The Royal Society of Edinburgh
and the French Embassy in London

Franco–Scottish Seminar

Robotics

Wednesday 15 November 2017

Report by Jennifer Trueland

The RSE and the French Embassy in London are working together in a programme of science events to explore and present areas of science in which both Scotland and France have a powerful presence. This day-long seminar, aimed at academics and practitioners, was on the topic of robotics, and was followed by an evening public lecture delivered by Professor Sethu Vijayakumar, Professor of Robotics at the University of Edinburgh.

The seminar was chaired by Professor Vijaykumar and Professor Jean-Paul Laumond, Directeur de Recherche, LAAS-CNRS.

Participants were welcomed by Professor Vijaykumar, FRSE, Professor of Robotics, University of Edinburgh, UK. Director, Edinburgh Centre for Robotics and Royal Academy of Engineering and Microsoft Research Chair in Robotics and, representing the French Embassy, Mrs Stephanie Dos Santos, Deputy Scientific Attachée: Information and Communication Technology. Professor Vijakumar said the collaboration is important because it means bringing together the best of both countries, and offering the opportunity to exchange ideas and ask hard questions. Mrs Dos Santos said she hoped it would encourage new collaborations between French and Scottish scientists.

Dr Ravinder Dahiya, Professor of Electronics and Nanoengineering, University of Glasgow

Robotic Tactile Sensing and e-skin Technologies

Dr Dahiya discussed the importance of electronic skin in robots and in prosthetics. Tactile sensing is important for humans, he said, (for example, to enable the grasping of objects), so why would we create robots without tactile feedback; that is, skin?

The importance of being able to feel things should not be underestimated. For example, keyhole surgery relies on visual feedback, but doesn't let you feel texture, such as how soft the tissue is. Tactile feedback isn't just about hands; it's a whole-body issue. Imagine how your legs feel when you jump off a wall, he said – we need sensors all over the body.

There are practical reasons why sensory feedback would be good for robots. In the car industry, for example, robots are kept in cages because it is too dangerous to have them working alongside humans. He gave the example of a factory worker killed by a robot – had the robot had sensors, it would have 'known' to stop.

Several technologies to create 'skin' for robots have been explored, and Dr Dahiya outlined some of the work in the field so far, and looked to the future. Several factors are important when developing electronic skin. These include flexibility – it has to be able to stretch as the robot moves – and readability, and the practical issue of low power consumption; this latter point can be a big problem as you increase the number of sensors.

He spoke about the work of the Bendable Electronics and Sensing Technologies (BEST) Group at the University of Glasgow, a multidisciplinary group which aims to produce cost-effective, high-performance flexible and large-area electronics and sensing systems. These include looking at off-the-shelf sensors and electronic components, integrated on flexible printed circuit boards, to obtain large-area tactile skin for robots and wearable systems. He also spoke about alternatives, including the printing of nanowires, ultra-thin chips and grapheme-based energy-autonomous skin.

He spoke of the advantages and disadvantages of using organic and inorganic semiconductors to create the effective ‘skin’. While organic materials have the benefit of being flexible, they are low performance; while inorganic materials, for example silicon, are high performance but have low flexibility. Finding a way of making silicon-based electronics flexible is key, he said. Graphene is attractive, but the technology isn’t mature, he added.

Dr Dahiya described several projects, including *Roboskin*, which can be used on any surface, and which uses triangular patches to create ‘skin’. Silicon-based chips can become flexible when you thin the silicon down, he said. He also described *Printskin*, a printable tactile skin which is also showing promise.

Dr Dahiya concluded that tactile skin is critical for robotics, and that e-skin technology could underpin advances in several application areas, such as social robotics, bionics, the internet of things, and smart cities. High-performance flexible printed electronics is an interesting direction for obtaining e-skin, and hybrid technologies based on silicon and organic semiconductors could enable new solutions.

Flexible electronics is a disruptive area that could add new dimensions to the electronics industry (which is largely based on Moore’s law, the observation that the number of transistors in a dense integrated circuit doubles around every two years). There could be an enormous impact on areas including consumer electronics, aviation and space electronics, robotics, life sciences, lighting, military applications and telecommunications.

