Foresight Cognitive Systems Project Research Review

Robotics and Cognition

Edited by Paul Newman and Lionel Tarassenko

This discussion document is part of a series of reviews of the state of the art in cognitive systems. It was prepared as background material for the Foresight Cognitive Systems Project. The objective of the project is to examine recent progress in two major areas of research – computer science and neuroscience (and their related fields) – and to understand whether progress in understanding cognition in living systems has new insights to offer those researching the construction of artificial cognitive systems.

Further details are available at the Foresight web site: http://www.foresight.gov.uk/

Contact the Foresight Cognitive Systems team at: Foresight Directorate Office of Science and Technology
1 Victoria Street
London
SW1H 0ET
Fax: 020 7215 6715
E-mail: cognitivesystemteam@dti.gsi.gov.uk

CONTRIBUTING AUTHORS

Dr Dave Barnes	University of Wales, Aberystwyth,
David Bisset	iTechnic Ltd
Dr Guido Bugmann	University of Plymouth
Professor Robert Damper	University of Southampton
Professor Kerstin Dautenhahn	University of Hertfordshire
Professor Brian Davies	Imperial College, London
Dr Andrew Davison	University of Oxford
Dr Alex Ellery	University of Surrey
Dr Dylan Evans	University of the West of England
Dr John Gan	University of Essex
Professor John Gray	University of Salford
Dr Tony Hague	Silsoe Research Institute
Professor Huosheng Hu	University of Essex
Professor John Illingworth	University of Surrey
Professor Colin McInnes	University of Glasgow
Professor Chris Melhuish	University of the West of England
Professor David Murray	University of Oxford
Dr Ulrich Nehmzow	University of Essex
Dr Paul Newman	University of Oxford
Dr Murray Shanahan	Imperial College, London
Dr Barbara Webb	University of Edinburgh
Dr Alan Winfield	University of the West of England

With great thanks to Michael Kenward for proof-editing and many constructive comments.

INDEX

INTRODUCTION

Project Aims

In this report, members of the robotics community in the UK describe some of the issues facing robotics researchers, and report on current research. This report, a part of the Foresight Cognitive Systems Project, is intended for members of the cognitive and life sciences community. The report set out to achieve the following goals:

- To report on the existing intersection between cognitive and robotics science in UK
- To report on the broader intersection between robotics and the life sciences
- To sample other work being undertaken by the UK robotics research community

The report sets out to fill a gap in the original series of Research Reviews commissioned by the Foresight Cognitive Systems Project. While at first it seemed likely that those reviews would deal with most of the issues germane to research in robotics, it became clear that while those reviews did indeed cover many of the underlying research themes of robotics, they did not address the issues involved in bringing together work on, for example, sensors and users interfaces. Nor is it obvious that the original reviews cover all of the areas where researchers in natural cognitive systems can contribute to robotics, which has its own special set of cognitive challenges.

When considering the intersection between research in natural and artificial cognitive systems enquiry, it is natural to ask how knowledge does, and might in the case of new collaborations, flow between the disciplines. For example, what can roboticists learn from cognitive scientists about perception and spatial reasoning? Can robots acts as test-beds for neurophysiological hypotheses? How might the life sciences enable biological 'food' to power mobile machines to achieve 'energetic autonomy'? How can we, as individuals, interface and relate to machines, and how does this tie in with cognitive models of language?

The following sections describe robotics research that is either already entwined with cognitive and life-science issues or asks questions that are likely to stimulate the interest of cognitive and life scientists – hopefully resonating across the fields. This report is not intended to be an explicit, encyclopaedic reference for robotics research in the United Kingdom, however it is does establish an implicit representation of the field through its contents and contributing authors.

Foreword

In the opening section Chris Melhuish, who heads the Intelligent Autonomous Systems Laboratory (IAS) of the University of the West of England (UWE), writes about the role and characteristics of artificial intelligence (AI) in mobile autonomy. He highlights clear shortcomings in the state of the art, especially when handling the 'torrential flow' of information that bombards 'embodied systems'. The discussion of autonomy leads to a tantalising description of work at UWE regarding the creation of robots that feed on biological food sources, "artificial digestion may be the prerequisite of artificial intelligence".

The next section is also concerned with autonomy, albeit from the perspective of navigation. Paul Newman from Oxford University asks how can a mobile robot be placed in an unknown environment and be expected to navigate within it? What problems of perception and data fusion are involved in answering the question “Where am I?”. This topic raises questions about understanding and modelling known and unknown environments, how sensor data can be interpreted and thinned. Andrew Davison, also from Oxford University, then carries the navigation and mapping theme onto a discussion about real-time, single-camera, navigation and mapping in common human workspaces.

Then comes a section on sensing and interpretation. Huosheng Hu and John Gan of Essex University discuss sensor design and deployment and the use of the measurements they produce. The sensors cover the whole spectrum from touch to sonar and laser to vision and highlight the difficulties caused by ubiquitous non-linearities in the real world. The authors identify the research occurring to overcome these difficulties. Murray Shanahan, of Imperial College London, finishes the section by discussing the use of abductive reasoning in sensor interpretation and cognitive and spatial reasoning.

In the following section, Barbara Webb, of Edinburgh University, considers the role robots can play as biological models: “Another view of these systems [*robots*] is that they are physical models of animals that we can use to address specific questions in biology. Like more conventional models and simulations in biology, they help to enforce clarity and specificity in hypotheses by requiring all the details to be made explicit. These models also aid in determining the consequences and predictions of complex hypotheses.” This is followed by a discussion by Robert Damper of Southampton University about the effect ‘embodiment’ has on research into bio-robotics.

Ulrich Nehmzow, of Essex University, begins the section on Interaction and Behaviour by discussing his research into analysis of behaviour through mobile robotics: “The first requirement for the analysis of mobile robot behaviour is that of establishing a ‘language’, a means of communicating properties of robot, task and environment precisely, quantitatively”.

Guido Bugmann, Kirsten Dautenhahn and Dylan Evans, of Plymouth, Hertfordshire and the West of England Universities, then follow and write about Human Robot Communication: “Effective communication between user and robot is essential to access the full capabilities of such robots. The capability for user-friendly communication will in itself contribute to the perceived value and social acceptability of the robot.” In this light they consider the disparity between the required and current sophistication of interaction between robots and humans in social, emotional and linguistic contexts.

The final section, Applications Research, focuses on application-motivated research. David Murray of Oxford University suggests that the “inexorable reductions in size, and gains in power efficiency, of electronic components, and progress in embedding and integration them within textiles, opens the way for clothing to acquire something of the nature of a second ‘perceptual skin’”. He goes on to consider how this might assist physiological sensing and how wearable cameras might enable inference user intent.

Alex Ellery, Dave Barnes, Colin McInnes, Alan Winfield and Ulrich Nehmzow, of Surrey, Aberystwyth, Glasgow, the West of England and Essex Universities, discuss the topical subject of space robotics and describe research in the UK. “Robotics is also the key to one of the most exciting scientific endeavours of the 21st century – the search for life elsewhere.”

Brian Davies of Imperial College London summarises the state-of-the-art in medical robotics and documents the progress towards reliable robotic based surgery, offering tangible improvements to health care. David Bisset of iTechnic Ltd considers the cognitive requirements placed on robots operating in the home. “The autonomous appliance will need, or appear to have, a good enough understanding of the cognitive context of the items it encounters. In essence, there must be a cognitive dialogue between the appliance and the user.” John Gray of Salford University briefly discusses the role of automation in food production before Tony Hague of Silsoe Research Institute finishes the report off with a discussion of the current and expected use of autonomous mobile robots in agriculture.

A recurring theme in this review is that the lack of adequate perception for robots blocks the path between the current state-of-the-art and the kind of robots we desire. Much of cognitive science is motivated by the desire to understand human perception. It could be said that robotics researchers want to construct a perceptual system and cognitive scientists want to deconstruct one. Where do these top-down and bottom-up approaches meet? How much can we learn about our own cognitive ability by trying to engineer a perceptual system from scratch? Conversely, how can reverse engineering of human perception influence the design and increase the value of robots to society? This report is timely precisely because these open questions are being asked in both the cognitive neuroscience and robotics communities.

The editors hope this report is a thought provoking and interesting read for roboticists and researchers in natural and artificial cognitive systems. Of course, the breadth of material covered combined with constraints on space does not allow in depth consideration of many key issues. However, we would consider the report a success if it enabled or prompted researcher in a cognitive systems – real or artificial – to open up a dialog with a robotics researcher, perhaps leading both parties to a deeper understanding than now exists.

Paul Newman and Lionel Tarassenko

Oxford, November 2004

AUTONOMY AND NAVIGATION

BUILDING AUTONOMOUS ROBOTS

Chris Melhuish

Introduction

In this section we focus on the idea of autonomy in robots. We address the nature of autonomous robot systems and how they are being developed. Examples of the challenges faced by autonomous robots, including the generation of their own energy, are briefly discussed. The report includes examples of current research in the UK.

Industrial robots perform their task in highly controlled environments, under direct or indirect human control. However, some robot systems, with one or more robots, will carry out missions in unpredictable environments without human supervision.

The term ‘autonomous robot’ is a generic term and encompasses a continuum – one from ‘dumb’ insect-like robots to highly sophisticated robots, interacting with humans, requiring social intelligence, theory of mind and an ‘awareness’ of self as distinct from other organisms and its own environment. The former, perhaps, do not require cognitive capacities, but the latter almost certainly will. A robot in this category may well require ‘predictive power’ that enables it to, in the words of Karl Popper, “allow its hypotheses to die in its stead”. Such a robot would have the capacity to predict what its state and its environment might be for sets of conditions that the robot may not have experienced.

Autonomous robots may be characterised, for example, as machines which, in some combination, “do things you don’t want to do, in a place where you don’t want to be, at a time you don’t want to or can’t be there”. Of course, one can imagine missions with varying degrees of external intervention – a continuum of autonomy. The vast majority of robots labelled as autonomous are therefore obviously not ‘fully’ autonomous but display a degree of autonomy. This is often referred to as ‘semi-autonomous’.

Examples of autonomous robots include deep-space missions, stealth reconnaissance and deep-sea mapping. However, it can be argued that there is a huge potential market for less exotic applications including domestic robots, security robots, robots for the leisure industry, such as gardening and care robots.

The public is well aware of the concept of ‘robot helpers,’ but the devices simply aren’t there for them to purchase. Apart from the cost issue – mass sales will reduce prices – the main barrier to building autonomous robots to meet such a demand is that they aren’t sufficiently ‘smart’ to be reliable. A key question is how do we make them smarter? Robots employing biologically inspired cognitive mechanisms may hold the key.

Living models

Animals and plants provide an obvious existence proof of ‘autonomous systems’. In Nature, individuals and groups of animals demonstrate abilities to navigate, move, sense, process and select appropriate actions, communicate, plan, repair and maintain themselves, supply their energy demands, interact with the environment and other animals, as well as the impressive feat of reproducing themselves. While self-replication may be further in the future, these are also the problems faced by roboticists interested in building machines that can carry out missions without human intervention.

Many roboticists are therefore naturally attracted to biology with its wealth of mechanisms and behaviours which endow animals with the abilities just described. Successful animals demonstrate the ability to solve the ‘action selection’ problem – doing the right thing at the right time, within the constraints of environmental conditions, resources and energy budget available.

Advances in autonomous robotics do not, however, have to be based on biological examples. Human designers are free to base their solutions on decidedly non-biological mechanisms. The point is that, to generate engineering solutions, it makes sense to be aware of and, if appropriate, to exploit, known biological principles.

THE ROLE OF ROBOTS

But what do we want robots to do and what approaches do we use to get them to do it? In very broad terms, robotics has experienced two waves; employing classical artificial intelligence (AI), and new AI (or behaviour based) strategies. Both are excellent for particular application areas. However, this section argues for a renewed focus on a cognitive approach.

Some researchers and technologists argue that for many applications the employment of classical cognitive AI is adequate – indeed excellent. At its core – and there are, of course a number of caveats – cognition was classically considered as rule-based manipulation of symbols. In this approach, a robot generates symbolic descriptions of its world and employs these to update and reason over an internalised world representation in order to select the next action. Although there is debate about how ‘cognitive’ this is, the approach has proven to be extremely successful for well defined tasks – for example, automated chemical analysis, car-assembly robots and smart missiles – and when the environment is strongly constrained and controlled, employing a significantly smaller degree of input than a biological system, typically within set boundaries often known in advance, operating within predictable environmental and sensor noise.

There are, however, environments where this otherwise successful classical approach has been found wanting, namely in unpredictable and unstructured environments. A new approach was heralded by the appearance in the 1980s of systems that generated robot behaviours as an emergent consequence of the interaction of simple, low level modules. The approach was aptly named ‘behaviour based’ and relied heavily on reactive mechanisms connected in such a manner that higher level modules were able to ‘subsume’ or inhibit the behaviour of lower level modules. There was also provision for information and control to filter up as well. Without subsumption or inhibition, each active module would compete for the ‘final common path’ to access the external actuators.

Some researchers claimed that this approach would generate “intelligence without representation” or “intelligence without reasoning,” perhaps at the level of insects. Some researchers would argue that this approach can be extended to ‘higher’ animals. Robots built on these principles were able to outperform competitors employing classical AI when faced with such problems as dynamic obstacle avoidance and corridor following.

Although this approach did not explicitly eschew the use of higher level planners, it has been argued that robots will require something different from and more than that which either classical AI or behaviour based systems offer. This is particularly true in domains where embodied systems require a cognitive sophistication capable of compressing a torrential flow of information in order to extract pertinent elements fundamental to the robot’s survival manifested in appropriate behaviour. This is precisely the challenge successfully faced by animals over the millennia. It is also the challenge faced by ‘truly’ autonomous robots. Furthermore, and perhaps more importantly, like their animal counterparts, autonomous robots will have to deal with the energy problem. Behaviour requires energy as does maintaining ‘metabolism’, communication and computation.

The history of robots shows an increasing ability to interact with objects. Early robots could deal with only static environments, clear of obstructions. Later robots could cope with static obstructions and, later still, roboticists developed robots that could keep out of harm’s way of moving obstacles. Humans, of course, represent dynamic obstacles too – but significantly more complex ones!

Certainly, a key challenge for robotics research will be creating autonomous robots that interact with people in non-trivial ways. Such robots will need to reason about their goals and what actions to take. They will be required to focus their perceptive capabilities and be sensitive to the cognitive states of humans and other robots.

As is always the case in research, there is debate on how this could be achieved. However, if robots are to share the same environment as humans, engage in dialogue with humans, be instructed by humans, offer advice to humans, cooperate in social tasks and so on, then designing robots based on knowledge of human cognitive processes offers a good starting point. Research will be need to construct a coherent theoretical and implementation framework to pull together reasoning, perception and action selection in the context of autonomy.

Autonomous Systems Research in the UK

Researchers are investigating many of the elements required by an autonomous robotic system, such as navigation, sensing and learning. (These are discussed in more detail elsewhere in this report.) The scale of the problems posed by attempting to build autonomous robots means that these elements are spread across the academic community engaged in robot research. Examples include the research at Essex University which encompasses a range of topics relevant to autonomous mobile robots, including hardware design, sensor signal processing and computer vision, navigation, robot learning, collective robotics, biologically inspired robotics and analytical robotics.

A group at Sussex University is exploring the use of neurological models of insects. These researchers argue that our understanding of the central nervous system should gain a great deal from attempts to model it in robots. A collaboration between the Universities of the West of England and Sheffield is exploring elements of mammalian neural architecture in an attempt to ‘embody’ exploratory behaviour in a robot. The approach focuses on hardware implementation of neural structures.

The Robotics Group at Imperial College is best known for its work in cognitive robotics – the application of formal methods of reasoning to robotics and software agents. While not generally intended as direct biological models – in contrast to some of the research at the University of Edinburgh, for example – the group’s work is both inspired and guided by biological and psychophysical evidence such as the work on robotic imitation, derived directly from neurophysiological evidence of ‘mirror neurons’ in primate brains.

