STEFAN SCHAAL
MAX PLANCK INSTITUTE FOR INTELLIGENT SYSTEMS, TÜBINGEN

Robots are becoming autonomous

Towards the end of the 1990s Japan launched a globally unique wave of funding of humanoid robots. The motivation was clearly formulated: demographic trends with rising life expectancy and stagnating or declining population sizes in highly developed countries will create problems in the foreseeable future, in particular with regard to the care of the elderly due to a relatively sharp fall in the number of younger people able to help the elderly cope with everyday activities. Autonomous robots are a possible solution, and automotive companies like Honda and Toyota have invested billions in a potential new industrial sector.

Today, some 15 years down the line, the challenge of "ageing societies" has not changed, and Europe and the US have also identified and singled out this very problem. But new challenges have also emerged. Earthquakes and floods cause unimaginable damage in densely populated areas. And ongoing global climate change could increasingly trigger further natural disasters. Another constant threat, that of rapidly spreading epidemics, has recently been highlighted by the Ebola crisis.

One hope for the future is that someday we will have autonomous technological helpers, i.e. autonomous robots that can provide assistance in all these areas. Of course, visions of such autonomous systems go well beyond the relatively obvious concepts of humanoid robots. They cover a broad range of scales – from nanorobots to large industrial robots. Toy robots, robots that can provide emotional support and miniature robots able to carry out clinical tasks directly inside the body are just some conceivable examples.

However, the current state of science and technology is far behind such societal aspirations. In the wake of the Fukushima reactor catastrophe, no robot was deployed that could provide useful help – despite the rapid development of mechatronics, which has spawned thousands of robot prototypes, especially in Japan. During the ongoing Ebola epidemic, people continue to be exposed to the contagious viruses during cleaning and clearance work. These are tasks that really could be performed autonomously without human intervention. And nanorobotics for clinical purposes is still firmly in the realm of basic research.

The first simple autonomous robots are being used today in private households (vacuum cleaners, lawnmowers), for military reconnaissance in rugged terrain, in driverless cars and in drones. These are essentially mobile wheeled or airborne robots without arms or grippers to perform manipulations. Robots that are able to run on legs across uneven ground or carry out complex manipulations are still in the research stage. Research into humanoid robots and assistive robots is being pursued around the world. Problems such as the complexity of perception, effective control without endangering the environment and a lack of learning aptitude and adaptability continue to confront researchers with daunting challenges for the future. Thus, an understanding of autonomous systems remains essentially a topic of basic research.

AUTONOMOUS ROBOTICS: PERCEPTION–ACTION–LEARNING

Autonomous systems can be generally characterised as perception-action-learning systems. Such systems should be able to perform useful tasks for extended periods autonomously, meaning without external assistance. In robotics, systems have to perform physical tasks and thus have to be realized physically. Sensors allow the system to perceive the environment and its own body. Learning and adaptation mechanisms enable it to adapt to changing environments or learn entirely new behaviours. The system must be able to respond to changing situations and disturbances robustly and without accidents.

It is difficult to assign autonomous robots to a specific discipline, such as artificial intelligence (AI), mechatronics research or machine learning (ML). Autonomous robotics requires a large number of components that must all be coordinated to collectively produce a robust behaviour, e.g. perception, control, planning and learning processes. This also raises a challenge regarding the reliability of such systems: for example, if ten components are concatenated, each of which has 99% robustness, the probability that the overall behaviour of the system will be reliable is only 90%.

Such a system would not be viable in daily use. However, given the current state of research, 100% reliability of the individual components is essentially only possible with analytical methods, i.e. methods that can be derived from accurate mathematical modelling. Empirical methods, namely data-driven methods that are derived from machine learning, barely achieve this level of reliability. This presents a daunting problem: complicated robotics systems have no reliable modelling, and empirical methods are not yet sufficiently accurate to bypass analytical models. A current focus of research is therefore to determine which system components can be realised with machine learning, as well as to find new analytical methods that are more robust than inaccurate models. As it is often very time-consuming for robots to learn new behaviours, discussions are also focussing on how a global database and central computing centres for robotics can be established by means of cloud computing. This appears particularly promising for perception tasks, but not so much for direct robotic control signals, which are very system-specific and require real-time processing.

In conclusion it can be said that autonomous robotics require a very broad knowledge base beyond the usual and narrowly defined standards of the contributing core disciplines. What is required is not so much specialists from one field but individuals who are multidisciplinary in their approach and are able to build bridges. In other words, scientists are called for who are simultaneously experts in mechanics, physics, electrical engineering, computer vision, control technology, ML, AI and software engineering. An interest in cognitive sciences would also be helpful when the aim is to mimic biological systems. To date, there are few researchers in the world who fit this profile and few educational options that convey such a broad knowledge.