Professor Michel de Mathelin, Strasbourg University

Endoscopic Surgery through Natural Orifices: a Perfect Case for Robotics in the Surgical Field

Professor de Mathelin reviewed a decade’s research and development in the field of using robots to enable advances in endoscopic surgery, considering advances to date, challenges to overcome and what the future is likely to bring.

He began by describing ICube, a multidisciplinary laboratory in Strasbourg that involves a wide range of expertise. It is more like an institution than a lab and, being based in the grounds of a hospital, is convenient for working with and training surgeons. The medical robotics research team includes staff and students with backgrounds in robotics, mechatronics, computer vision, MR imaging, radiology and many other areas. It has a medical robotics platform with rapid prototyping facilities and an MRI 1.5 T platform shared with the hospital for MRI guided procedures.

One of the main research areas is robotic assistance to minimally invasive surgery. Professor de Mathelin spoke about the trend towards ‘no-scar’ surgery – which he said should more accurately be called ‘no-visible-scar’ surgery – and how robotics could help with this.

Laparoscopic robotic surgery already happens, albeit arguably only for the wealthiest. But it fits with the trend towards less and less invasive surgery. Single-port-access surgery – where there is one incision, for example through the umbilicus – minimises scarring, but is difficult in comparison to laparoscopic surgery; not least because the surgical workspace is more limited. Natural orifice transluminal endoscopic surgery (NOTES) uses the body’s natural orifices (such as the gastric system, colon or vagina) to reach the required area. This has challenges, including the need for at least two surgeons for one surgery; the flexibility of the instrument also means it's not possible to exert force other than retraction.

He described an early trial on the development of instruments for NOTES. The current prototype is the Anubiscope, a 2-DOF main endoscope with a diameter of 18mm; two instrument channels with good triangulation and instruments with 1-DOF flexion. More than 50 no-visible-scar surgeries have been performed in Strasbourg. Essentially, it's possible to remove the gall bladder through the patient's mouth.

In summary, there are limitations with manual instruments, including stability – disturbances can be caused by physiological motion – limited kinematics and the difficulty of synchronising endoscope and instrument motions. The surgery takes longer and at least two surgeons are needed. Robotics solutions could tackle many of these issues.

For example, active stabilisation using external sensors such as the endoscopic camera is possible; and the development of a telemanipulated robotised system would mean that only one surgeon would be needed for endoscope and instrument motions. Manipulability could also be improved with a telemanipulated robotised system that could synchronise endoscope motions and instrument motions for a specific task, and could provide different mappings between 'master' motions and 'servant' motions.

Experimentations on the first prototype have shown that telemanipulation works, but with a number of caveats. For example, at least 8-DOF should be actuated and preferably 9-DOF or 10-DOF; the rotation around its own axis of a two-way instrument is preferable to a four-way flexible instrument without rotation; and master and servant have different kinematics and it is not clear what is the best mapping. Work has been done to develop a simulation protocol, and two further proptypes, STRAS #1 and STRAS #2, have been developed with significant improvements. A fourth version is currently being developed for human use.

It has been a long process, said Professor de Mathelin, and has relied on a number of crucial factors. These include a close relationship between engineers and medical doctors, bringing the lab to the hospital and involving a large multidisciplinary team, and keeping the doctor – as a key opinion leader – involved. It's a complex issue and it will be necessary to keep developing new iterations of the equipment; there are also specific challenges, such as safety and asepsis, that have to be borne in mind. In the long run, however, it will be worth it, he concluded.

Questions

Asked about the value of automating the process, Professor de Mathelin replied that in medical robotics, you have to keep the surgeon in the loop – the robot should be the assistant to the humans. He also stressed that it has to be useful in the real world. "It has to make surgery better – if it's just the same, what's the point?"

Professor Helen Hastie, Professor of Computer Science, Heriot-Watt University

Establishing Trust in Autonomy Through Explainable AI

Autonomous systems are hugely useful, especially when they can be deployed in areas that are hazardous or inaccessible for humans. But there are questions around how much we trust robots.