The universities of Plymouth and Hertfordshire focus on several key issues in the design of domestic and helper robots including artificial vision for object recognition and navigation, action planning and natural language instructions dialogues between user and robot.

Researchers at the University of Wales, Aberystwyth, and their partners SciSys Ltd, the University of Leicester and Joanneum Research, Austria, are developing a demonstrator imaging and localisation package (ILP) for an autonomous Martian balloon. Aerobot technology is generating a good deal of interest in planetary exploration circles.

Balloon based aerobots have much to offer the Aurora programme of the European Space Agency (ESA), which aims to formulate and then to implement a European long-term plan for the robotic and human exploration of solar system bodies holding promise for traces of life. For example, the benefits include high resolution mapping, landing site selection, rover guidance, data relay, sample site selection, payload delivery, and atmospheric measurement.

Smart imaging and localisation is a key enabling technology for remote aerobots. Given the current lack of comprehensive localisation and communications systems, it is important that aerobots have the ability to determine their location, with respect to a planet’s surface, to a suitable accuracy and in an autonomous way. The availability of a variety of salient terrain feature (e.g. crater) detection, feature tracking, and image compression algorithms means that such a self-reliant system is now achievable.

If we are to develop robotic outposts for human habitation of planets such as Mars, then autonomous robotic systems are a major research area. The British 2003 Beagle 2 Mars lander by way of comparison, had little onboard autonomy. It was planned that mission engineers and scientists back on Earth would undertake most of the engineering and scientific cognition. Future robotic systems will have to survive for extended periods with few, or no, opportunities to “phone home,” hence the need for greater autonomy.

UNDERWATER SENSING

Unmanned autonomous underwater vehicles (AUVs) are being developed to carry out underwater sensing and survey work. The University of Southampton has developed Autosub a long-range, deep-diving underwater vehicle that can carry a variety of physical, biological and chemical sensors. The research team claims that this technology provides the ability to monitor the oceans in ways not possible with conventional research ships. This robot has completed over 270 missions, has operated in force-6 weather conditions and is rated to a depth of 1600 metres.

The Ocean System’s Laboratory at Heriot-Watt University also has an impressive portfolio of AUVs. The European collaborative project ARAMIS, for example, will include automatic obstacle avoidance, automatic recognition and tracking, path planning, texture analysis of seabed images and concurrent mapping and localisation. This will provide a highly automated robotic tool for benthic investigations in deep and shallow water.

ENERGY NEEDS

All robots require energy to carry out their tasks. Some systems, labelled as autonomous, AUVs for example, operate in an autonomous manner as long as their onboard battery can sustain them. However, many tasks will require periodic recharging of batteries, no matter what advances are made in battery technology. Some robots will therefore need to be able to find and dock with a charging station. Examples of this class of robot exist and one can readily imagine that in the future there will be types of robots, such as domestic care robots, which employ existing electricity outlets in the management of their energy budget. However, for sustained autonomy where there is no recourse to such infrastructure, robots will need to extract energy from their environment. Solar power has been employed to generate electricity but there will be many applications for autonomous systems with no access to solar energy.

Research at the Intelligent Autonomous Systems (IAS) laboratory at the University of the West of England is looking into ways to extract energy from raw substrate in the environment, that is, robots that feed. The group has constructed a mobile robot which, by employing novel microbial fuel-cell technology, can convert sugar into electricity, move to a target and transmit sensor information. The group is now concentrating on the generation of sugars and equivalent input molecules from raw feedstocks such as plant material. In this way artificial digestion may be a prerequisite of artificial intelligence for some autonomous robots!

The management of energy opens a new dimension of control for autonomous robots, leading to devices that will have to balance the energy requirements of their missions and the energy requirement of how they find food sources, convert and accumulate their onboard energy as well as maintain their internal 'metabolic' systems. The IAS laboratory is also looking into how multiple robot applications might exploit energy sharing.

THE RESEARCH COMMUNITY

There is an active community of robotics researchers interested in the interface between life sciences and the physical sciences. The Biologically Inspired Robotics Network, Biro-net, is a British academic network, centred at the University of Wales, Aberystwyth, and funded by the EPSRC. Biro-net aims to establish a network for advancing interdisciplinary research in the UK which will further our understanding of the underlying mechanisms that allow natural and robotic agents to adapt and survive in uncertain and dynamic environments.

The aims and objectives of biro-net, which has a membership of some 65 academics and industrialists, are:

- to bring together academic and industrial researchers in biologically inspired robotics as a community for the exchange of ideas, technologies and contacts for collaboration;
- to provide an informal forum for communication and exchange of ideas through organised symposia and workshops;
- to expand the initial core of network members; and
- to advertise the field more widely across related disciplines to stimulate interest, to aid the cross fertilisation of ideas, and to stimulate research in complementary areas.

With thanks to: Dave Barnes (Aberystwyth), David Bisset (iTechnic Ltd), Dylan Evans (UWE), Ulrich Nehmzow (Essex) And Mark Witkowski (Imperial College).

NAVIGATING IN UNKNOWN ENVIRONMENTS

Paul Newman

Introduction

Navigation lies at the heart of mobile robotics. The very act of moving from one location to another inevitably requires some kind of answer to the question "Where am I?" Of course, there are many ways to answer this question, and with varying precision. But fundamentally, spatial awareness is crucial.

This section of the report looks at how robotics researchers have tried to answer the "Where am I?" question by using techniques centred in estimation and probability theory. They are engineered solutions and do not have a significant intersection with the life sciences. However, it is interesting, and perhaps no surprise, that the conclusion is that the single and greatest impediment to robust autonomous navigation is a problem that has challenged cognitive scientists since day one – perception.

Three Navigation Sub-Problems

Consider the case of a mobile robot equipped with only onboard sensors: an external observer cannot tell the robot where it is. We can imagine using cameras, global positioning systems (GPS), compasses, lasers, wheel-rotation sensors, accelerometers, sonars, radar etc. The question is, how can we make the robot navigate?

The term 'navigate' covers a spectrum of tasks and skills. It will be useful to consider the following sub-problems:

- Localisation – producing an estimate of vehicle position
- Mapping – producing a representation of the operating environment.
- Simultaneous Localisation and Mapping – (SLAM)

This intentionally narrow definition of navigation does not include tasks such as obstacle avoidance, path planning and exploration which are commonly lumped in with the 'navigation' problem. The point here is that frequently, with the notable exception of purely reactive behaviour based schemes, these tasks need an ongoing answer to "Where am I?".

Solutions to any of these sub-problems require some representation of space. Commonly, the locations of vehicle and workspace features are described using Euclidean co-ordinate frames. Although we shall go on to discuss mobile robot navigation using the Euclidean descriptions, it is not the only approach available. Good progress is being made in the use of qualitative descriptions such as topology – for example the Spatial Semantic Hierarchy of Ben Kuipers at the University of Texas. Currently, however, Euclidean representations dominate the field and so will be discussed exclusively here. Furthermore, we acknowledge that for many commercial, military and industrial applications a metric idea – with respect to some fixed reference – of vehicle location is essential and further motivates, from an engineering perspective, the choice of global, Euclidean spatial descriptions.

LOCALISATION

Cartesian localisation – the task of determining the robot's position – always requires some knowledge of the robot's workspace. This could be in the form of the existence and use of infrastructure such as a constellation of GPS satellites or some kind of prior map.

Of the three navigation subtasks under consideration, localisation is by far the easiest and is to a large degree considered a solved problem in robotics. The map may be as simple as a set of surveyed coordinates of beacons – deployed for example on the ocean floor to enable localisation of a remotely operated vehicle (ROV). A sensor on the vehicle makes noisy relative measurements, perhaps range and bearing, to the beacons: with well known estimation techniques it can then deduce its position.

Beacons are designed to be easy to spot. Measurements derived from them are, for all intents and purposes unambiguous. Beacon-based localisation is an engineered solution and is suitable for deployment in many common environments for example warehouses, and offshore maintenance.

Things become much harder when the map provided is simply a description of the workspace – an architectural map of a museum or hospital, for example. Suddenly measurements are far more ambiguous. For example, consider a robot equipped with a laser scanner which returns a swathe of ranges and angles to nearby objects. Given an architectural map of the workspace, it would seem sensible to try and extract ‘wall measurements’ from this data. These measurements could be matched to the map and the vehicle’s location derived. However, problems arise when more than one vehicle pose (position and orientation) could generate these measurements. For example, anywhere along a corridor would generate two identical parallel-wall measurements.

DATA ASSOCIATION

Before data from sensor on a robot can be processed, a tough decision needs to be made. Should the data be interpreted as measurements of known artefacts? Or should it be treated as measurements of new, as yet unknown, objects and environmental features? Finally there is the possibility that the data is at worst spurious and at best highly ambiguous. The problem of making the right decision is central to the navigation task and is known as the Data Association (DA) problem. This is, of course, a perception issue. The DA problem is hard, very hard. From a navigation perspective it is the biggest single impediment to the deployment of robust mobile robots. We shall return to the DA problem, after continuing the discussion of mapping and simultaneous localisation and mapping (SLAM), albeit with the ogre of the DA problem hovering in the wings.

MAPPING

Mapping is the conjugate of localisation but with fewer real-world applications. Here we are given the location of the vehicle and as the vehicle moves around we use relative measurements of the vehicle world to build an internal representation of the workspace. The constructed map may be an occupancy grid, a mesh of cells, each annotated with a probability of containing some or part of an object. A common approach is to build a map consisting of a list of locations and parameterisations of geometric features – planes, points, corners and arcs for example. Alternatively it can be the past vehicle locations themselves that constitute the map – a crumb trail through space. Figure 1 shows such a map created using a technique discussed in the next section, SLAM. The sensor data, in this case laser data, is used in its raw unprocessed form and the vehicle poses are ‘adjusted’ to make the data ‘maximally align’ from pose to pose.

Figure 1: A SLAM-built map of the Information Engineering Building at Oxford University. Small triangles mark the path of the robot around the building. The map, built from scratch, develops as the vehicle moves. The same map is used for navigation – a chicken and egg problem.

SIMULTANEOUS LOCALISATION AND MAPPING

The solution to the problem of simultaneous localisation and mapping (SLAM) is in many respects a Holy Grail of research on autonomous vehicles. The ability to place an autonomous vehicle at an unknown location in an unknown environment and then have it build a map, using only relative observations of the environment, and then to use this map simultaneously to navigate, would indeed make such a robot 'autonomous'. Thus the main advantage of SLAM is that it eliminates the need for artificial infrastructures or *a priori* topological knowledge of the environment.

Solving the SLAM problem would be of inestimable value in applications where absolute position or precise map information is unobtainable. These problems include, among others, autonomy in planetary exploration, sub-sea vehicles, air-borne vehicles and all-terrain vehicles in tasks such as mining and construction.

The SLAM problem has been the subject of substantial research since the inception of research in mobile robotics, indeed before this in areas such as manned vehicle navigation systems and geophysical surveying. A number of approaches have been proposed to address the SLAM problem. One of the most popular employs a common estimation tool called a Kalman filter, which is simply a recursive, online algorithm for applying probabilistic inference when the probability distributions involved are Gaussian.

This approach is popular for two main reasons. Firstly, it directly provides both a recursive solution to the navigation problem and a means of computing consistent estimates for the uncertainty in vehicle and map landmark locations on the basis of statistical models for vehicle motion and relative landmark observations. Secondly, the autonomous-vehicle community can draw on a substantial body of method and experience in aerospace, maritime and other navigation applications. However, the big draw back to the approach is that the computation required to estimate simultaneously both vehicle pose and map grows with the square of the number of mapped features.

This quadratic scaling means that, notwithstanding issues of perception, the vehicle cannot operate in an arbitrarily large environment. As the map grew, the requirements of the estimator would soon swamp whatever computational resources are available. We shall return to this problem shortly.

The Role of ProbabIlity

At this point it is useful to discuss the role of probability theory applied to robot navigation. Every measurement made is subject to error – the only unknown is how much noise, and hence uncertainty, is associated with the measurement. The mathematical apparatus for dealing with uncertainty is probability theory. It comes as no surprise, then, that the manipulation and estimation of probabilities is ubiquitous in contemporary navigation algorithms. Indeed, the navigation problem, in each of the three guises mentioned above, can be thought of as the estimation of probabilities distributions over various domains.

Take, for example, localisation. We want to estimate the probability of the vehicle being in every single pose *given* all the measurements and the map. We would then be inclined to define the vehicle ‘estimate’ as the pose for which this conditional probability is maximum.

Similarly, we can think of mapping as the process of finding the most likely world description, given the vehicle’s position and measurements. Almost without exception, cutting-edge navigation algorithms use Bayesian methods to estimate the underlying probability density functions or suitable parameterisations of them, such as mean and variance. Bayes’s rule comes into play because the sensor we use can be thought of as producing samples from an underlying probability distribution – the probability of each measurement given the state of the world. Bayes’s rule allows the inversion of this to yield something more useful – the probability of the state of the world given the measurements.

The Scaling Problem

The big challenge in SLAM is overcoming the scaling problem. How do we avoid the quadratic increase in computation with map size, resulting from a naïve implementation of Bayes’s rule in a Kalman filter. The unbounded growth of computation with map size essentially prohibits large-scale sustainable mobile autonomy. Because of the difficulties encountered by SLAM algorithms when applied to larger environments, map scaling is a key issue for research in this area. Over recent years, there has been great progress in this area.

Davison’s and Knight’s [U. Oxford] ‘postponement method’ and later Guivant and Nebot's [U.Sydney] ‘compressed filter’ allow computational resources to focus on maintaining a representation of the local area, postponing the computation required to manage all other mapped features. However, eventually the same ‘quadratic-in-map-size’ computation must be completed – once again placing a limit on the size of environment in which the algorithm can operate. In a similar vein, the constrained sub-map filter [Williams, U. Sydney] and the Geometric Projection Filter [Newman] seek to delay the full computation. Other techniques such as decoupled stochastic mapping [Leonard, MIT] and Sparse Extended Information Filters [Thrun, Stanford] achieve constant time performance, but make approximations that require empirical testing to verify state estimation consistency. Then In 2002, Montemerlo and Thrun of Stanford published the FastSLAM technique. It was the first consistent technique with sub-linear scaling. The cost is logarithmic in the number of features mapped. Finally, in 2003, Newman and Leonard (U. Oxford and MIT) published the Constant Time SLAM algorithm (CTS) that offered the first provably convergent and consistent SLAM at a computational cost independent of the number of mapped features.

Even with the scaling problem overcome we still do not have SLAM-enabled robots capable of substantive deployments. Why not? The remaining problem lies in perception. The Data Association is far from solved. All the SLAM algorithms mentioned require an oracle of some description to associate each measurement with either an existing feature or an as yet unmapped feature. If this process goes wrong, terrible things happen. Vehicles get lost, maps get corrupted and new features are mapped when in fact they are observing previously mapped features. Most SLAM algorithms fail catastrophically under poor data associations.

The Open Problem – Building an Oracle

The big question now is how to improve data association. One option is to maintain a competing set of hypothesis about the world. This approach is, to some degree, intrinsic to, and is one of the great strengths of, the FASTSlam approach. Another is to exploit pattern recognition techniques developed for image understanding and computer vision. Remarkably, as yet most SLAM algorithms use geometric maps and use geometric tests only to determine measurement-feature associations. It is clear that data association could be far more robust if the process included other cues – for example, texture and colour as seen through a camera *in addition* to the geometric representations derived from an active sensor such as a 3-D laser scanner. This is the motivation behind research at the Oxford University Robotics Research Group

It is in the search for good, well-engineered data-association methods using images or other data, that there could be an intersection between the robotic navigation community and life-sciences. It is clear that we rarely have problems with data-association. We can be transported blindfold to a new location and can instantly classify our new location as novel or somewhere we have been before.

