARPA and the Georgia Institute of Technology that introduced the 3DEXPERIENCE cloud-based curriculum for systems design and engineering into 200 universities across the United States.

Adopting a holistic approach that connects disciplines together to understand the challenges faced by our existing industries can create smarter, more sustainable systems. Europe has advantages through its high value in art and scientific inspiration, but needs to be aware that some of its traditional industries lag behind global competitors. Due to China's less risk-averse attitude, Dassault Systèmes have been able to introduce new systems for home design that offer digital continuity without human intervention and allow mass customisation at the cost of mass production.

> "The future of RAS is not to replace humans. The future of RAS is to reposition human action within the production system."
>
> **Bernard Charlès**

Charlès considered the much cited fear surrounding job losses caused by automation, arguing that it is important to look at the unfulfilled needs of our future society – the ageing population, future cities, healthcare, and in industries such as mining. His conviction is not to look at the economy as it is, but to look at the economy of tomorrow, which calls for simultaneous innovation in products and business models. There remains much to do in terms of technological development to bridge the gap between our future requirements and our current capabilities.

The position of the UK and how it can benefit from the technology

> “There is a big difference between invention and innovation. While invention takes money and turns it into ideas, innovation takes those ideas and turns them back into money.”
>
> **Professor David Lane**

Smart procurement

The 2014 UK RAS 2020 Strategy[2] outlines how the UK government, the research base, innovation agencies, companies and their customers can help the UK capture some of the value from its science base by employing smart procurement strategies, informed by lessons learned regarding the real needs of each horizontal industry sector.

From 2012 –2014, the Robotics and Autonomous Systems Special Interest Group (RAS-SIG), led by Professor Lane, investigated how the UK could make the most of its excellence in RAS. The resulting five-stranded vision, as laid out in the RAS 2020 strategy, requires a focus on coordination, assets, grand challenges, clusters and skills.

While there are opportunities in the private sector, there are also opportunities in the public sector, where government becomes the customer. In particular, Lane believes both providing assets and setting out challenges are important for the translation of research into commercial opportunity.

Assets and challenges

Both intangible and tangible assets are critical. Intangible assets, such as a favourable regulatory environment, will help make the UK an attractive location for large global businesses. Tangible assets are places where developers can test and de-risk new systems in real environments, building trust with customers and demonstrating what the technology can do. These assets can be anything from a hamlet of houses to test driverless cars, to farms, factories and specialist environments such as the Schiehallion deep-sea oil bed and Boulby Mines.

Finding ways to get users and technologists working together, e.g. through Grand Challenge competitions, is important to the future development and translation of RAS technologies. Using assets as hosting grounds, these competitions focus on real scenarios across vertical markets and bring together industry, academia, finance and innovators to stimulate competition, identify new possibilities, excite the public and support an innovation pipeline. Solutions can be developed in de-risked environments, abiding by regulations specific to each market to deliver genuinely useful capabilities.

In the US, for example, funded challenges such as the Defence Advanced Research Projects Agency (DARPA) Robotics Challenges, and annual competitions, such as AUV Fest, allow technology experts to work alongside users over a sustained period, generating a “requirement spiral.”

---

[2] Robotics and Autonomous Systems Special Interest Group (RAS-SIG): RAS2020 Robotics and Autonomous Systems (July 2014). https://connect.innovateuk.org/documents/2903012/16074728/RAS%20UK%20Strategy?version=1.0 (accessed February 2016)

Training ‘innovation ready’ RAS researchers

The EPSRC Robotics and Autonomous Systems Network (UK-RAS) is bringing together the UK’s core academic capabilities in robotics innovation under national coordination for the first time. It will encourage academic and industry collaborations that will accelerate the development and adoption of RAS.

Training ‘innovation ready’ researchers will be at the heart of the new £35 milion Edinburgh Centre for Robotics. A joint venture between Heriot-Watt University and the University of Edinburgh, and supported by the EPSRC and 30 industrial partners, it will graduate over 100 PhD students in advanced robotics research topics.

Summary

RAS hold the promise of revolutionising how people live their lives, making us more productive, more energy efficient and relieving us of the unpleasant "dirty, dangerous and dull" jobs. RAS pull together art, science and technology, take inspiration from nature and activate social innovation. There are, however, specific challenges that must be addressed before RAS can safely, reliably and usefully integrate with people.