If we look at depictions of robots in popular culture, we would probably trust C3PO from Star Wars to translate something into French, or R2D2 to unscrew a hatch door: we know what their capabilities are and that they have a good reputation, so we trust them. But what if it goes wrong, for example, like Hal in 2001: A Space Odyssey? How do you regain that trust? It's also hard to trust things where the reputation isn't already there – such as trusting an automatic car not to crash.

Professor Hastie talked a little about what trust actually is, and said there are many models of it. She quoted a Dutch proverb that 'trust arrives by foot but leaves by horseback' – in other words, it's hard to win, but easy to lose. Trust is the willingness of an individual to be vulnerable to the actions of another, based on the expectation that the other will perform the action without being monitored or control, she added.

Trust can be formed through predictability, for example, through reputation, and can be based on the trustor's idea of the trustee's ability, integrity or benevolence. But humans have varying disposition to trust. When it comes to humans trusting automation, generally we trust it initially: if we hire a car we expect the sat-nav to get us to our destination. But as the relationship with the system develops, dependability and predictability replace faith as the primary basis of trust.

In cognitive theory, mental models provide a view on how humans reason, whether functionally or structurally, that is, understanding what the system can do, or understanding how it works. What the user believes they know about a system strongly influences whether they use it, whether they trust it, and how they use it. That means that designers need to make a user interface communicate the basic nature of the system sufficiently, so that users form reasonably accurate mental models.

When humans and robots are collaborating, it requires trust between team-mates, social cohesion and task cohesion. Research also suggests that there is lower trust if the humans and robots are not located in the same place. There are three main mechanisms for establishing and maintaining the correct level of trust. These are communication, so the user can understand what the system is doing; transparency, so the user can understand how it works and why it is doing what it is doing; and shared autonomy, so the human feels more in control.

Professor Hastie described the MIRIAM project, (Multimodal Intelligent inteRactIon for Autonomous systeMs). MIRIAM enables operators to interrogate the actions and status of autonomous underwater vehicles (AUVs) in real time, through natural language. As well as updating status and progress towards meeting objectives, MIRIAM can give warnings and highlight faults. A short evaluation has shown that using MIRIAM improves situation awareness and user satisfaction.

This is a project that has been ongoing for more than 20 years, so pre-dates Alexa and Siri and other 'personal assistants', but the key thing is that you have to teach it what it can talk about. There are specific challenges; for example, the interface needs to be able to explain what it is doing; the mission might change; the environment might change, which means there is no predictable course of events.

There are also challenges in evaluation; for example, multiple factors might influence task success. However, evaluation of MIRIAM has shown that the more that people use the chat interface, the higher the user satisfaction and the greater the situation awareness. It also showed that users prefer short, chatty interactions. A second evaluation looking at trust and the mental model for AUVs showed that people with a clearer mental model generally have a higher level of trust in the vehicle, and that people who can see as well as hear what is going on have a clearer mental model.

Transparency is also important. Explanations help the user understand how the system works and what its capabilities and limitations are. Knowing how much trust to place in a model's predictions is important, especially when it affects decision making. For example, if there's an explanation on a website giving personalised movie recommendations, people are more likely to trust it and have a positive user experience. It's important to get the design of explanations right, however, as people won't read them unless they deem it worth the effort.

Professor Hastie then spoke about explainable AI (artificial intelligence) in the context of human–robot interaction. Explainable AI essentially involves the system to explain its reasonings and actions to improve trust. This is important, because trust is a vital commodity in allowing the human to feel confident with working with robots. The trade-off, Professor Hastie explained, is between control and trust in the system, and finding the optimum balance is key for the future. Trust and control can complement each other, but it's a dynamic relationship. Shared autonomy is important so that the human feels more in control: the more you control something, the less you trust it, and vice versa.

Human–robot interaction (HRI) is one solution. HRI through natural language enables truly collaborative interaction and can include advice, new instructions, repair, recovery, and much more. It can also lead to formation of trust through interaction, in much the same way as it occurs through human–human interaction.

There are ethical questions, however. For example, what if a system doesn't take advice and carries on with its previous plan, and who is to blame if something goes wrong in a human–machine team? These are all big questions as we move forward in this exciting area.