What is particularly impressive is our response when presented with an unfamiliar view of an otherwise well known location, for example, recognising our homes from an aerial photograph. Achieving such very wide baseline correspondences would be of immense use in mobile robotics. It would greatly aid in the so called ‘loop closing’ problem. In this a vehicle exploring and mapping a new area suddenly recognises it has arrived back at a location that it has already mapped. The exploration stops, no new features are added, and subsequent observations of walls etc. are matched to those already in the robot’s map of its workspace. The hard part is recognising the intersection of current and past locations when the pictures painted by the raw sensor data can be so different.

VISION-BASED LOCALISATION AND MAPPING

Andrew J. Davison,

For many reasons, vision is an attractive choice of sensor for robot navigation. Cameras are compact, accurate, non-invasive and well-understood – and today cheap and ubiquitous. However recent progress in real-time robot navigation, and specifically simultaneous localisation and mapping (SLAM), has primarily built on specialised sensors such as laser rangefinders. There has been less emphasis on using cameras for SLAM, despite the strong computer-vision research community in the UK and elsewhere.

Vision also has great intuitive appeal. It is the sense that humans and animals primarily use to navigate. The reason for the slow uptake of vision in SLAM is perhaps that cameras capture the world's structure only indirectly through photometric effects. It is more difficult than with other sensors reliably to extract mappable geometry from images.

Work by Andrew Davison and colleagues in the Active Vision Laboratory at the University of Oxford has focused on demonstrating that vision can be the primary sensor in real-time SLAM problems.

The first system, developed in 1998, used a 'mechatronic' binocular camera platform, or 'active head', mounted on a wheeled robot. By detecting and serially fixating its active cameras on arbitrary salient image regions in the surroundings, the robot built a sparse map of landmarks as it moved through an indoor environment while simultaneously localising itself. This was the first correctly formulated approach to SLAM using vision.

More recently, the laboratory's visual SLAM research has concentrated on relaxing several assumptions made in this and most other SLAM systems – that robots move on flat ground planes, make controlled motions with odometry available, and use sensors that directly return both the range and bearing to features – stereo vision achieves this using triangulation.

By 2004, advances in monocular landmark initialisation, efficient feature search and dynamic motion modelling allowed progress to a single-camera SLAM system in full 3-D, capable of real-time localisation and mapping in room-sized environments. Operation is at 30Hz, with all processing on a standard Pentium processor.

CAMERA VISION

A single camera is a truly cheap, compact and flexible SLAM sensor. Weak motion modelling characterising only the bounded accelerations of smooth movement means that this algorithm is not restricted to typical wheeled-robot navigation but can apply where the camera is mounted on a wearable collar, humanoid robot, or even waved in the hand.

There are also signs that such research on the boundary of mobile robotics and computer vision is bringing the two fields closer together again, after a period in which their goals diverged. This is driven first by a desire in the robotics community to achieve robot navigation in more challenging domains.

Current challenges in robot SLAM include the move from 2-D to 3-D, operation in self-similar scenes where the basic geometry can be ambiguous and loop closing in large environments. In all of these cases, extra richness of visual data should permit progress. Secondly, developments in computer vision, particularly in object recognition and invariant feature detection, look promising for use in SLAM.

Fundamental new types of features with non-fixed scales and orientations offer the chance of efficient visual scene mapping in a much more comprehensive manner than with previous approaches. It is hoped that developments in robotics will very much interact with new vision research and aim to take visual SLAM to the next level of applicability.

SENSING AND INTERPRETATION

SENSORS AND DATA FUSION

Huosheng Hu and John Q. Gan

Introduction

Without sensors and data-fusion algorithms, the development of autonomous and intelligent robotic systems remains in the realms of science fiction. Sensors provide a robot with the ability to sense its environment and to handle environmental uncertainty. Data-fusion algorithms reduce the inaccuracy of data from sensors and eliminate possibly false sensory information.

Sensor Technology

Many different types of sensors have been developed in robotics, driven mainly by the need to deploy mobile robots in unstructured environments or to coexist with humans. These sensors are embedded in robot systems. Their functions can be divided into two categorises: internal-state sensors and external navigation sensors.

Internal-state sensors mainly measure and monitor the internal states of a robot. For example, they detect its velocity, acceleration, attitude, current, voltage, temperature, pressure, balance and attitude so that static and dynamic system stability can be maintained, and potential robot failures detected and avoided. There are two kinds of internal-state sensors, contact and non-contact sensors.

CONTACT SENSORS

Contact-state sensors detect if a robot is touching something. They involve direct physical contact with the objects of interest. Such sensors include micro-switches, touch, force, tactile and potentiometers. Contact sensors are typically employed on robotic manipulators to handle objects, reach a specific location, or to protect the robot from colliding with obstacles.

Contact-state sensors are cheap, have a fast response and are easy to construct and operate. For instance, micro-switches can be attached to the robot's grippers or 'hands' to operate as a binary sensor to detect the presence or absence of an object. Tactile sensors can adopt different technologies, such as magnetic, capacitive and piezoelectric effects. A 2-D tactile sensor can provide information on size, shape and position of an object. This technology is beginning to mature and many commercial devices are available.

However, contact state sensors provide limited information. They also have a short life span because of their frequent contact with objects. Also, in many circumstances it is difficult to develop effective algorithms to interpret data from these sensors.

NON-CONTACT SENSORS

Non-contact state-sensors include synchros, resolvers, proximity, accelerometers, tilt, compasses, gyroscopes, optic encoders and so on. Since they are not designed for contact with any object of interest, these sensors have, in theory, unlimited life.

Synchros and resolves are rotating electromechanical devices that measure angular position information with great accuracy. They are normally very large and heavy devices.

Gyroscopes have two types, mechanical and optical. Mechanical gyroscopes operate on the basis of conservation of linear or angular momentum. In contrast, optical gyroscopes have no moving parts and are virtually maintenance free. Mechanical and optical gyroscopes are widely used to maintain the stability and attitude of robotic systems, especially useful to unmanned flying robots, underwater vehicles and space robots that navigate in a 3-D environment.

Optical encoders are the most popular non-contact sensors in mobile robots. There are two types of optic encoders: absolute and incremental. Absolute optical encoders normally measure and control the angle of the steering wheel in a wheeled robot for path control. In contrast, incremental optic encoders measure and control speed and acceleration in a mobile robot

EXTERNAL NAVIGATION SENSORS

The purpose of external navigation sensors is to measure and abstract the environment features – e.g. range, colour, gap and road width, room size, object shape, etc. – so that the robot can correct errors in the world model, detect environment change, and avoid unexpected obstacles. External navigation sensors could be roughly divided into two types, vision-based and non-vision based sensors.

Non-vision based navigation sensors

Various types of non-vision navigation sensors are based on different physical principles, such as force, magnetic, sound, smell, infrared, optic, acoustical, laser, radio frequency, proximity, satellite and radar. Among these force and magnetic sensors are passive navigation sensors that do not generate electromagnetic waves. Most non-vision based sensors are active sensors. They emit some kind of energy that travels between transmitters and receivers, and thus clutter the environment.

Both force and magnetic sensors have proven very reliable and produce little noise. The data measured from these passive sensors is normally easy to understand and to interpret. For instance, force sensors can monitor whether there is an object in contact with the robot for the purpose of collision avoidance.

Active sensors measure the change in emitted-energy properties, such as frequency and phase, or make simple time-of-flight calculations. These sensors normally provide 1-D data at a high rate, and demand less computing power than vision-based sensors. Active sensors have been widely used for different navigation tasks of mobile robots such as following a safe path, reaching a goal, avoiding obstacles and mapping an environment.

Ultrasonic or sonar sensors are cheap and easy operation makes them popular in mobile robotics. They also provide direct range information which is very convenient for many real-world applications in which range information is essential.

A problem with ultrasound sensors is false range readings caused by specularities resulted from the long wavelength of sound waves. Many algorithms have been developed to reduce uncertainty in sonar range readings such as EERUF (Error Eliminating Rapid Ultrasonic Firing) algorithms developed by Borenstein and Koren. Since EERUF has a fast sampling rate, a mobile robot can travel safely in a densely cluttered environment at a high speed of 1 m/s.

Satellite based GPS systems and RF positioning systems are widely used in robot localisation, object and human tracking. The use of GPS in outdoor localisation of mobile robots is common as navigation satellites are always available. The absolute 3-D position of any GPS receiver is determined through simple triangulation based on time-of-flight radio signals that are uniquely coded and transmitted from the satellites.

The main problems of GPS systems include:

- time synchronisation between satellites and the receiver;
- precise real-time location of satellites;
- difficult to measure signal propagation time; and
- electromagnetic noise and other interference.

A possible solution is to integrate GPS and other navigation sensors, such as inertial sensors in the system to fuse data and reject noise.

The use of odour by insects has motivated researchers to develop odour detection to assist robot navigation. In earlier work, Deveza, et al. developed an odour sensing system that allows a mobile robot to follow trails of volatile chemicals on the floor. An improved odour sensing system was developed. In this the sensor draws odour laden air over the sensor to increase its response speed. Odour markings on the floor can then be reliably detected and accurately localised. It remains to be seen if odour sensing can be effectively applied to useful robotic tasks, especially in hazardous environments and disaster rescue sites.

Vision based navigation

Robots must be able to see if they are to perform specific tasks such as assembly, inspection and recognition. Vision based navigation sensors, one of the most powerful sensors in mobile robotics, mimic our eyes and can provide huge amount of information. However, the visual information obtained from a vision sensor needs three processing stages: image transformation; image segmentation and analysis; and image understanding. These are extremely time consuming and difficult to achieve in real time in many circumstances, especially when colour image data is considered.

In general, vision sensors in mobile robots are active or passive charge-coupled-devices (CCDs). Active vision uses some kind of structured lighting to illuminate the scene and to enhance the area of interest to speed up image processing – only the data in the enhanced area is processed. For example, by projecting a pattern of light strips into the scene, the depth information can be obtained by looking for discontinuities or deformation in the resulting line image. Researchers at NASA's Jet Propulsion Laboratory used a light striping system to avoid obstacles in the Sojourner Mars rover, with five lines of light strips projected ahead of the rover to detect unexpected obstacle.

In contrast, passive vision works under normal ambient illumination and has to deal with difficult problems associated with shadows, intensity, colour, texture and specularities. The large quantity of image data has to be processed in real time to abstract useful information for navigation.

Passive vision sensors have been widely used in both indoor and outdoor navigation for different kinds of mobile robots, such as wheeled, legged, tracked, underwater and flying. They are mainly used for object and landmark recognition, line following, and goal seeking. For instance, Saripalli et al. adopted a vision sensor to implement visually guided landing of an unmanned helicopter. In their system, the helicopter updates its landing target parameters based on vision data and uses on-board behaviour-based controllers to follow the path toward the landing site and land on the target with 40-cm position error and 7 degrees orientation error.

Minten et al. developed a docking behaviour for a pair of mobile robots, mother and daughter, based on low-level-complexity vision sensor and an unobstructive artificial landmark. The vision sensor directly detects and recognises the landmark on the mother robot so that a micro rover can return to its mother body from an approach zone with a 2-m radius.

The Centre for Vision Speech and Signal Processing (CVSSP) at the University of Surrey is one of the UK's largest research groups devoted to analysis and understanding of sensory data. It is a major international player in research in computer vision and pattern recognition and has participated in a series of EU funded research projects in the cognitive systems areas. These have investigated image interpretation systems, researching the means and methods of representing prior and/or contextual knowledge together with principled techniques for combining expectations with incoming sensory data to yield consistent and persistent world interpretations. This activity continues with projects into video surveillance and intelligent video database search.

A related series of projects has pioneered techniques for integration of multi-camera data with the aim of constructing visually and geometrically accurate models of objects and environments. Recent work has been targeted for use on mobile robots.

SENSOR INTERPRETATION

Different sensors provide different kinds of information, including range, size, shape, colour, angle, force, etc. Therefore, each sensor obtains only a partial view of the real world, and reacts to a certain stimulus from it. It has been a challenge for robotics researchers to develop algorithms to interpret sensory data before they can be used in control and navigation of mobile robots. There are a number of challenging issues in this process.

Better sensors

No sensor works well in all situations. Every sensor suffers some drawbacks. The performance of sensors may degrade after a limited life span. Many sensors have been developed for indoor navigation and have limited range and resolutions. For instance, although sonar sensors are widely used in mobile-robot navigation, robotics researchers have to spend tremendous efforts interpreting sonar data to abstract useful information and separate it from noise. In contrast, a SICK-laser scanner (made by a company called SICK) produces range data with much higher angular and range resolution, which can be easily interpreted and used for navigation tasks.

A crucial task is to develop better sensors for robot navigation. As mobile robots move outdoors, underwater and into space they demand more advanced sensors with a long range and that can work well at extreme conditions.

Better sensor models and data fusion

To interpret sensory data, the operation of the sensor as a function of its surroundings must be understood, it must be modelled. Sensor models should be based on adequate data sets sampled by sensors in the real world. Basically, sensor models present a useful description of sensor's abilities and limitation, such as accuracy and variance. Sensor models may have to be adaptive if sensors are to operate in diversified environments that change over time. Some kind of learning is required to achieve autonomous adaptivity. It remains a challenge for robotics researchers to build good and adaptive sensor models which could be updated in real time and handle non-linearity.

Different sensors provide different kinds of information that must be ‘fused’ to obtain a complete picture of the real world. More specifically, multi-sensor data fusion aims to overcome the limitations of individual sensors and produce accurate, robust and reliable estimate of the world state based on multi-sensory information.

OPEN PROBLEMS

Different sensors provide different kinds of information. No sensor works perfectly in all real-world applications. How to effectively utilise the positive side of each sensor and avoid its negative side becomes critical for the deployment of mobile robots in the real world. To reach this goal, sensor technology and data fusion algorithms have been a hot research topic and have played a key role in the acquisition of more accurate and reliable information for the past two decades. However, a number of open problems in both sensor technology and multi-sensor data fusion algorithms remain to be answered. We list only some of them here.

Handling non-linearity and ambiguity

Bayesian methods provide a rigorous general framework for dynamic-state estimation problems such as positioning based on sensory information in robotics, in which the probability density function (PDF) or ‘posterior’ of the robot’s location is constructed based on all the available sensor measurements. If the sensors are linear – a change in the world, say distance to a wall or wheel speed, causes a proportional change in the sensor output – then tools like the ubiquitous Kalman filter can combine multiple measurements and come up with a single, optimal estimate of vehicle state.

However, the real world is overwhelmingly non-linear, and linearisations are required before tools like the Kalman filter can be used again. The resulting estimator is called the extended Kalman filter or EKF. These linearisations impinge upon both robustness and accuracy or estimation. Alternative solutions are required. For example, a series of neuro-fuzzy approaches have been proposed at Southampton University to local linear modelling and neuro-fuzzy Kalman filters for non-linear state estimation and multi-sensor data fusion. Although it has been demonstrated that the neuro-fuzzy Kalman filter has better performance than the EKF, almost all the EKF-like methods failed in some state estimation problems with strong non-linearity, such as the bearing-only target tracking problem.

Another crucial property of Kalman-based estimators is that they implicitly assert that the world is unambiguous. They produce a single estimate of vehicle state. Although this seems attractive at first glance, it soon becomes problematic. The real world is ambiguous and frequently more than one the sensor data could be ‘explained’ by more than one vehicle location. For example, imagine trying to estimate vehicle location from measurements to two orthogonal walls in a room of known dimensions. The measurements say “there is a wall *x* metres away and another *y* metres away and perpendicular to the first”. In terms of estimating vehicle location, this means there is disjoint set of likely vehicle locations near each corner of the room. However, ‘uni-modal’ approaches like the Kalman filter cannot represent this multi-modal distributions. This is a real problem.

Gordon et al. proposed a new way of representing the PDF of the state. They developed a bootstrap filter to recursively update the PDF representation. This was shown to be far superior to the standard EKF in the bearing-only target tracking problem. The power of this filter stems from the key idea of representing the required PDF as a set of random samples (particles) with importance weights, rather than as a function over the state space, which is updated recursively by updating these samples and weights. The above idea and method, now commonly known as particle filters, have drawn much attention recently in non-linear state estimation.