AN OUTLINE OF THE CURRENT STATE OF RESEARCH

Mobile systems: When it comes to achieving high levels of autonomy, mobile robotic applications are certainly the most advanced. Mobile robots – with few exceptions all of them mounted on wheels, not on legs – are standard research topics in many computer science labs. This topic was particularly popular in research in the 1990s and experienced a breakthrough around 2005 thanks to the DARPA Grand Challenge, a competition in autonomous vehicle navigation. The fact that six teams reached the finish line in the Grand Challenge with rather minor time differences shows that research and technology in this field is quite advanced. The focus was more on good system integration, software engineering and quality control rather than on algorithms; most of the components for autonomous navigation were already well known at the time. It is therefore surmised that autonomous navigation on wheels has gradually left the research labs and has now been taken up by industry through technology transfer.

Assistive robots and humanoid robots: The situation is quite different with regard to autonomous manipulation robots, i.e. robots with arms, hands and legs, which include humanoid robots. Navigation in two-dimensional space with vehicles is a low-dimensional problem and is highly structured by the environment (roads, corridors, traffic signs, walls, etc.). The aim of navigation is therefore defined in rather simple specifications, namely to reach a destination while avoiding accidents. Such scenarios can be effectively handled today by machine learning and artificial intelligence methods. By contrast, manipulation robots are high-dimensional systems, often with more than 50 control dimensions, and non-intuitive geometric spaces for planning and control. For example, a humanoid robot often has seven degrees of freedom per arm and leg plus 10–20 additional degrees of freedom for the body, fingers and head. Many machine learning and artificial intelligence methods were developed for low-dimensional spaces and do not scale up well to high-dimensional spaces, i.e. the calculations required to plan and optimise in these spaces grows exponentially with the number of spatial dimensions. In many cases it is not even possible to collect enough data in a sufficiently short time to be able to process them empirically in these spaces: the process of data collection in real-time is simply too slow to even approximate high-dimensional spaces, even after 100 years of uninterrupted operation of a robot. Hence, a mixture of analytical and empirical methods is currently the most efficient way to achieve competent behaviour. If good mathematical models exist, it is often possible to work more efficiently with those models than with models that first have to be learned empirically through data acquisition.

FREQUENTLY, IT IS NOT CLEAR HOW THE GOAL OF A BEHAVIOUR CAN BE MATHEMATICALLY FORMULATED. FOR EXAMPLE, WHAT DOES "STIRRING SOUP WITH A SPOON" MEAN IN MATHEMATICAL TERMS?

The current state of research is that complex behaviour such as manipulation tasks, locomotion on legs and whole-body manipulations still requires a great deal of manual programming. For a robot to fetch a cup from a kitchen cabinet, for instance, scientists divide the overall task into individual subtasks (e.g. "go into the kitchen", "find the cabinet", "open the cabinet door", "grasp the cup", etc.) and define the sequence in which those actions must be performed. The robot's behaviour in the individual subtasks is then achieved by means of programming and/or learning methods. In many cases, robustness problems occur in perception, in the control of the robot and in the algorithms used, and additional behaviours are required to correct for errors. Frequently, it is not clear how the goal of a behaviour can be mathematically formulated. For example, what does "stirring soup with a spoon" mean in mathematical terms? Quickly, overly specific solutions for a behavioural task arise that are neither robust nor transferrable to other tasks or other robots. Small changes in the environment or calibration of the robot can quickly lead to completely incorrect behaviour.

Truly autonomous behaviour of manipulation robots is therefore still very far from reality. Certain autonomous subsets of behaviours can be achieved, for example balancing and walking on two legs on level ground or movements for grasping relatively arbitrary objects. The term "supervised autonomy" is often used today. This means that a short movement is performed autonomously, after which a human determines the next goal and the appropriate behaviour. The hope of constructing autonomous robotic assistants in society, a vision described earlier, is therefore still a long way from realisation.

THANKS TO THE ALMOST COUNTLESS IMAGES ON THE INTERNET, AN INEXHAUSTIBLE RESERVOIR OF LEARNING DATA IS NOW AVAILABLE.

Micro- and nano robotics: Reducing the linear scale of robots to centimetres, millimetres, micrometres and even nanometres produces a whole new set of challenges. Different physical laws apply at these scales than at the scale of mammals. Whereas the mechanics of a human-like system is largely characterised by inertial forces, at the centimetre and millimetre scales frictional forces and surface tension play a much greater role. That is why, for example, some insects, unlike humans, are able to walk on water. At the micrometre scale, this means that an intelligent swimming technique in liquids with a corkscrew-like motion is much more efficient than, say, the fin-driven motion of fish. Moving down even further on the length scale, stochastic influences dominate, such that behavioural goals can only be achieved in swarms of many systems – and then only to a certain degree of probability.