Professor Vincent Hayward, Professor (on leave), Université Pierre et Marie Curie (UPMC), Paris

Tactile Perception In and Outside our Body

We often underestimate the importance of our sense of touch, said Professor Hayward, but it is important – just think of whether you can manipulate an object when your hands are cold. Indeed, touch is trusted more than vision. Our primal sense of touch has evolved to give us a lot of information, such as what things are made of. For example, formica might look like wood, but it's clear from touching it that it feels like plastic.

Vision and hearing (audition) are much studied, as is taste, but that's not so much the case for touch. But it's complicated – and it's possible to 'fool' our sense of touch with illusions (for example, crossing our fingers, putting a pen between them, and 'feeling' it on the wrong side of the finger). It's all to do with how the brain solves problems and it might not be so much an illusion as a short cut.

Professor Hayward talked about tactile perception as a computational experience. The physics of contact, and what it means for our tissues, differs from the physics of acoustics and optics, and the way we experience touch is likely to have had an influence on the way our sensory systems have evolved at every level.

In his talk, Professor Hayward gave examples of how surface physics shapes the messages that are sent to the brain. One example is around how the brain perceives glass – although we think we are touching it, it takes a while for skin to make true contact with an object; this is because, as humans, we are essentially exuding water all the time. This has a dramatic effect on the skin and alters the mechanics. The field of haptics – and associated modelling – provides completely new opportunities for applications at the human–machine interface.

Dr Zhibin Li, Lecturer, University of Edinburgh

Reaching the Unreachable: The Science of Robot Locomotion

Many years of research has focused on robot locomotion, with the aim of creating walking, leaping, climbing and running robots that can operate on a variety of terrains. The advantages of such robots would be manifold, both on this planet and beyond. For example, terrestrial uses could be in scientific expeditions, or to inspect and maintain machinery, or conduct rescue missions in disaster scenes – or even to deliver goods to our homes. Such robots would also have many applications in space programmes; for example, in making Mars a habitable planet. Despite decades of research, however, robot locomotion that comes close to replicating what humans can do has not been seen outside the lab.

Progress is being made and moving robots are being used in a number of settings, but they have limitations. For example, the Mars rover, a robot with simplified 'legs', can be used to reach and operate at sites of interest on Mars, but wheeled robots can only operate in limited areas. Wheeled robots do fall and they can't step over obstacles.

These compromises are why legged locomotion matters, said Dr Zhibin Li. As humans, we can walk, step over things, leap and climb – which means we can traverse a large variety of terrains. Humans have survived for millions of years because we can do rock climbing, he said. With legs, you can almost reach every corner of the planet, places that are unreachable by machines. That's why the quest to 'reach the unreachable' is so urgent. The ultimate goal is an all-terrain mobility solution; legged locomotion complemented by wheels.

Dr Li spoke a little about robots for disaster response and the Defense Advanced Research Projects Agency (DARPA) Robotics Challenge, which focused on this area. The top three robot winners all had legged designs. Wheels are good on flat areas, but legs are good at stairs, he said. But there are limitations. For example, if a human operator makes a mistake, the robot falls. The robot makes no decision, so there is no autonomy at a local level.

Lessons learned from this include the separation of locomotion and manipulation. The scientific challenges include motion planning in a complex environment, and how to allow a certain amount of autonomy with minimum human supervision. Whole-body motion planning considers manipulation and locomotion. He showed an experiment to plan for a robot to move to an object and grasp it, and discussed work to achieve multi-contact locomotion in uneven terrain.

The main lesson from the DARPA Robotics Challenge was that robots need to be more autonomous – we need to give them the skills to deal with difficult situations on the ground. But how do you fill the gap between high-level human supervision and low-level control for task execution, particularly where there is limited communication between the two?

Dr Li described research into deep learning for robot control, where the machine learns to balance and control its locomotion. This proof-of-concept research used deep deterministic policy gradient (DDPG), a model-free, actor–critic reinforcement algorithm. The actor network learns the policy that generates the optimal action, while the critic network evaluates the performance of the policy. Deep learning for rough terrain locomotion, for example, gives feedback on body orientation, speed and velocity and other factors, and also helps the machine to cope with obstacles in a modelling scenario.