Research on particle filters for coping with the non-linear and non-Gaussian situations in mobile robot positioning and navigation is being carried out at Essex. The increasing role of particle filtering in autonomous navigation is discussed further in the Navigation section earlier in the document.

Characterisation of uncertainty

Characterising the uncertainties in sensor measurements remains a challenge. This is firstly because there is no general analytical solution to non-linear and/or non-Gaussian situations, and secondly because in many practical applications the environment and sensor working conditions vary in time. Monte-Carlo methods provide a novel approach to non-Gaussian distribution approximation. Multiple models plus adaptive model switching methods provide a divide-and-conquer approach to handling complicated uncertain situations. Fuzzy reasoning as a general tool for coping with uncertainty could be useful in characterising sensor uncertainties.

Fuzzy local linearisation (FLL) has emerged as a useful approach to non-linear process modelling. It distinguishes itself from traditional piecewise linearisation by fuzzy (soft) input space partitioning. Although new methods have been developed for crisp input space partition in piecewise linear modelling, applications are restricted due to the inherent ambiguity or fuzziness in the input space partitioning based upon its local linearity. FLL provides a potential way to resolve this problem.

In an FLL model, local linear models are constructed on local regions generated from the input space partition by fuzzy sets and are combined by membership function weighting. In general, the membership functions that define the fuzzy sets and the corresponding local linear models need to be identified, using optimisation techniques such as least squares (LS) and least mean squares (LMS), based on observational data and/or fuzzy rules.

Expectation maximisation (EM) is a general technique for maximum likelihood or maximum *a posterior* estimation. It has become an alternative to LS and LMS techniques in solving many estimation problems as the EM technique can provide covariance information about model mismatch.

LS, LMS and EM are frequently used for local model identification. For input space partitioning or membership function construction, evolutionary or growing-pruning algorithms based on optimal criteria, such as structural risk minimisation (SRM), have been developed. For instance, the adaptive spline modelling (ASMOD) algorithm based on the analysis of variance (ANOVA) decomposition has been widely used in spline models to combat the curse-of-dimensionality in high-dimensional system modelling.

In many applications where a priori knowledge is insufficient, it is highly desirable to automatically partition the input space using fuzzy sets in an effective (parsimonious) manner. Aiming at resolving this problem, researchers at Southampton decomposed an FLL model into sub-models in an ANOVA form and developed a modified ASMOD (MASMOD) algorithm for automatic fuzzy partitioning of the input space, including automatic determination of the number of local regions. This model construction algorithm essentially decomposes the underlying non-linear process into an additive set of low dimensional sub-models which are individually parameterised, avoiding the curse-of-dimensionality. However, it remains to be seen how it can be effectively used for robot navigation.

Data fusion

In most real-world robotics applications, the environment is uncertain and dynamically changing. Adaptive multi-sensor data fusion is essential in these applications. Adaptive techniques are required to decide which sensors should be involved in sensor fusion and which fusion method to adopt. A key issue is when and how adaptation should take place. This requires effective performance feedback signals. Due to the efficiency of the EKF-like methods, they are still the main approach to the complex positioning and navigation problem of mobile robots or autonomous vehicles, such as the SLAM problem.

However, there are realistic problems – such as mobile robot navigation with initial robot position unknown or global position estimation – in which the EKF-like methods offer very poor performance. Compared to the EKF-like methods, particle filters do not need the initial robot position and the linearisation of the state-space model.

To the best of our knowledge, there is no work in the UK on particle filters for mobile-robot localisation. Unsolved problems in this area include; how to choose and adapt the number of required samples, how to avoid the collapse in the number of distinct values in the sample set, and how to reduce the computational cost.

To overcome the vulnerability of centralised data fusion systems, decentralised data fusion architectures have been investigated In a decentralised architecture, there is no global fusion centre. Local fusion centres provide information with equal importance. Decentralised architectures are flexible, fault-tolerant and suitable for parallel implementation. However, complicated communication and control mechanisms are needed to coordinate the sensor network – both wired and wireless.

CONCLUSIONS

Sensors, and the algorithms used to interpret them, play an important role in the development of autonomous and intelligent robotic systems. A huge variety of sensing technologies have been developed, both in the UK and abroad. No sensor is perfect and suits all the applications. We expect robotics researchers to develop better sensors and better data fusion algorithms. We predict not just a quantitative expansion of different kinds of sensor technologies, but also qualitative advances.

REFERENCES

R. Deveza, R.A. Russell, D. Thiel, and A. Mackay-Sim, Odor Sensing for Robot Guidance, International Journal of Robotics Research, Vol. 13, No. 3, 1994, pages 232-239.

J. Borenstein and Y. Koren, Error Eliminating Rapid Untrasonic Firing for Mobile Robot Obstacle Avoidance, IEEE Transactions on Robotics and Automation, Vol. 11, No. 1, 1995, pages 132-138.

S. Saripalli, J.F. Montgomery and G.S. Sukhatme, Visually Guided Landing of an Unmanned Aerial Vehicle, IEEE Transactions on Robotics and Automation, Vol. 19, No. 3, 2003, pages 371-380.

B.W. Minten, R.R. Murphy, J. Hyams and M. Micire, Low-Order-Complexity Vision-based Docking, IEEE Transactions on Robotics and Automation, Vol. 17, No. 6, 2001, pages 922-930.

N.J. Gordon, D.J. Salmond and A.F.M. Smith, Novel Approach to Non-linear/Non-Gaussian Bayesian State Estimation, IEE Proceedings-F, Vol.
140, No. 2, 1993, pages 107-113.

C.J. Harris and Q. Gan, State Estimation and Multi-sensor Data Fusion using Data-based Neuro-Fuzzy Local Linearisation Process Models, Journal of Information Fusion, Vol. 2, No. 1, 2001, pages 17-29.

H. Hu and D. Gu, Concurrent Navigation and Map Building for Mobile Robots, Proc. of CACUK'98, pages 207-212, Leicester, ISBN 0 9533 8900 6, 19-20, September 1998

H. Hu and D. Gu, Landmark-based Localisation of Industrial Mobile Robots, Int. Journal of Industrial Robot, Vol. 27, No. 6, pages 458 - 467, November 2000

SENSOR INTERPRETATION VIA ABDUCTIVE REASONING

Murray Shanahan

Research in *cognitive robotics* aims to endow robots with high-level cognitive skills, such as planning, reasoning about other agents, and reasoning about their own knowledge. The field has its roots in classical artificial intelligence (AI), and deploys the traditional AI concepts of representation and reasoning.

Building on his earlier research on reasoning about action using a logic-based formalism called the *event calculus*, Murray Shanahan developed a formal, theoretical framework for robot perception, which pins down the process by which raw sensor data is transformed into meaningful symbolic representation. This account, casts perception as a form of *abductive* inference. Shanahan's early work (1996) along these lines was carried out in the context of mobile robot map-building. Since then, this theoretical framework has been considerably extended, partly in the context of a series of projects, funded by the Engineering and Physical Sciences Research Council, at Imperial College.

The first of these projects, Cognitive Robotics, demonstrated the feasibility of high-level robot control through logic programming. The architecture developed on the project combined sensor data assimilation, hierarchical planning, and reactivity in a single framework based on abductive logic programming.

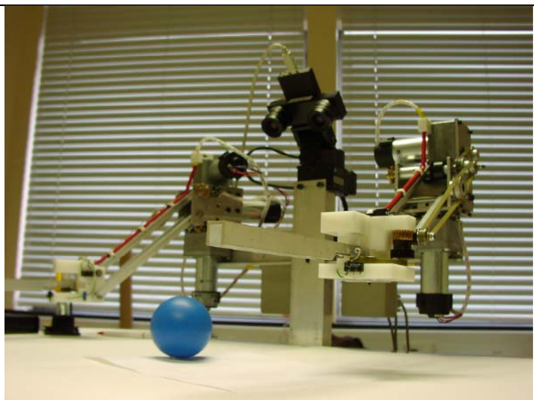


LUDWIG in the laboratory and at the Royal Institution

Abduction is often characterised as "inference to the best explanation", and is the process of finding the hypothesis (or set of hypotheses) that most satisfactorily accounts for a set of data. In terms of mathematical logic it is the inverse of deduction.

In the follow-on project, Cognitive Robotics II, the emphasis was on scaling up the earlier results and applying them to richer sensory modalities, in particular to vision. With the appointment of David Randell, whose work on qualitative spatial reasoning has been highly influential in the knowledge representation community, the Imperial College research group could explore the interface between robot perception and spatial reasoning. This led to original work on two fronts: a logic-based calculus for reasoning about spatial *occlusion*; along with methods for automated reasoning about the relations it deploys.

The earlier abductive account of sensor data interpretation was extended to handle *top-down information flow* through expectation. As in the preceding work, logical abduction was used to form a set of initial hypotheses to explain the low-level image data. But in the new approach, a deductive component subsequently computes the expectations of each competing hypothesis. The raw image is then re-consulted to check which expectations are fulfilled, resulting in the confirmation or disconfirmation of each hypothesis. A logic programming implementation was also developed that has demonstrated the computational feasibility of the technique.

In the ongoing project, Spatial Reasoning and Perception in a Humanoid Robot, this work on top-down information flow is being further extended to accommodate *active perception* in the context of an upper-torso humanoid robot named LUDWIG (see above), that can nudge objects to see them from different angles. Again, this is facilitated by the top-down influence of a reasoning component, which uses the influence of its actions on the scene to confirm or disconfirm an ongoing interpretive hypothesis.

REFERENCES

Shanahan, M.P. (2002). A Logical Account of Perception Incorporating Feedback and Expectation. In Proceedings 2002 Knowledge Representation Conference (KR 2002) (pp. 3–13).

Shanahan, M.P. & Randell, D. (2004). A Logic-Based Formulation of Active Visual Perception. In Proceedings 2004 Knowledge Representation Conference (KR 2004) (pp. 64–72).

Shanahan, M.P. (2005). Perception as Abduction: Turning Sensor Data into Meaningful Representation. Cognitive Science, to appear.

BIOROBOTICS

ROBOTS AS BIOLOGICAL MODELS

Barbara Webb

Introduction

In the following, I will briefly review the current state of biorobotics, that is, robotic systems built with some intent to test or explore hypotheses about the functioning of biological systems.

There is a long history of attempts to build robotic machines inspired by biological behaviour, including Greek automata, the clockwork devices of the 17th and 18th centuries, and early work in cybernetics such as Grey Walter's celebrated tortoises.

Biology can be a source of existence proofs for the kinds of capabilities we would like to achieve in robotics, and a source of ideas for how we might achieve them, (e.g. Ayers et al 2002, see also special issues on this theme in Philosophical Transactions of the Royal Society: Mathematical, Physical and Engineering Sciences 361 (1811); Robotics and Autonomous Systems 30 No 1-2, International Journal of Robotics and Automation November 2002). We should keep in mind, however, that solutions from nature may not always be the best alternative from an engineering perspective. Sometimes we may extract a few useful principles but the implementation may be very distant from biology; alternative solutions may work better than any biologically inspired option.

The constraints determining biological designs – such as evolution, development and self-sufficiency – can differ substantially from those driving robot applications. Nevertheless, for those capabilities of animals that we cannot yet emulate, it is reasonable to look to biology for solutions.

Another view of these systems is that they are physical models of animals that we can use to address specific questions in biology. Like more conventional models and simulations in biology, they help to enforce clarity and specificity in hypotheses by requiring all the details to be made explicit. These models also aid in determining the consequences and predictions of complex hypotheses. A unique feature of robotic implementations is that they can test hypotheses under more realistic bodily and environmental conditions, helping to characterise the problem that the animal actually needs to solve, and the plausibility of proposed solutions.

Testing control algorithms

Although it seems that the interaction of biological and physical sciences is often characterised as revolving around neuroscience, many of the most effective interactions are between robotics and behavioural biology, which can overlap with psychology. There are many animal systems for which the understanding of neural mechanisms is limited, yet the behaviour has been studied in some depth, and can be addressed and tested at the level of possible control algorithms, or as dynamical system of interaction between the animal and its environment.

The Sahabot, Saharan Ants and Robots, studies provide a good example of this kind of research (Lambrinos et al. 2000). The work aimed to reproduce the homing behaviour seen in desert ants Catyglyphis. These animals have been the subject of extensive study because of their impressive ability to use path integration and visual homing to return rapidly to their nests after long excursions. In path integration, they use the polarization pattern of light in the sky as a compass cue.

The first Sahabot robot replicated this sensor mechanism and showed that it could be used continuously in updating the home vector, improving on the 'scanning' hypothesis that had been previously suggested for the ant. The second Sahabot adopted the 'snapshot' model of visual homing proposed for insects by Cartwright & Collet, and showed that it could be effectively used to take current visual and derive a movement direction towards a remembered scene.

Further consideration of this mechanism led to a more efficient version of the algorithm that stored an average landmark vector rather than a snapshot. This mechanism was implemented in hardware on a third robot and shown to reproduce the search behaviour of bees when landmarks are moved.

The interest here is that one motivation for this hardware implementation was to assess the plausibility of the proposed algorithm as a system that could be implemented in the animal's nervous system. In general, it seems a good strategy, if the intention is to tie the results back to biology, to keep in mind such plausibility issues, even if no explicit neural mechanism is being suggested. For example, the behaviour should not rely on a solution that requires highly accurate calculations.

Infant Models

At the other end of the biological scale, there have been similar studies of algorithmic methods aimed at replicating the shared attention mechanisms of human infants. For example, Scassellati (2001) has implemented Baron-Cohen's model of the development of joint attention processes in infants, using a humanoid robot, Cog. This is a very high level model, in which eye contact and gaze-following lead to the development of shared attention and thus 'theory of mind'. The robot implementation needs to spell out all the underlying capabilities, such as the visual processing necessary to recognise faces, extracting the angle of gaze, extrapolating the gaze to the object of interest, and re-orienting the robot's gaze. This raises useful issues, such as the potential difficulty in determining the distance of the object along the line of gaze. This turns out to correspond to a staged development of competence at gaze-following in the child.

An important issue in this work is that of 'underdetermination'. Even if we can reproduce the behavioural capability of the animal with a robot implementation, this does not prove that the animal actually uses the same means to obtain this behaviour. As a consequence, we require appropriate experiments to determine how well the behaviour matches, and if it can predict new observations. It is also productive to compare different hypotheses. We can often learn, as much if not more, from a robot model that fails to reproduce the behaviour. This points to essential gaps in our knowledge of the biological system.

Neuroanatomy

Although the work I have described so far has been characterised as operating at a level above any neural detail, there is a continuum between this and the work I will now describe. In this, the suggested algorithms might be partially based on known neural properties, or mapped on to possible functional roles for certain neuroanatomical constructs, but without the intention of detailed copying of the neurophysiology. Some examples of this kind of work are robot controllers that draw on the 'mirror neuron' concept, and the hypothesised role of the basal ganglia in behavioural choice.

Much human capability is acquired through learning by imitation. Thus for some time researchers have pursued the idea of having robots learn by imitating a human demonstrator.

A possible neurological underpinning for this capability was discovered when it was observed that specific neurons in the premotor cortex would fire both when a monkey performed an action and when it saw another agent perform the action. More generally, motor areas in the brain are implicated in the process of recognising actions – in both primate neurophysiological studies and in studies on humans with positron emission tomography (PET) and functional magnetic resonance imaging (fMRI). There are now several robot implementations that attempt to replicate this functional architecture, at different levels of abstraction, and to use them to perform imitation in humanoid robots (e.g. Demeris & Johnson, 2003; Elshaw et al. 2004).