Newman summarised the core capabilities required by robots to succeed autonomously (see Figure 10).

Fully autonomous robotics can be problematic due to ambiguity resulting from noisy sensors, lack of safety guarantees, restricted environments and un-modelled user intentions. As a result, there is a shift from isolated decision-making systems to those that share control, with significant autonomy devolved to RAS, leaving end-users to make only high-level decisions. Shared autonomy will demand the closing of the semantic gap between human and machine.

The journey from prototype to commercial system, whether it be a drone or a robot surgeon, involves many small steps and necessitates paying attention to and understanding the requirements of industry, allowing these to inform development. Being able to operate in realistic environments is very important to the transition of technology.

RAS researchers will need to use systems thinking and be "innovation" ready. It is necessary that they get out of the lab and work with those people who will eventually use these technologies, in order to discover the real needs and real issues associated with their application.

FIGURE 10

The core "spectrum" of capabilities required for autonomy.

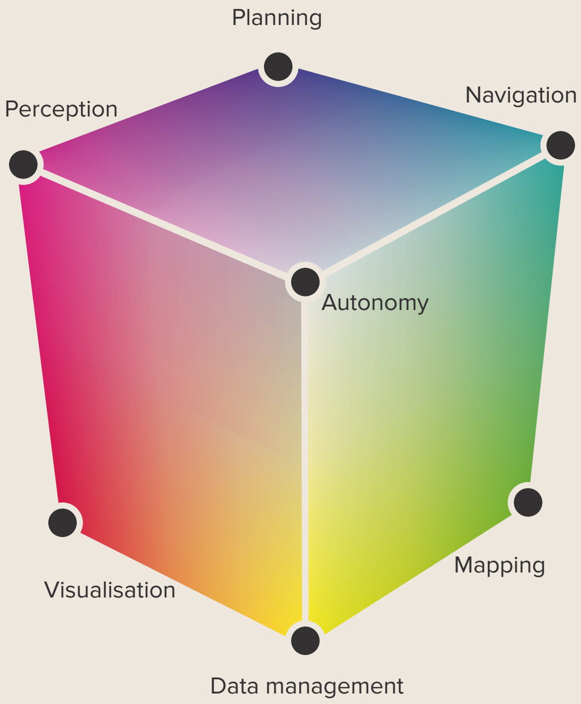

Image used with permission from Professor Paul Newman.

Acknowledgements

Organisers

Professor Peter Bruce FRS
The University of Oxford

Dr Ralf Speth KBE FREng
Jaguar Land Rover

Professor David Lane CBE FREng
Heriot-Watt University

Professor Paul Newman FREng
The University of Oxford

Chairs

Professor Peter Bruce FRS
The University of Oxford

Dr Andy Wright
BAe Systems

Professor Carole Mundell
University of Bath

Dick Elsy
High Value Manufacturing Catapult

Speakers

Bernard Charlès
Dassault Systemès

Professor David Lane CBE FREng
Heriot-Watt University

Professor Andrew Blake FREng FRS
The Alan Turing Institute

Professor Sethu Vijayakumar
The University of Edinburgh

Professor Katherine J. Kuchenbecker
University of Pennsylvania

Professor Nick Roy
MIT

Professor Peter Corke
Queensland University of Technology

Professor Paul Newman FREng
The University of Oxford

Professor Bradley Nelson
ETH-Zurich

The Royal Society is a self-governing Fellowship of many of the world's most distinguished scientists drawn from all areas of science, engineering, and medicine. The Society's fundamental purpose, as it has been since its foundation in 1660, is to recognise, promote, and support excellence in science and to encourage the development and use of science for the benefit of humanity.

The Society's strategic priorities emphasise its commitment to the highest quality science, to curiosity-driven research, and to the development and use of science for the benefit of society.
These priorities are:

- Promoting science and its benefits
- Recognising excellence in science
- Supporting outstanding science
- Providing scientific advice for policy
- Fostering international and global cooperation
- Education and public engagement

For further information
The Royal Society
6 – 9 Carlton House Terrace
London SW1Y 5AG

T +44 20 7451 2500
W royalsociety.org

Registered Charity No 207043
Issued: May 2016 DES4182