A number of enabling technologies are needed, such as tactile sensing and robot skin for physical interactions, machine learning and explainable AI for robot autonomy, and shared control so that humans can handle what robots cannot do. The applications are potentially worth a lot of money; for example, the UK market for offshore wind turbine operation and maintenance is valued at £2 billion by 2025.

There would be a number of positive impacts if we "reached the unreachable", Dr Li concluded. These include improved safety, because humans would not have to undertake risky activities; improved efficiency and higher productivity; better jobs for humans (because they are freed from having to do dirty, dull and dangerous jobs); and proactive action to conduct more frequent inspections of difficult-to-access locations, catching defects early and preventing escalation of damage.

Dr Jean-Pierre Merlet, Senior Research Scientist, INRIA (French Institute for Research in Computer Science and Automation)

Robotics: The Can Do and Can't Do

Dr Merlet's talk was historical and prospective, looking at development of robotics to date and considering where it is likely to go next, and the social and ethical implications of this.

In the year 2000, robots were working in industry, but they were 'caged' and kept apart from humans by physical barriers. The major change today is the loss of that barrier – there are even robot toys and machines such as robot vacuum cleaners are in our homes.

The appearance of robots that look like humans has led to a perception that intelligent robots are on their way, giving rise to old fears of robots taking over the world. To be realistic, however, current robots are not that smart, said Dr Merlet. In a very structured context, they might be faster and more efficient than humans, but if they are taken out of that context then they become ‘pretty dumb’. Asking where we are going next, Dr Merlet questioned whether the future would involve a single universal robot that could do every task, or a fleet of relatively specialised robots that together could come close to the ‘universal’ robot.

Dr Merlet looked at the practical issue of assistance robotics to explore these questions. He spoke about the HEPHAISTOS project, an INRIA team focused on assistance to frail people, such as older people and those with disabilities. When the project started in 2006, they had almost no knowledge about the needs of this population, so they conducted in-depth research with a number of stakeholders, including end users, local authorities, medical staff and caregivers (nurses and family). Over a three-year period, they conducted some 200 interviews. The information from the interviews allowed the researchers to determine priorities such as mobility assistance and medical monitoring (especially at home) and also come up with guidance on how it should be done; for example, what is ethical, and whether something is intrusive or not, and the cost both of the object and of the energy required to run it. The interviews revealed what really mattered to people – for example, one person said a real change in their life would be being able to go alone to the toilet. Mobility is essential for autonomy, Dr Merlet said. He talked in particular about mobility issues of transfer (e.g., from sit to stand) and walking assistance.

There are several solutions available to help people to transfer, including hoists and harnesses. These tend not to be used because they are cumbersome and time-consuming, Dr Merlet asserted. So, what would a robotic solution look like? A humanoid robot would not be the answer, he said. These are energy inefficient, transfer is mechanically demanding, and safety would be a concern – the robot would have to be reliably stable. This trajectory would lead to a large (intrusive) robot with low autonomy. He gave the example of the RIBA robot (Robot for Interactive Body Assistance), developed in Japan, which is intrusive, costly and requires a (human) helper.

A more specialised robotised solution is a cable-driven parallel robot (CDPR), a platform connected to the ground through a set of independent cables. This is low cost and non-intrusive and has a high lifting capacity. CDPR can also be used to prevent mobility problems, such as by preventing falls (for example, to pick up things from the ground to prevent challenging posture and risking falling). CDPR could also be used for mobility monitoring, which is currently assessed via clinical tests, which are not objective, can often be inaccurate or incomplete, and are not universally available. The CDPR solution with an instrumented platform provides synthetic walking indicators that are robust and medically pertinent.

Dr Merlet then looked briefly at whether mobility aids such as walkers or canes could be robotised for assistance and falls management. Researchers took a commercially available walker and added sensors to assess walking – in experiments, interestingly, elderly people walked faster than young people, but took a less direct route. Similar things can be done with a cane. This technology also has other applications, such as helping people to navigate a city by telling them where lowered kerbs are, for example.