Another range of robot models emulates the place-cell and head-direction cell systems of the rat's hippocampus. In all these systems there are two aspects to the work: neuroanatomy may inspire a new approach to the problem; or the robot implementation may suggest new interpretation of the biological data. However, a number of different implementations can all claim to be consistent with the neurophysiological data, hence it can be difficult for work at this level to close the loop between robotics and biology.

Identified neuron circuits

While in some cases in the previous section, there is quite extensive physiological data, the precise neural circuitry that leads to the observed physiological properties is often speculative. The models incorporated into robots focus on copying the functionality with an abstract neural representation. This can be compared to other systems – typically for initial sensory processing and/or in simpler animals – where the actual neural circuitry is known and can be copied in some detail, to directly determine whether it does support the hypothesised functions.

One example of this is work on the 'looming detector neuron' LCMD in the locust. This single neuron is well characterised anatomically and physiologically. The specificity of its response to expanding retinal images is believed to be a function of two inhibitory mechanisms, one comprising lateral connections between adjacent motion sensitive units, the other a feed-forward inhibitory response to overall motion. Tests of this system in a robot have shown that the model avoids collisions (Blanchard et al 2000). The results have been useful both in providing real-world validation for the model and in suggesting a useful device for collision avoidance in transport systems.

Another example of an interaction between detailed knowledge of a biological system and robotics is the development of robot models to investigate the sound-localising behaviour of the cricket (Reeve & Webb, 2003). Again this is a well explored system in biology, with many of the critical neural connections mapped. Thus the underlying circuit that enables the cricket to recognise and move towards sounds can be implemented in a robot and tested under the same stimulus conditions, including outdoors. Results to date have shown that dynamic synapse properties are likely to play a significant role in the tuning of the circuit to particular temporal patterns.

One way in which work in this area is moving forward is the attempt to combine different biologically inspired systems, to obtain multimodal capabilities on the robot. For example, the sound localisation behaviour has been combined with visually mediated correction mechanisms. This raises many interesting issues, such as whether the different behaviours can be smoothly integrated or whether they potentially interfere with one another, and what mechanisms might be found in biology to solve these problems.

In human neurobiology, we can investigate few systems in this detail. However, one example of a well understood pathway is the vestibular ocular reflex, in which inertial sensors in the semicircular canals detect head movement and cause rapid compensatory eye movements to maintain a stable gaze. This capability has been replicated on a number of robot systems.

Understanding sensor and stimulus properties

Although many of the systems I have described might also be investigated using simulations of the neural circuits, a robotic approach is particularly suitable for addressing physical aspects of the sensory and motor interaction of biological systems with their environments. One example of this is the examination of the properties of whisker sensors.

Whiskers are important sensors that allow rats to navigate in the dark. The animals can recognise features and textures from the patterns of vibration when actively 'whisking' their environment. Several groups are investigating these capabilities in robot systems. For example, they are addressing the issue of the layout of whiskers that provides the best information, and how to obtain texture discrimination.

Antennae are a related form of sensor, but with several more sophisticated features such as fine motor control. Antennae also have many different kinds of tactile receptors, as well as chemical receptors. Some of these features have been used in robot antennae, but there is still much room for development.

Another way of considering the capability of physical sensors is to test the nature of the stimulus source with a robot system. For example, attempts to emulate chemical tracking in animals have helped in understanding the nature of the plume and what directional information can and cannot be extracted from it. This has been seen both in underwater systems such as the RoboLobster and in walking/flying systems such as robot moths.

Understanding actuators and substrate properties:

As well as enabling the exploration of real sensory interactions with environments, robot implementations of biological models also require actuators that deal with real substrates. Although many of the systems discussed thus far use simple actuation (i.e. wheels) to broadly copy the body motion of the animal, some are nevertheless tested on the same environmental substrate. There is also substantial work under way on more biological forms of locomotion, including swimming, flying, worm and snake motion, and walking.

I will briefly describe some recent examples of walking robots for which the design and control is largely biologically based. Important issues include materials, mechanical design, actuation, sensing and control.

The Sprawlita robots, based on principles of cockroach walking, are a good illustration of this combination of issues (Cham et al 2002). The hexapod system is designed to have a self-stabilising posture, with a wide base of support and low centre of mass (the 'sprawl'). The leg movements operate as thrust of prismatic actuators, with the angle of thrust determined by hip joint actuation to produce acceleration or deceleration.

A feature of the system is the use of shape deposition manufacturing using viscoelastic materials to obtain compliant actuation. This enables the system to cope mechanically with traversal of rough terrain using a relatively simple feed-forward control scheme.

This uses the principle of the spring loaded inverted pendulum that has been proposed as a common model for running in a variety of animals. The behavioural dynamics of this robot can be compared to the cockroach, with the similarities and differences providing useful information.

There are other examples of six-legged robots closely influenced by biology. There is less work on four-legged robots. The best known is probably Sony's toy robot dog Aibo. Another interesting example is Tekken (Fukuoka et al. 2003). This uses the coupled dynamics of the neural system and the mechanical system to successfully generate adaptive gaits in a cat-like robot. The neural system uses biologically based central pattern generators (CPGs) and sensory reflexes. A virtual spring-damper system mimics muscle properties in actuating the limbs, and the oscillations this produces are mutually entrained with the CPG oscillations. Some reflexes are implemented mechanically, such as the flexor reflex at the ankle joint. Others, such as the response to pitch and tilt, are detected by appropriate sensors and used to influence the CPG output appropriately.

There is also extensive research into bipedal robots, some humanoid and some resembling other animals, which also draw on, and aid the understanding of, biological walking. Actuation in robots can be for manipulation as well as locomotion, and there is also much ongoing interaction in the biological and robotic study of arms and hands .

Robots as stimuli in animal experiments

An additional area of research that involves the interaction of robots and animals is the investigation of interactions between animals and robots. In this case, the robot's behaviour – which may or may not be controlled by biologically inspired algorithms – serves as a controllable stimulus cue for the animal.

Two simple examples are the use of a robot bee to test the dance communication of honeybees and the use of a robot female bowerbird to elicit courtship dances in male bowerbirds. Work on the robot sheepdog aimed to invert flocking algorithms to derive robot movements that could herd a flock of ducks (Vaughan et al 1998). In recent research, the behaviour of cockroaches has been studied using a miniature robot. In this case, one aim is to find the physical and behavioural features that enable real cockroaches to react to the robot, as though it was one itself (http://leurre.ulb.ac.be/Descript.html).

Again, at the other end of the scale, researchers use robots as stimuli in human experiments, for example, they examine what cues appear to be lifelike, or lead to natural interactions such as pointing, talking and so on. This forms an important part of the emerging field of human-robot interaction (see Robotics and Autonomous Systems vol. 42, issue 3-4, special issue on this theme).

REFERENCES

Ayers, J., Davis, J.L. and Rudolph, A. (2002) Neurotechnology for Biomimetic Robots. MIT Press

Blanchard, M., Rind, F.C. & Verschure, P.F.M.J. (2000) Collision avoidance using a model of the locust LGMD neuron. Robotics and Autonomous Systems. 30:15-36.

Cham, J. G., Bailey, S. A., Clark, J. E., Full, R. J. and Cutkosky, M. R (2002) Fast and Robust: Hexapedal Robots via Shape Deposition Manufacturing. International Journal of Robotics Research. 21:??

Demeris, Y. & Johnson, M. (2003) Distributed predictive perception of actions: a biologically inspired robotics architecture for imitation and learning. Connection Science, 15:231-243.

Elshaw M., Weber C., Zochios A., Wermter S. (2004) An Associator Network Approach to Robot Learning by Imitation through Vision, Motor Control and Language. Proceedings of the International Joint Conference on Neural Networks, Budapest, Hungary.

Fukuoka, Y., Kimura, H. and Cohen, A.H. (2003) Adaptive dynamic walking of a quadraped robot on irregular terrain based on biological concepts. International Journal of Robotics Research 22:187-202

Lambrinos D., Möller R., Labhart T., Pfeifer R., Wehner R. (2000) A mobile robot employing insect strategies for navigation. Robotics and Autonomous Systems, 30:39-64

Prescott, T. J., Gurney, K., Montes-Gonzalez, F., Humphries, M. and Redgrave, P. (2002).The robot basal ganglia: Action selection by an embedded model of the basal ganglia. In L. F. B. Nicholson and R. Faulls (Eds.) Basal Ganglia VII. New York: Plenum Press.

Reeve, R. and Webb, B. (2003) New neural networks for robot phonotaxis. Philosophical Transactions of the Royal Society B 361:2245-2266

Scassellati, B. (2001) Investigating models of social development using a humanoid robot. In Webb, B. & Consi, T. (eds) Biorobotics.MIT Press Walter

Vaughan, R.T., Sumpter, N., Frost, A. and Cameron S. (1998) Robot Sheepdog Project achieves automatic animal control. in Proc.Int. Conf. Simulation of Adaptive Behaviour, Zurich, Switzerland

Webb, B. (2000). What does robotics offer animal behaviour? Animal Behaviour, 60 (5), 545-558.

Webb B. (2001) Can robots make good models of biological behaviour? Behavioural and Brain Sciences, 24:1033-1050.

THE IMPORTANCE OF EMBODIMENT

Robert Damper

The relationship between cognitive systems and cognitive science is akin to that between engineering and science. Cognitive science – at least, as the subject is generally conceived – seeks to understand high-level brain function, usually in humans, and consequent intelligent behaviour via the ‘brain as a computer’ metaphor.

There is a long history of scientists and philosophers thinking about the brain by analogy to the most advanced technology of the day. There are some reasons why the computer metaphor might not be entirely empty. Chief among these is the universal programmability of computers which appears to mirror the flexibility of the human species in solving problems set by our environment. Further, the central metaphor of cognitive science provides a bridge, or common language and purpose, between computer science and neuroscience which has demonstrably benefited both in recent years.

Cognitive systems as a subject is perhaps less firmly wed to the brain/computer metaphor. It seeks straightforwardly to exploit our emerging understanding of high-level brain function and intelligent behaviour in the design and implementation of engineering artefacts. The Foresight Cognitive Systems Project defines cognitive systems as “… natural or artificial information processing systems, including those responsible for perception, learning, reasoning, decision-making, communication and action”.

Biorobotics is a key discipline for cognitive systems. It is set apart from some components of cognitive systems by the strong biological motivation – which has the potential both to exploit the latest knowledge from biology and to contribute to the study of biology – and by the central role that embodied implementations play, as opposed to ‘mere’ simulations.

Embodiment plays at least two important roles. First, from the practical perspective, biorobots can perform real-world tasks in hostile environments, such as under the sea, in space, areas where there are hazards such as radiation, chemicals or explosives, in cramped confines and so on.

Remote tele-operation is a realistic possibility for some applications. For others – space exploration at great distances from Earth, for example – a high degree of autonomy is essential. For this we need robot cognition of a high order.

Embodiment is also important for those researchers in artificial intelligence and cognitive science – sometimes called the embodied AI movement – who hold that “intelligence requires a body”. Intelligence is about interacting with the environment, so only if an agent can act on the world to influence it, receiving sensory feedback about the consequence of its actions, can it be said to display intelligence. This school believes that “good old fashioned AI” stalled just because the concentration on symbolic, disembodied ‘brain in a vat’ systems solving purely intellectual puzzles such as playing chess is essentially irrelevant to the real biological business of surviving and prospering in the physical world. If the new AI – of which cognitive science is arguably a part – is to deliver on this promise, then biorobotics is going to be a necessary piece of the jigsaw.

OPEN QUESTIONS

So what at this point seem to be the key problems confronting biorobotics and its intersection with cognitive systems if this promise to be realised? At least some of the issues appear clear.

UNDERSTANDING BIOLOGY

If we are truly to take inspiration from biology, and to produce artificial systems with new and useful capabilities, then we need to understand how biology functions at a systems level. What is the right level of description? To take a concrete example, is the neuron and its pattern of firing the right level at which neuroscience should attempt to understand the basis of intelligent behaviour?

Roger Penrose believes this is already too high a level, and argues for concentration on microtubules. Gerald Edelman believes it is too low a level and prefers to think in terms of neuronal groups. If the level of description is right, things should fall more easily into place. If it is wrong, our biological motivations will be illusory.

The importance of embodiment, as mentioned above, is something of an article of faith in the new AI. This has to be transformed from a matter of faith to a matter of fact. What is needed is a thorough, on-going study of how bodily form affects function. Precisely how much of ‘intelligence’ can be abstracted from the possession of a body? Is the answer to that really “all of it”, as the proponents of embodied AI would have us believe?

UNDERSTANDING EMERGENCE IN COMPLEX SYSTEMS

Many workers in cognitive science and AI argue that the reductionist methodology may have served us so well in the study of physics but it is ill-suited to the study of living systems, where we need a new ‘science of complexity’. The notion is that many of the brain’s wonderful properties, and that we would like to emulate in artificial systems, are collective and emergent. They depend on the interaction of large numbers of component parts. Just as viscosity emerges as a bulk property of a liquid, where it makes no sense to talk of the ‘viscosity’ of a component molecule, perhaps intelligence, consciousness, and so on, are emergent properties of brains. Unfortunately, at our current state of knowledge, this notion is something of a mantra. It needs to be fleshed out and better understood, which brings us back to the first issue of “understanding biology” at the right level of description.

Scalability in intelligent systems

The short history of AI has taught us that small, tantalising successes come cheap. The challenge is in scaling from small successes in toy domains to real achievements in real tasks. An intriguing aspect of William Grey Walter’s work with his ‘tortoises’ in the early 1950s was their ability to avoid obstacles and to perform simple path planning in a reactive way. It is salutary to reflect that after half a century modern biorobotics has yet to progress much beyond simple obstacle avoidance.

We need to focus on more challenging domains if we are to scale up from small successes to real and meaningful tasks. Luc Steels has argued for a careful structuring of task difficulty – in terms of problems posed by the environment – as a way of gradually increasing agent capabilities in a controlled fashion. Rodney Brooks, of MIT’s Computer Science and Artificial Intelligence Laboratory, on the other hand, has called for the community to focus on harder tasks of real value, such as manipulation, and to leave behind the old toy problems.

INTERACTION AND BEHAVIOUR

SCIENTIFIC ANALYSIS OF BEHAVIOUR THROUGH MOBILE ROBOTICS

Ulrich Nehmzow

Introduction

How do Arctic terns manage to fly from one pole of the Earth to the other, and find their destination reliably? How do desert ants return home, having moved a distance of a million times their body length from their nest, in a desert with no landmarks? How do capuchin monkeys learn to perform a sequence of abstract actions? How do humans visualise their environment, orient in it and plan complicated paths? These questions have captured human interest for millennia. Providing answers continues to be the subject of intensive research and requires the analysis and precise description of the behaviour of 'behaving' agents. This is the focus of this section.

THE USEFULNESS OF ANALYSING AND MODELLING BEHAVIOUR

As an economy, we are interested in improving industrial processes, developing technologies to support production, maintenance and security of manufacturing and business processes. A prime example is automation of transportation within factories. In the past, factories used fixed installations such as conveyor belts and other rigid transportation methods. We can now use intelligent, autonomous transportation systems that require fewer modifications to buildings, can deal with certain problems (e.g. blockages) autonomously, and can quickly be re-used for different transportation tasks.

To give a second example, research in mobile robotics has also resulted in high-tech products for domestic use that are beginning to change the way we operate our houses. Autonomous lawn mowers are commercially available, as are robotic vacuum cleaners (see the section later in this report on Domestic Robots). Entertainment robots, through their ability to interact with the human and to learn, make stimulating toys, for adults as much as for children.

Thirdly, ergonomics addresses man-machine interaction and the advancement of industrial process through improved interfaces, where mobile robotics research and analysis of (human) behaviour give valuable insights. In all of these examples, a precise understanding of behaviour – precise in the sense of analysing a transparent model of the behaviour in question – forms the foundation on which progress is achieved.

THE ROLE OF MOBILE ROBOTICS

By mobile robot we mean a physical device, or 'embodiment,' that can perceive its environment through on-board sensors, can reason about the perception, and can move in its environment. Typically, there is an interaction between the mobile robot and the environment, in that the environment's features influence the robot's actions. The robot's actions may also change aspects of the environment, 'situatedness'. Robots that carry all necessary components for the execution of this process on board – power, computation, perception and locomotion – are referred to as autonomous.

Behaviour, for instance the behaviour of animals or humans in a certain environment, entails the perception of that environment through the agent's sensors (visual, tactile, auditory, olfactory, etc.), processing of the signals received (e.g. is this object threatening?, "Will I succeed in picking that object up?" etc) and actuation (e.g. control of limbs). While some aspects of this process are open loop, for example throwing a stone, the behaviour as described above is essentially a closed-loop process in which the agent continuously perceives its environment and responds to the stimuli it receives.

Autonomous mobile robotics can play an important part in analysing and modelling behaviour. The mobile robot essentially possesses all the processes the behaving agent possesses – power source, perception, computation and actuation – and offers the distinct advantage that its controls mechanisms are known and can be changed at will. There is, therefore, a long tradition in using autonomous mobile robots in the analysis of behaviour. Research goes back to the beginning of the last century [Hammond 18]. W. Grey Walter, the pioneer of modern mobile robotics as a tool in analysing behaviour, conducted experiments in machine learning and navigation, using autonomous mobile robots, in the early 1950s at the Burden Neurological Institute in Bristol [Walter 53].

To date, most robotic implementations of intelligent behaviour are 'insular' solutions, i.e. they are 'existence proofs' that demonstrate that a particular behaviour is achievable. Increasingly, however, the argument is made within the mobile robotics community that the time has come to develop a theoretical understanding of robot-environment interaction, to analyse robot behaviour quantitatively rather than qualitatively, and to model robot behaviour precisely, rather than approximately. In other words, a science of mobile robotics is emerging. The group at the University of Essex is one proponent of this view, and investigates scientific robotics in its RobotMODIC project (Robot Modelling, Identification and Characterisation).

Case Study

The RobotMODIC

The first requirement for the analysis of mobile robot behaviour is that of establishing a 'language,' a means of communicating properties of robot, task and environment precisely, quantitatively. One way to achieve this is to apply dynamical systems theory and chaos theory to the analysis of robot behaviour [Nehmzow 03, Nehmzow & Walker 03b, Nehmzow & Walker 03a]: the interaction of a mobile robot with its environment becomes "measurable" through quantitative measures such as Lyapunov exponent and correlation dimension [Kaplan & Glass 95].

Once such quantitative description of robot behaviour exists. It can be used further, most importantly, in modelling robot-environment interaction. Using polynomial models to represent input-output relationships [Korenberg et al 88], it is possible to describe a robot's perception at certain locations in the environment (simulation scenario), to express a robot's response to sensory stimuli analytically (programming scenario), to 'translate' sensor modalities from one to another (for instance, translation of laser perception in to sonar perception), or to self-localise, using sensor perception (navigation scenario).

SUMMARY

Apart from scientific interest in the fundamental processes involved in intelligent behaviour, the analysis and theoretical understanding of agent-environment interaction will allow the design of sophisticated machines to advance industrial processes (e.g. transportation, maintenance and surveillance), machines for domestic use (e.g. lawn mowers and floor cleaners), and improvements in ergonomics in general (e.g. man-machine interface).

Similarities between living beings and mobile robots – notably autonomy, sensory perception of the environment, reasoning capability and the ability to move – make mobile robots are ideal test-beds to investigate hypotheses about intelligent behaviour. The internal processes of robots are analysable and easily modified. They are relatively cheap, and pose few ethical issues when used in research. For this reason, mobile robots have been used in behavioural research for almost a century, and have gained in importance. The 'identification' of these processes – their expression as analysable mathematical functions (Robot Modelling, Identification and Characterisation) – is one step towards better understanding of robot-environment interaction and a science of mobile robotics.

REFERENCES

[Korenberg et al 88] M. Korenberg, S.A. Billings, Y.P. Liu and P.J. McIlroy, Orthogonal parameter estimation algorithm for non-linear stochastic systems, Int. J. Control, Vol. 48 No. 1, pp. 193-210, 1988.

[Hammond 18] John Hays Hammond and B.F.Miessner, An artificial heliotropic machine, in Jacques Loeb, Forced movements, tropisms and animal conduct, Philadelphia: J.B. Lippencott, 1918.

[Kaplan & Glass 95] Daniel Kaplan and Leon Glass, Understanding nonlinear dynamics, Heidelberg: Springer 1995.

[Nehmzow & Walker 03a] Ulrich Nehmzow and Keith Walker, The Behaviour of a Mobile Robot Is Chaotic, AISB Journal, Vol 1 No 4, 2003.

[Nehmzow & Walker 03b] Ulrich Nehmzow and Keith Walker, Investigation of robot-environment interaction using chaos theory, AISB quarterly, No. 114, 2003

[Nehmzow 03] Ulrich Nehmzow, Quantitative analysis of robotenvironment interaction—towards "scientific mobile robotics", Int. J. of Robotics and Autonomous Systems, Vol 44, pp. 55-68, 2003.

[Walter 53] W. Grey Walter, The living brain, London: Gerald Duckworth & Co, 1953.

HUMAN-ROBOT COMMUNICATION.

Introduction

The development of effective human-robot communication (HRC) systems is an important factor in the service robot market. Service robotics is predicted to overtake industrial robotics and to gain a dominant share of the estimated $59 billion robotics market in 2010 (Key market drivers report 2004, Future Horizon).

A service robot can be a collection of intelligent devices integrated in an intelligent house, or it can take the more classical form of a mobile robot performing a variety of tasks, from domestic help to entertainment and companionship. Effective communication between user and robot is essential to access the full capabilities of such robots. The capability for user-friendly communication will in itself contribute to the perceived value and social acceptability of the robot.

Service robotics poses specific and new problems of HRC that are very different from those in industrial robotics. In industrial robotics, human-robot interfaces are designed to enable *skilled* workers to generate and activate programs for the robot. In service robotics, robots operate in an uncontrolled environment: they have autonomy, are multifunctional, must be adaptable to the user's needs. Users of such devices are definitely unskilled in robot programming. Thus the problem becomes very much one of providing the robot with communication capabilities that are compatible with human-to-human communication in everyday situations, referring to everyday tasks.

SOCIALLY INTELLIGENT ROBOTS

Kerstin Dautenhahn

Socially intelligent robots are robots that show aspects of social intelligence and social behaviour that we can see in humans and many other social animals. Such robots can interact socially, either with each other or human interaction partners.

As an example, let us consider a robot in the home. Whether people will accept a robot in the house, and will be willing to pay a non-trivial sum of money, depends on the robot's role and how it can express this role in social interactions. To give an example: A robot in the role of a 'servant' needs to:

- possess useful skills to solve tasks that it is desirable for robot to carry out, e.g. vacuum-cleaning, doing the dishes, mowing the lawn, tidying up the children's room etc.,
- it needs to receive from a human commands, by dialogue or possibly through a simple keyboard-like interface, about tasks that it should carry out
- needs to possess skills to carry out these tasks to the satisfaction of the 'master,' or the family, who might or might not be present while the robot executes these tasks.

Importantly, the robot must be aware of when and how to execute tasks, and to deal with conflicting. Thus, even the 'servant' paradigm for a robot in the house involves highly sophisticated problems of robot perception, action selection, and coordination of behaviour in the presence of humans.

A home robot could have other roles – as an assistant, a friend or a companion – that require more sophisticated means of perceiving and responding to the social environment. Thus, scenarios in which robots enter the private and social sphere of humans raise important issues of believability, acceptability, and 'naturalness' of robot-human interaction.

A socially, and human centred perspective in robot-human interaction fundamentally requires an interdisciplinary perspective with input from e.g. psychology, ethology, and social sciences.

CHALLENGES IN ROBOT-HUMAN INTERACTION RESEARCH

- Difficulties in a priori definition of design requirements Systematic user studies are necessary to explore the design space of social robots (see e.g. DiSalvo et al. 2002). People's attitudes towards robots will provide input for understanding the role of robots in human society (see e.g. Friedman et al., 2003).
- Robots that learn and adapt their social behaviour and communication and can deal with changing environments (e.g. Weng et al., 2001).

EMOTIONAL INTERFACES

Dylan Evans

Science fiction is full of machines that have feelings. These stories achieve their effect in part because the capacity for emotion is often considered to be one of the main differences between humans and machines. This is certainly true of the computers and robots we know today. Our interactions with these machines are rather dry affairs, mediated largely by keyboards and alphanumeric displays. Most robots today neither recognise human expressions of emotion, nor produce their own.

The new field of affective computing has already made progress in building primitive emotional machines. However, critics argue that a machine could never come to have real emotions like ours. At best, they claim, clever programming might allow it to simulate human emotions, but these would just be clever fakes (Evans, 2001).

Currently, advances in affective computing do not depend on resolving such deep philosophical disagreements. Researchers need not set themselves the far-fetched goal of endowing robots with real emotions. They can focus on more practical but still challenging objectives, such as producing robots that people perceive as being emotional. This is not so difficult to achieve, as people already ascribe emotions and other internal states to inanimate objects such as cars and electrodomestic appliances. As the Foresight Cognitive Systems Project report on 'Applications and Impact' states, people "are ready to take things at 'interface value'' (Sharpe, 2003).

Rosalind Picard, a computer scientist at the MIT Media Laboratory in Boston, has proposed dozens of possible applications for computational systems that can recognise human emotions and respond appropriately (Picard, 1997). These include:

- Artificial interviewers that train you how to do well in job interviews by giving feedback on your body language

- Affective voice synthesisers that allow people with speech problems not just to speak but to speak in genuinely emotional way
- Frustration monitors that allow manufacturers to evaluate how easy their products are to use
- Wearable computers (‘intelligent clothing’) that give feedback on your emotional state for therapeutic or other reasons (see the section in this report by the Professor David Murray’s on Wearable Computing)
- Robotic companions for elderly people
- Robotic pets and toys such as Sony’s AIBO robot dog, which has a limited range of ‘emotional expressions’.

In this context, the 2003 World Robotics survey, issued by the United Nations Economic Commission for Europe, revealed that entertainment robots now account for over 90 per cent of the number of robots installed worldwide. If this trend continues, it is likely that the leisure industry, rather than manufacturing or military requirements, will drive developments in emotional robotics, with some demand from the health sector. The analogy here is with home computer development : watch->game and watch, Sinclair Spectrum ->Atari -> PC etc..

At present, Sony’s AIBO robot dog is the only commercial off-the-shelf platform that is suited specifically to the demands of emotional robotics. A few labs are adopting this platform now that Sony has put the source code for AIBO’s operating system, OPEN-R, in the public domain.

It is more common, however, to find researchers constructing their own bespoke ‘robot heads’. The most famous of these is without doubt Kismet, the robot head and torso designed and built by Cynthia Breazeal and colleagues at the MIT. Kismet has moveable eyelids, eyes and lips. Kismet has a limited range of emotional expressions, but they are convincing enough to generate sympathy among humans who interact with him. Breazeal invites human parents to play with Kismet on a daily basis. When left alone, Kismet looks sad, but when it detects a human face it smiles, inviting attention. If the carer moves too fast, a look of fear warns that something is wrong. Parents who play with Kismet cannot help but respond sympathetically to these simple forms of emotional behaviour.

RESEARCH CHALLENGES

Among the various technological components that need further development are:

- Image-processing: including better face-recognition and face-tracking software, and more sophisticated algorithms for analysis of facial-expression in real time.
- Sound-processing: including algorithms for extracting affective information from intonation patterns and pitch of spoken language and from nonverbal utterances.
- Actuators: including novel actuators that might be used to deform plastic surfaces (‘robot skin’) with fast response times and low latency.

In addition to these technological developments, advances in emotional robotics will also require a deeper theoretical understanding of how emotions work in people. If robots are to discern a person’s emotional state, for example, they will require a valid taxonomy of human emotions. Likewise, robot designers will need to know what visual and auditory cues people most rely on when ascribing emotions to other people and to animals. Collaboration with psychologists, neuroscientists and other life-scientists will be essential to develop this understanding.

SPEECH INTERFACES

Guido Bugmann

Speech recognition and generation will constitute the main mode of communication between a service robot and its user. There will be no keyboard, screen or specialised robot language to learn. A spoken interface gives freedom of movement, enables collaborative work and is a natural communication method. Spoken interfaces, however, pose complex problems, not only of speech recognition, natural language processing (NLP) and dialogue design, but also of robot design.

In computer software, it is recognised practice to specify the user interface early in the design process and then to design the software around the interface. In robotics, this is a new concept. Spoken interfaces were very much seen as the last component to be added to a robot. This traditional approach then automatically requires the user to learn the specific language and keywords prepared by the robot’s designer. However, if one expects the robot to understand unconstrained spoken language, then the question of interface will need to be considered prior to robot design.

To illustrate this, let us assume that a user of a domestic robot cook needs to give an instruction involving the expression “a pinch of salt”. This will clearly exert constraints on how to design the robot’s manipulators. Similarly, if a mobile robot needs to understand the command “turn right at the blue sign”, it will need colour vision. Note that “turn right” is highly under-specified for a robot, with no details on what actuators must do. Hence service robots must gather missing information from the environment and make autonomous decisions, such as recognise the layout and plan a trajectory.

Thus, to understand natural language, a robot needs a high level of functionality. In fact, utterance like “clean this window” or “hang up the washing” make demands on robot design and control that are beyond current knowledge. There are also examples where particularities of human language – such as references to combinations of procedures – exert more subtle constraints on aspects of robot design, such as its computational architecture (Bugmann et al., 2004).

Verbal interaction with robots also sets constraints on the natural language processing components of the system. Robots are designed with a limited vocabulary corresponding to their action capabilities. In principle, this simplifies the design of NLP components and improves the performance of speech recognition. However, users do not know the limits of the domain of competence of the robot and often address the robot with utterances that it cannot understand.

We can reduce the frequency of such cases by tuning an NLP system on the basis of a representatively large corpus of utterances natural to users in that domain, but no matter how advanced a robot is, it will have limited linguistic and functional competence. When the user steps out of this domain, communication usually breaks down. The robot should be able to help the user to discover its domain of competence. An impractical alternative would be to ask the user to undergo long and detailed training sessions on the robot’s capabilities.

Verbal communication is a powerful tool for expression rules and sequences of operations. However, it is less expressive for shapes, locations and movements. Natural spoken communication is usually supported by gestures such as pointing to an object or a direction. Many tasks that cannot be explained and are best demonstrated.

This has long been recognised and research in speech interfaces is a part of the wider area of multi-modal communication.

Case Study

Instruction-Based Learning

In the Instruction-Based Learning (IBL) project, supported by EPSRC at Plymouth and Edinburgh universities, the aim is to teach robots new procedures using spoken natural language. For example, a user can teach a small mobile robot how to travel from A to B in a miniature town. This takes the form of a spoken dialogue between the robot and the user.

For instance:

User: Go to to the post-office.

Robot: How do I get there?.

User: Take the first right and take the second exit at the roundabout. Then it is on your right.

After this dialogue, the robot converts the instructions into an internal program in machine language that controls the displacements of the robot.

The next time the user asks for the same destination, the robot can retrieve the learnt procedure and execute the command immediately. So, the more the user instructs the robot, the more competent and personalised the robot becomes. The design of this system is based on the analysis of human-to-human route instructions and requires the implementation of high-level robot primitives.

SUMMARY OF CHALLENGES IN SPEECH INTERFACES

Dialogues are full of misunderstandings. The ability to overcome these makes human-human communication so effective. In this respect, human-robot communication is poor. A large number of problems remain, such as error detection, error repair, learning new words and actions, informative dialogues, etc. Such research is guided by findings and methods in psychology.

Overall, speech interfaces require a high level of functional competence from the robot, as humans refer to high-level functions in their everyday language. As robots are expected to mimic many characteristics of human listeners, several areas of life science can support the development of human-robot spoken interfaces, in particular psychological interaction studies and neuroscience of the auditory system.