He concluded by saying that the last 20 years have seen huge progress in robotics in terms of hardware, sensors, computers, dedicated software, actuation and algorithms. But we are still very far away from an ‘overall smartness’ for our robots. Social robots remain a dream that might occur by the time we are very old, unless there is a drastic change in computing technology. But there are also ethical issues: closeness between humans and robots raises issues such as responsibility when things go wrong; users and families must always be aware that nothing will be 100% bug-proof. And what will happen to collected data, and to whom does it belong?

Other questions are around whether a robot can disobey a human (for example, for the sake of the human's safety), and autonomous robots would potentially have to make difficult decisions; for example, would an autonomous vehicle make the choice to kill the driver or pedestrians, and how would this happen?

The responsibility of roboticians is to present all the possible futures of robotics so that society can decide.

Questions

There was some discussion around ethical issues, with some audience members suggesting that these would be harder to solve than scientific problems. Questions included who would be responsible if a vacuum cleaner robot killed an older person. One audience member said that it would be difficult to make robots know what to do in different contexts – and asked who would be responsible if an automous car killed someone – would it be the car manufacturer or the company that provided the autonomy technology? Professor Laumond remarked that the mobility assessment technology would also be useful in rehabilitation, where there were few objective measures. Objective assessment would make the rehabilitation process more successful.

Professor Jean-Paul Laumond, Directeur de Recherche, LAAS-CNRS

Bipedal Walking: A Multidisciplinary Perspective

After citing some of the amazing things that robots can do – for example, play the trumpet – Professor Laumond talked about their limitations. He spoke about the high numbers of failures at DARPA, for example, showing robots falling over and failing to complete tasks. Nevertheless, humanoid robots are a nice platform for fundamental research: in trying to find out how to make them walk, we can learn more about how we actually walk, which means that experimentation can benefit other scientific areas.

Movement can be considered from a robotics and a biomechanical perspective, he said. It is complex – if you tell a robot to take a glass, it has to negotiate physical, sensory and motor space. A human has a brain to do this, but the robot has programmed software that has to take into consideration many complex factors such as motion, placement and posture. A better understanding of the way that humans actually do tasks, such as walking, should help reach the goal of designing better bipedal robot architectures. Again, humans have the edge, with 700 muscles to control three variables, compared to the robot's 30 motors.

Professor Laumond spoke about the differences between walking and rolling, and paid tribute to the wheel: put it on the ground and it transforms, covering the distance on a trajectory. Similarly, when we walk, we do so from front to back, not sideways – following the trajectory that we believe is optimum. This hypothesis is wrong, however, because if it was the optimum trajectory, we would take the same route in reverse – and experiments show that we don't. Instead, we head for where we can see the goal. There is no miracle in this, but it involves including the bearing angle, which changes the perspective.

Professor Laumond spoke about what happens when humans walk without thinking – they are more likely to stumble than if they are taking care and mapping their steps. Passive robot walkers can be modelled by a simple rotating, rimless wheel, and we know that humans and animals stabilise their heads while moving. This led to the hypothesis that a wheel with a stabilised mass on top of it could be a plausible model of bipedal walking. This is called the 'yoyo-man'.

He cited research showing the mechanical effect of head stabilisation – it contributes to locomotion robustness. Modelling shows that the centre of mass follows a tronchoid – essentially the fixed point of a trajectory on a wheel. This suggests that taking a new 'downward' approach to humanoid robot design, through exploring the computational foundations of walking like humans, could be effective.

Questions/discussion

Discussion ranged around whether carrying things on your head makes your walking more stable – this could potentially fit with the rolling man hypothesis. Asked if the same is noted when people are watching their step, Professor Laumond said that it only happens when people are walking without thinking.

The seminar was followed by a public lecture from Professor Vijayakumar. A summary report of this can be found on the RSE website at: www.rse.org.uk/event/franco-scottish-lecture-shared-autonomy-the-future-of-interactive-robotics/

Opinions expressed here do not necessarily represent the views of the RSE, nor of its Fellows.

The Royal Society of Edinburgh, Scotland's National Academy, is Scottish Charity No. SC000470