Asoh H., Motmura Y., Asano F., Hara I., Hayamizu S., Itou K., Kurita T., Matsui T., Vlassis N., Bunschoten R., Krose B. (2001) Jijo-2: An office robot that communicates and learns. IEEE Intelligent Systems, 16:5, 46-55

Bugmann G., Klein E., Lauria S. and Kyriacou T. (2004) Corpus-Based Robotics: A Route Instruction Example. Proceedings of IAS-8, 10-13 March 2004, Amsterdam, pp. 96-103.

K. Dautenhahn, A. Bond, L. Canamero, B. Edmonds, Eds. (2002) *Creating Relationships with Computers and Robots*. Kluwer Academic Publishers.

K. Dautenhahn (2003): Roles and Functions of Robots in Human Society - Implications from Research in Autism Therapy, *Robotica* 21, pp. 443-452.

C. F. DiSalvo, F. Gemperle, J. Forlizzi, S. Kiesler (2002) All robots are not created equal: The Design and perception of humanoid robot heads. DIS2002, London, ACM Press, pp. 321-326.

Evans, D. 2001. *Emotion: The Science of Sentiment*. Oxford and New York: Oxford University Press.

T. Fong, I. Nourbakhsh and K. Dautenhahn (2003): A Survey of Socially Interactive Robots. *Robotics and Autonomous Systems* 42(3-4), 143-166.

Friedman, B., Kahn, P. H., Jr., and Hagman, J. (2003). Hardware companions?: What online AIBO discussion forums reveal about the human-robotic relationship. *Proceedings of CHI* 2003 (pp. 273 - 280). New York, NY: ACM Press

Garrod S, Pickering MJ (2004) Why is conversation so easy? Trends in Cognitive Sciences. 8 (1): 8-11.

Picard, R. 1997. Affective Computing. Cambridge, MA and London: MIT Press.

Sharpe, W. 2003. Cognitive Systems Project: Applications and Impact. London: DTI.

J. J. Steil, F. Rothling, R. Haschke, and H. Ritter (2003) Learning Issues in a Multi-Modal Robot-Instruction Scenario. Proc. IROS, 2003

Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., & Thelen, E. (2001). Autonomous mental development by robots and animals. *Science*, 291, 599-600

APPLICATIONS RESEARCH

WEARABLE COMPUTING

David Murray

The inexorable reductions in size and gains in power efficiency of electronic components, and progress in embedding and integration them within textiles, opens the way for clothing to acquire something of the nature of a second ‘perceptual’ skin. The idea of clustering sensing around the individual is complementary to the approach in ubiquitous computing, with sensing and computational power spread throughout the environment.

Wearable computing has three research goals. The first, and the most surely realisable, is the development of wearable computing substrates – the power supplies, processors, memory and interconnects required to support a range of input sensors, output devices and interfaces. The second goal is development of the sensory input and output interfaces themselves, devices that give a first-person perspective on the wearer's interaction with the environment, and to interact closely and in a socially acceptable manner with the wearer and others.

Although the success of the first and second will allow degrees of self monitoring by the wearer, or remote monitoring of the wearer, the third goal of research in wearable computing is one of closing a high-level cognitive loop around the wearer and wearables, to autonomously augment the wearer's capabilities, by offering contextually pertinent advice, warnings, and so on.

Realisation of the first goal is well advanced. In the US, for example, Jayaraman's group at Georgia Tech has developed ‘sensate liners’ [1,3] which provide a substrate for passive embedded biological and physical sensors, aimed at medical and military application. Carnegie Mellon University is undertaking device oriented work [2, 3]. In Europe, Infineon (Munich) and the Troestner's Wearable Computing Lab at ETH, Zurich, have used conductive textiles as a substrate for electronic components and as interconnection to the interface circuits.

The ETH group has developed its work in collaboration with MIT into commercial system called WearArm [4]. In the UK, the University of Bristol's Cyberjacket is a more conventional garment, loaded with interfaces, sensors and computers. Philips Laboratory, Redhill, takes a similar, but nearer-market, view of wearables. Its jacket connects a mobile phone and a MP3 player using textile cables.

Wearable computing will not develop if it remains at the level of sewing miniature versions of conventional input/output (I/O) interfaces into garments. The second goal is to develop person-oriented interfaces. Sensors and I/O devices are conveniently classified into those sensing the wearer and those sensing the wearer's surroundings.

SENSING PHYSIOLOGY

Physiological signals that have been proposed in the wearable context include; electromiography (EMG), electroencephalography (EEG) and electrocardiography (ECG) to determine muscle, brain and heart activity, respectively, galvanic skin conductance, and respiration and blood pressure sensing. Beyond straightforward monitoring, their use systematically to moderate human-computer interaction in so-called affective computing (for example, the work at MIT [5]) remains experimental, perhaps because the interpretation of the incoming signals is still a research issue.

Environmental sensors proposed for wearables include cameras, microphones, gyroscopes, GPS, compasses, ultrasound sensors to measure the temperature, barometer pressure sensors, humidity sensors and accelerometers. All these sensors continuously deliver differently structured data at rates from a few bytes/s for GPS, say, to several Mbytes/s for video. Handling, synchronising, segmenting and recovering salient signals from these data is difficult in itself, but the key challenge is in combining them to provide broad contextual information.

MOBILE VISION

A reflection of the substantial increase in mobile computing power is that 2-D and most recently 3-D vision has become realistic prospect for wearables. Vision fits well into the wearable domain. It offers many non-intrusive ways of recovering the cues that allow us to navigate, to follow action, to understand gesturing and signage, and to react to events.

A number of researchers, notably at MIT [6], have deployed static cameras on a wearer's clothing to recover ambient contextual information, such as lighting levels, colour histograms and so on. Microsoft Cambridge is working on similar ideas for memory augmentation. Others have linked cameras to the wearer's attention by attaching them to the head or hands, but often the views are too susceptible to the vagaries of the wearer's changing stance.

A different approach, taken at the University of Oxford, uses a miniature directable camera to allow exploration of the environment independent of the wearer's motion. The Oxford group has also demonstrated single-camera simultaneous localisation and mapping on its wearable, demonstrating a transfer of leading-edge robotic techniques to the wearable domain [7].

AUTONOMY IN WEARABLES

These advances begin to show a way towards establishing self-contained reliance or autonomy in wearables, the third core research goal. However, the potential of wearable systems for truly cognitive, multi-modal processing has barely been addressed, with only global contextual aspects extracted using simple sensors. Often reasoning has been based directly on raw sensor data, without intermediate processing. Mono-modal data processing has dominated so far.

Examples of application scenarios for wearable computing lie in:

- *Memory augmentation*, where textual, audio and visual input is compared with a database to retrieve related information. There are applications not only in technical fields, but also as aids for the elderly.
- *Surroundings awareness,* which seek fast ways of retrieving broad indications of location, such as outside/inside, quiet/noisy and so on. This provides information that should modify the behaviour of other wearable equipment. This category might also include the recognition of useful objects and people in the surroundings.
- *Enhanced self-awareness,* where rather than just monitoring bio-signals, the systems attempts to match effect and cause, and to proffer remedial advice.
- *Navigational aids,* applied to, say, the emergency services, tourism and disability care.

- *Remote collaboration*, where wearable computers gives the opportunity to follow people to where they move, and to share first-person views with remote collaborators or experts.
- *Augmented reality,* where the aim is to annotate the environment around the wearer.

REFERENCES

[1] C. Gopalsamy, S. Park, R. Rajamanickam and S. Jayaraman, "The Wearable Motherboard: the first generation of adaptive and responsive textile structures (ARTS) for medical applications", Journal of Virtual Reality, 4:152-168, 1999.

[2] P. Stanley-Marbell, D. Marculescu, R. Marculescu and P.K. Khosla. Modeling, Analysis, and Self-Management of Electronic Textiles. IEEE Trans. on Computers, 52(8):996-1010, 2003.

[3] D. Marculescu, R. Marculescu, S. Park and S. Jayaraman, Ready to Ware. IEEE Spectrum 40(10): (2003) Online at www.spectrum.ieee.org/WEBONLY/publicfeature/oct03/ware.html

[4] P. Lukowicz, U. Anliker, G. Troester, S. J. Schwartz, R. W. DeVaul, The WearARM Modular, Low-Power Computing Core. IEEE Micro 21:16-28 (2001)

[5] R.W. Picard, E. Vyzas, and J. Healey, Toward Machine Emotional Intelligence: Analysis of Affective Physiological State IEEE Trans on Pattern Analysis and Machine Intelligence, 23(10): (2001).

[6] B. Clarkson, K. Mase and A. Pentland, Recognizing user context via wearable sensors Proc 4th Int Symp on Wearable Computing, Atlanta GA Oct 2000 69-75.

[7] W.W. Mayol, A.J. Davison, B.J. Tordoff and D.W. Murray, Applying Active Vision and SLAM to Wearables. Robotics Research: Proc 11th Int Symp on Robotics Research, eds R Chatilla and P Dario. Springer Verlag 2004.

SPACE AND PLANETARY ROBOTICS

Alex Ellery, Dave Barnes, Colin McInnes, Alan Winfield, Ulrich Nehmzow

Introduction

The application of robotics in space are unique in that it forces the robot to survive and function without human assistance. In acting as a platform for the projection of human capabilities to remote and hostile environments, the robot provides an existence proof for understanding our own cognition. Robotics is also the key to one of the most exciting scientific endeavours of the 21st century – the search for life elsewhere.

Rationale and Relation to Life Sciences

Robotics has much to contribute to space exploration. Generally, space robotics follows the division of robotics into manipulators and mobile robots. It has applications in the following areas:

- planetary exploration by surface rovers – be they wheeled, tracked or legged, ground-penetrating moles, aerobots, hydrobots, cryobots, snakebots etc.

- space-mounted manipulators for planetary deployment or robotic on-orbit servicing tasks on ailing satellites

Of particular interest here is the role of robotics in exobiology/astrobiology. A primary role for robotics in planetary exploration is in exobiological prospecting. Until humans arrive on planets such as Mars, Europa and Titan, robotics is the key to finding evidence of former or extant prebiotic or biotic species.

Robotics determines the selection and sampling capability of scientific targets during a space mission. Robotic manipulators, drills and rovers are the only tools we have for performing in-situ exobiology.

Robotics is a key element in the deployment of scientific instruments and is intimately part of scientific instrument design. The deployment of such instrument to otherwise inaccessible targets is the key to maximising the probability of success. This is where robotics comes in. Astrobiology as a science needs an experimental component focussed on extraterrestrial environments. In a very real sense, planetary robotics is the experimental component of astrobiology.

Space Manipulators

The University of Wales Aberystwyth developed control software for the Beagle 2 manipulator arm which was to import stereo camera system images of the Martian terrain and convert these images into 3-D terrain models

Space-based robotic manipulators also provide the basis for on-orbit servicing of satellites, through the replacement of equipment modules. In the past five years or so, the incidence of on-orbit failures has reached epidemic proportions. In addition, astronauts are now forbidden to perform certain servicing missions that have been the key to the success of space telescopes such as the Hubble Space Telescope. Indeed, the Hubble telescope faces decommissioning unless robotic servicing spacecraft can meet its servicing requirements.

Path-planning to generate safe collision-free trajectories for free-flying robot manoeuvring in proximity to other satellites is an important part of robotic servicing missions (Figure 2b, University of Glasgow). Such techniques may allow co-operative control of multiple free-flyers for on-orbit assembly problems. The dynamics and control of such manipulators, particularly with regard to the control of actuation forces between the servicer and the target spacecraft whilst the servicer grapples the target, is a particularly thorny issue (University of Surrey).

The German Aerospace Centre, Deutschen Zentrum für Luft- und Raumfahrt (DLR), is working with the University of Surrey to introduce robotic servicer spacecraft within the next five years. Such robotic servicers represent an order of magnitude increase in complexity for spacecraft as most spacecraft are designed as sensor platforms rather than as actuation systems. The on-orbit servicer must be designed to collide with its target in a controlled manner and perform useful and complex tasks.

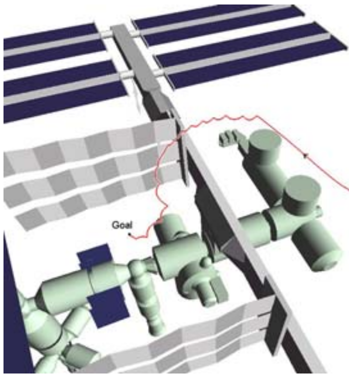

Figure 2. a) (left): Advanced TeLerobotic Actuating Servicer (ATLAS) concept for on-orbit servicing of satellites – robotic on-orbit servicing is the key to further infrastructure development in space, offering significantly reduced costs to the construction of spacecraft and enhanced performance (University of Surrey).

b) (right): Potential field path-planning of collision-free trajectories for robotic freeflyer satellites around the International Space Station (ISS), a cluttered environment (University of Glasgow)

PLANETARY ENTRY DESCENT & LANDING

Both the University of Surrey and the University of the West of England have been concerned with entry descent and landing of robotic devices onto the surfaces of other planets. Autonomous planetary landing systems for small solar system bodies (asteroids) require autonomous mapping of the asteroid on arrival, selection of landing sites, planning and execution of the descent and landing phases (University of the West of England). Researchers at the University of Surrey have focussed on planets with atmospheres and the control of descent available using different techniques such as inflatable structures.

PLANETARY ROVER ‘TERRAINABILITY’

The environment in which the autonomous planetary rover operates is largely unknown. It is hard to install adequate control mechanisms *a priori*. The group at the University of Surrey is concerned with the need to understand such environments in which rovers will operate and to provide sufficient robustness to survive. Planetary surfaces must be understood and negotiated effectively in order to traverse such hostile terrains. The University of Surrey team is interested in the interaction forces between the robot and its hostile terrain environment. This is a problem of chassis design and we have been championing a new chassis concept called the elastic loop mobility system which is a variant on the track to maximise terrain traversability with minimum chassis complexity (Figure 3).

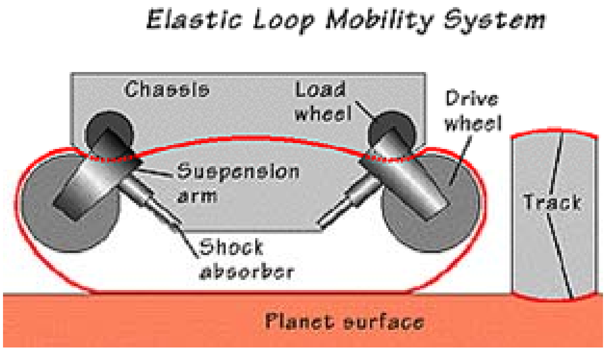

Figure 3. Elastic loop mobility system (ELMS) – Mars offers a hostile rocky terrain but which places a premium of rover chassis design which must be able to traverse this terrain autonomously and reliably

The Surrey group has developed in-house software tools to predict the traction performance for any robotic vehicle on planetary terrain. We have used these tools in the Phase A study of ESA‘s ExoMars rover, with EADS Astrium UK as the prime contractor.

We are further developing these tools for ESA for general-purpose use to assess the performance of any rover design. Based on our analyses in understanding the Martian terrain, we are utilising this experience to develop more robust autonomous navigation algorithms that take into account the terrain of the planetary surface to allow automated throttling of the robotic rover in response to obstacles and changing soil conditions. To that end, we are developing the behaviour-based potential field technique to provide autonomous navigation with robust obstacle negotiation capability.

PLANETARY AEROBOT NAVIGATION

A group at the University of Wales, Aberystwyth, has focused on the use of balloon-type planetary aerobots to acquire images of the planetary terrain to reconstruct accurate models of the surface of the explored planet, and accurately locate of the balloon with respect to the Martian surface. They include extracting naturally occurring terrain features from topographical maps and updating terrain data for the purpose of navigation. Furthermore, 3-D Navier-Stokes simulations in conjunction with terrain data, provides predictions of realistic Martian wind conditions.

PLANETARY ROVER INTEGRATION

The University of Surrey has amassed much experience in the interface between astrobiology and the robotics component of planetary exploration. The group’s particular interest is in infrared laser Raman spectroscopy to detect biomolecules. It is also interested in the robotic deployment of rover-mounted ground-penetrating radar to detect sub-surface water on Mars and select drilling sites for sample extraction. It has been developing a small micro-rover concept called Vanguard (Figure 4).

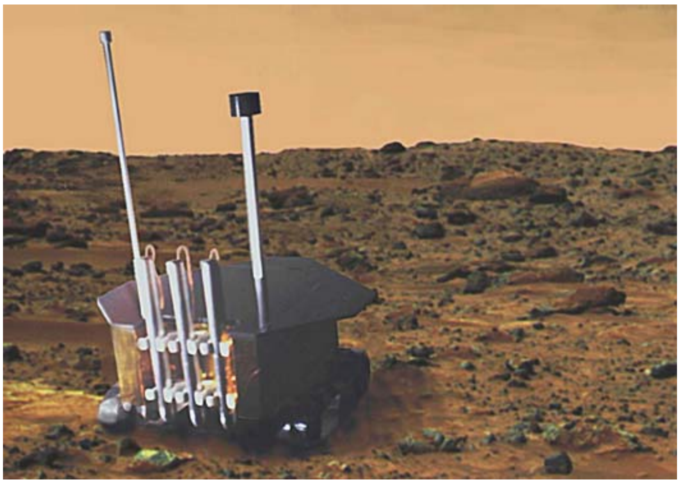

Fig 4. Vanguard micro-rover proposal with three vertically-mounted moles to search for evidence of former life on Mars to depths of up to 5m below the surface.

Integration with a focus on the scientific objectives and requirements of a space mission are the key to the success of space robotic systems. A group at the University of Glasgow has analysed the integration of robotic payloads on large science missions with the combined use of manipulators and small rovers to maximise the scientific return for a mission to Mercury to bring back samples.

SPACE APPLICATIONS OF BIOMIMETICS

Groups at the University of Surrey, University of Wales Aberystwyth and the University of the West of England are looking to the next generation of robotic rovers. The University of Surrey's group is interested in exploiting biomimetic techniques where appropriate to future rover system to enhance their capabilities. It is currently studying this issue for ESA with a view to developing a prototypical legged rover. In particular, we have established an interest (with Surrey Satellite Technology Ltd) in the use of electro-active polymers as miniaturised artificial muscle actuators. The group's specific interests are in compliant control of actuators, electro-active polymers, and genetically-evolved neural networks, with the University of Sussex, and their *integration* into rover platforms to provide greater autonomy and robustness to planetary explorers. Integration is the key to the biomimetics approach to planetary robotics.

Conclusion

The broad themes of space robotics research in the UK are clear. The single most important theme that runs throughout all the activities is *autonomy*. The needs of space science dictate that robustness is the key to this autonomy. It is specific in that autonomy refers to the need for survival and scientific function of planetary robots in real world, hostile environments – namely, planetary environments – for which there is limited scope for human intervention. This type of autonomy is *robotic* autonomy, dealing with the real and uncompromising world. This type of autonomy yields tangible problems with tangible metrics for performance – successful rover missions.

The scientific community, astrobiology in particular, requires robotic platforms to meet the challenge of searching for evidence of extraterrestrial life, one of the most important scientific questions of this century.

Faced with the same stark challenges as the biological world, we must overcome similar problems and find similar and/or appropriate solutions. It is no wonder that biomimetics is of interest to many of us. The transfer of such technology to other terrestrial applications will be of immeasurable benefit to these applications, in oceanography for example. But the technology transfer will be much wider. As our capabilities for achieving autonomy advance, so will the diversity of terrestrial applications.

ROBOTICS IN MEDICINE

Brian Davies

The UK is relatively strong in medical robotics. Two companies focus on the supply robots for surgery; Armstrong Healthcare Ltd, and The Acrobot Co. Ltd.

Armstrong has an endoscopic mover Endo-assist and is researching systems for neurosurgery and orthopaedic surgery. Acrobot is nearing the conclusion of randomised trial by the Medicines and Healthcare products Regulatory Agency of 15 robotic and 15 conventional operations. This activity focussed on minimally invasive uni-condylar knee replacement surgery.

In addition to these two companies, orthopaedic research is under way at Imperial College, on knee, hip and spine. Loughborough University has also developed a special purpose robot for leg trauma, particularly femoral nail fixation. Hull University has primarily focussed on computer-aided surgery techniques for interventions, but has also used a special purpose haptic system for knee surgery training. More recent haptic research at Imperial College is on an arthroscopic knee surgery simulator and trainer, in conjunction with Sheffield University which supplies the knee imaging and simulation aspects.

Soft-tissue robotic systems are also researched at St Mary’s Hospital, Paddington, part of Imperial College. The group there has a Da Vinci robot for closed heart surgery. This has also been used clinically for a range of other procedures such as gall bladder removal and radical prostatectomy. The group is are also researching changes to the user interface and also general haptics investigations.

The Mechatronics in Medicine group at Imperial College also has a Da Vinci slave robot which it is using to investigate the implementation of haptics onto the Da Vinci robotic system. This group is probably the largest researching medical robotics in the UK. It has previous activity in: a prostate robot for transurethral resection of the prostate, the first active robot in the world to be used clinically for tissue removal, in 1991; a special purpose robot for neurosurgery, a robot for spine surgery; a robot for brachytherapy of the prostate; and bloodbot for finding veins and taking blood.

A further activity has been the use of high intensity focus ultrasound (HIFU) in which a robot moves the focus of a multiple probe system in an attempt to treat liver tumours on an outpatient basis. There have also been a number of computer-aided surgery activities –including a training system for prostate surgery and a system for hip resurfacing surgery.

A recent project supported under the New and Emerging Applications of Technology Programme (NEAT) of the Department of Health is to devise an MRI compatible robot, initially intended for prostate biopsy. This group also gave rise to the spin-off business, the Acrobot Company Ltd, in 1999. The Guy's Hospital urology unit recently purchased a further Da Vinci robot has for radical prostatectomy.

Another group concerned with haptic robots for medicine is at Salford University as part of an EC Framework 5 project Multisense. This uses a haptic system to demonstrate hip replacement surgery. A further medical robot is about to undergo preliminary clinical trials in robotic stapedotomy at Aston University

Open questions

Three main aspects are important drivers for surgical robots; accuracy, minimal invasiveness, and application in areas of restricted access, such as within the closed bore of CT or MRI imaging devices

The attempt to provide accuracy has primarily been demonstrated in orthopaedic surgery, in which bones can be clamped and treated as fixed objects in a quasi Computer Numerical Control (CNC) milling procedure. Soft tissue, in which the target moves, is more difficult. Open loop telemanipulator systems are used so that the surgeon at the master can track online soft tissue as it deforms and distorts.

Most soft-tissue activities attempt to undertake minimally invasive surgery. Rotation centres are generally provided at the skin or the skull surfaces. This constrains the robot mechanism and provides difficulties for haptics. Most robotic systems employ excellent imaging, but have virtually no haptic sense. Recent attempts to image a patient and to use a robot to take biopsy samples try to use online imaging to compensate for soft-tissue motion. Quality 3-D real-time ultrasound probes would make this online imaging process much cheaper and simpler.

Most recently, small and dexterous stents and biopsy probes have allowed keyhole surgery behind objects such as the heart. The tracking of the beating heart and breathing, while not compromising safety, is another issue in medical robotics

Robots in medicine are expensive to develop and to produce. The cost/complexity versus clinical benefits need to be carefully balanced. Probably the biggest difficulty in implementing medical robotics is the gulf between projects in the lab and systems that are safe to implement in the operating room. Further problems are concerned with the unclear regulatory requirements in this rapidly developing area, which is also prone to litigation.

Collaboration with medics

Clearly there are considerable interactions with a wide range of life sciences. The surgeon and radiology community are heavily involved. More recently, those concerned with precision placement of material within the body are also becoming interested. This is centred around the desire for precise placement of genetic material and tissue scaffolds at targeted sites, for example. The area of rehabilitation robots has long been a concern of physiotherapists and occupational therapists, but has more recently become the domain of stroke therapists in treatment of the aged and other stroke patients.

DOMESTIC ROBOTS

David Bisset

Autonomous domestic vacuum cleaners exist. Every major manufacturer of domestic appliances has announced that they have a prototype, but real products have been far slower to materialise. One reason for this is simple, building a reliable robot, even with the limited autonomy needed for domestic vacuum cleaning, requires a high level of technology and new skills not normally employed by domestic appliance manufacturers.

The first generation of domestic vacuum cleaners are little more than ‘helpers’ requiring significant human intervention to obtain an acceptable outcome. Able to remove some of the surface dirt, their random bouncing around a room is likely to leave significant areas uncleaned or worse leave the robot stuck behind sofas or under tables. The limited ability of robots to map an unknown space and work out where they are, leaves them unable to detect accurately how much of a room remains to be cleaned or even to work out how to get back to where they started.

At $200, the Roomba from iRobot Inc has been largely successful but it is limited, and requires user intervention and careful use. Other machines on the market – from Electrolux and Karcher for example – cost five times as much but are not five times more effective.

First generation autonomous vacuum cleaners bypassed the issue of user-to-machine communication because they are single-function machines that operate in a specific way when turned on. They need little more than an on/off button. The lack of user interaction is an inevitable consequence of their single-function nature and their limited knowledge of their own location. Such autonomy can be characterised as ‘constrained autonomy’. The devices are autonomous only within a limited space – often defined by physical barriers or light beams – and do not make significant decisions about the nature of the task they execute.

Future generations of vacuum cleaner will need to be more autonomous. They will need to know *where* they are and will need to make significant task decisions based on *what* they sense in the environment. These improvements represent significant technical steps.

For ‘place autonomy,’ an appliance needs to know its own location, possibly only relative to a start point rather than in absolute terms, and be able to execute actions on identifiable areas, often without physical boundaries: “Clean up the mess in front of the sofa”.

The autonomous appliance will need, or appear to have, a good enough understanding of the cognitive context of the items it encounters. In essence, there must be a cognitive dialogue between the appliance and the user. There must be both fulfilled expectations and comprehensible communication. The comprehensibility of the communication must be such that the user can communicate commands in an easily understood and unambiguous way without resorting to a complex user interface: “Clean the lounge and the hallway, but not the kitchen”.

In a conventional appliance, the buttons and controls have a one-to-one relationship with the functions that the appliance can achieve: “Press the record button to record”. On an autonomous appliance the actions required to carry out the command “clean the lounge" will be different in each house, not only in terms of the shape of the room but in terms of where the user perceives the lounge to be. The accurate communication of this perception to the machine is an essential part of its operation or installation.

The need for a common cognitive dialogue between user and appliance is different from the problems that are encountered by fully autonomous machines operating without human interaction. Fully autonomous robots would be free to interpret sensory signals in their own way. By definition they are not required to hold or create a cognitive dialogue with a human. As a consequence, such machines may have limited utility.

For machines that must interact, this essential requirement for a common cognitive dialogue with humans presents a significant technical challenge. Interestingly, humans often find communication between themselves difficult and error prone.

Domestic vacuum cleaning represents a micro world in which autonomous systems must operate. And yet that micro world reflects many of the major issues that face the designers of autonomous machines. The issues of a common cognitive representation and its communication are core to the study of cognitive science, neurophysiology and philosophy.

ROBOTICS AND AUTOMATION IN FOOD MANUFACTURING

John Gray

The food industry is the largest sector in the UK economy and is valued at approximately £100 billion per year. In regions such as the Northwest of England, the food sector accounts for about 45% of the manufacturing base and about 14% of the GNP. Traditionally, the industry has been a labour intensive sector, both at the raw material end of the supply chain and in the assembly of food products. In 1992, a DTI report highlighted the sector as a possible growth area for robotics and automation with the UK.

In the intervening years, there has been little progress in bringing robotics to the food sector. However, recently there have been significant drivers for change. These include financial pressures and demands on hygiene, quality and consistency as well as regulation on health and environmental factors, not to mention traceability requirements. The technical challenges are also recognised. Food automation requires grippers for soft, fragile sticky/slippery, irregularly shaped objects, instrumentation and inspection, fast and reconfigurable automated lines to supply rapidly changing markets and a philosophy of hygienic design.

Many of the standard characteristics of current industrial robots and their components are unsuitable for deployment in the food sector. There is a need for a new paradigm to define the requirements of a general purpose robotic device for the sector. This will involve specified performance envelopes, new construction materials and innovative mechanical design.

An ideal is emerging which suggests a 'dark factory' concept, with production facilities compressed into compact, sealed, modular units with savings on labour, wastage, factory space, cleaning costs, energy consumption and the vast infrastructure support for the current labour intensive manufacturing procedures. Clearly this approach poses major issues of reliability, manufacturing flexibility and rapid reuse and on-product inspection. There are also research issues on overall process modelling, optimisation and control.

Clearly not all food products will be suitable for such an approach. However, the major supply chains will demand the sort of consistency, quality and hygiene (aseptic) of outputs that this type of system can deliver. Given present market demands, and recent changes within the European Union, the maintenance of a labour intensive food industry within the UK may not be a viable option beyond a timescale of five to 10 years from now. Robotics and associated automation technology has thus a vital role to play in the maintaining this industry with the UK, with opportunities for robotics research.

ROBOTICS IN AGRICULTURE

Tony Hague

Applications of robotic technology in agriculture have, on the whole, focused on two main areas, automated harvesting of produce and automatic guidance of agricultural machines. The aim of the ‘driverless tractor’ has been around for some time. Mechanically guided, and later leader cable guided, prototypes appeared as early as the 1950s. It is only recently, however, that technology has become sufficiently mature for automatic guidance to be a realistic possibility, with the existence of GPS and affordable hardware for machine vision.

The best known research project is perhaps the Carnegie Mellon Demeter automated harvester. This achieves automatic guidance using a fusion of GPS derived position with machine vision to detect the cut edge of the previous pass. The initial aim is commercialisation as a driver assistance system, with autonomous operation as a long-term goal.

Vision guided weed control

VISION GUIDED WEED CONTROL

In the UK, the Silsoe Research Institute (SRI) of the Biotechnology and Biological Sciences Research Council demonstrated an autonomous crop treatment vehicle. This achieved navigation by fusion of machine vision to detect crop row with inertial and odometric sensing. Crop plants were detected using the planting pattern, allowing the crop to receive targeted treatment, and the vehicle to navigate the crop rows. Fully automatic operation, including turning at the end of rows, could be achieved.

Current commercial products use some of the key technologies developed for autonomous navigation, but as a driver aid. Automatic steering control for tractors and harvesters, based upon GPS, is available from manufacturers such as John Deere and as retro-fit kits. These systems take over the task of straight-line steering. SRI's autonomous vehicle research has produced a commercial system for precision guidance of weed control machinery mounted on a manually driven tractor (see photos). Here machine vision locates the crop rows and a hydraulic control system positions cultivator blades between the rows to an accuracy of 15 millimetres.

As yet, there is no commercially available fully autonomous agricultural machine. More development is needed in a number of areas, starting with robust solutions for navigation. Current high-precision, real-time kinematic GPS systems perform poorly under, or near to, trees, presenting problems around field margins, where visual navigation is also less reliable.

Secondly, although navigation has been widely studied, there has been less attention to the automation of other functions of the driver, as a monitor of operations, and the function of the machinery itself (e.g., detection of blockages, etc). Indeed, the success of commercial guidance systems as a driver aid has been in part due to the benefit of freeing the operator to deal with other tasks.

Finally, improved visual perception might permit an autonomous platform to carry out tasks that are not easily performed manually. The automated detection and mapping of weed species has been approached using a variety of techniques, including spectral and shape analysis. To date no method delivers robust operation under field conditions.