Sensors and motors are also different at small scales. Sensors are often realised by way of chemical processes. Motor function and energy supply become complex problems, so that an external source is often needed in research, for example magnetic fields or an atomic force microscope. Batteries are utterly inefficient at the micro and nano scales. And it remains unclear what data-processing methods are even possible at such scales. Autonomous robotic systems at the micro and nano range are still firmly in the domain of basic research.
Autonomous perception

Recognising and understanding the world and deriving behavioural possibilities are important aspects of perception. Computer vision (the simulation of vision by computers) has made great strides in recent years – on the one hand, as low-cost camera systems are now available, so that anyone who wants to work with cameras can easily do so; and on the other, because computers have become so fast that even complex image processing can be performed on a laptop.

Scene categorisation and object and face recognition are very far advanced. Thanks to the almost countless images on the Internet, an inexhaustible reservoir of learning data is now available. If this can be combined with large computer networks, it would be possible to train learning systems to a high level of quality. This is currently being intensively studied with deep learning, a new approach to machine learning that uses multi-layered neuronal networks and improved learning methods which profit from large computing clusters. "Action recognition" is another topic that is also attracting much attention. Specifically, it consists of recognising human motions and then understanding the intention behind those motions. It is relatively easy to recognise locomotion on two legs; however, recognising whether a person is playing soccer or just jogging is a much more complicated task. What researchers hope to achieve is an understanding of every kind of action – for example, grasping a cup, eating with a fork or playing with building blocks. This skill is essential if autonomous robots are to interact with humans in a human environment.

Of course there are also other perception sensors: tactile sensors, force sensors, microphones, etc. can provide valuable data to complement other sensory modes. Human perception is probably so successful because it is multi-modal, such that we rarely have to rely on just one modality. This fusion of sensors is an area of research, but is unfortunately not very pronounced in current projects.

In any case, a major problem remains: making the information processing robust enough so that variables – a change in lighting, the noise of a construction site, the tactile difference between a plastic cup and a porcelain one, or the like – do not overwhelm the perception system. Robust autonomous perception is therefore certainly one of the most daunting problems facing autonomous robotics.

Learning systems, planning systems and artificial intelligence: There are now a large number of algorithms to facilitate planning and learning tasks. Only experienced experts are able to judge which algorithm is applicable, what its weaknesses and strengths are and how efficiently it can be used. Of course, the hope of science is to produce generic black-box systems, i.e. systems that function with 100% reliability without the normal user having to understand what takes place inside the system. That has not yet been realised and perhaps never will be, as universal learning systems (systems that can learn any task) are theoretically impossible. Nevertheless, the hope remains that the creation of such black-box systems might be possible, at least in certain restricted domains. This would require specialists who are very well versed in AI, ML, robotics and perception. Unfortunately, specialists with such broad knowledge are rare. Universities continue to focus on educational specialisations, such as control theory, computer vision and machine learning. Few faculties and institutes stress a broad education in "perception-action-learning" systems.

Hardware developments: Hardware realisation is a prerequisite for autonomous robots. Working with the right robots and right sensors in real and complex environments poses problems that cannot be reasonably simulated. Of course, the quality of the hardware plays an important role. Hardware problems often make it necessary to improve algorithms or adopt completely different strategies of autonomous robotics. For example, the position sensors of some robots are not accurate enough to determine where exactly the hand is in 3D space. Computer vision algorithms that locate the hand with the help of 3D camera images are then needed to help out. Future developments in mechatronics are required to produce reliable high-quality, high-performance robots.

For robots on the nanometre to millimetre scale, the fabrication technology often represents the dominant research topic. Using material research methods, it is possible to develop techniques that generate thousands of robots at once. However, it is up to basic research to determine exactly what functions can be integrated into these miniscule robots. Investments often centre on the nano process technology that underlies such fabrication methods and has to run through many iterations until suitable systems can be synthesised.

Software developments: Appropriate hardware is clearly a prerequisite for autonomous robots; however, autonomy resides primarily in the software systems. And precisely those software systems are one of the core problems in autonomous robotics, as there are few standardised software systems. It is relatively easy to program a robotic function "somehow" in a feasibility study, such that a successful behaviour is produced at least once. This is what research labs and universities mostly achieve. But a robust autonomous system with reproducible behaviour, with perception, action and learning components and with complex algorithms quickly becomes too large to be realised and maintained by "software amateurs". True software engineering requires suitable personnel, i.e. researchers familiar both with software engineering and the problems of autonomous robotics. Software errors can have catastrophic consequences in robotics, for which reason entirely different and more conservative programming methods are called for. Unfortunately, there are few experts who combine all this knowledge.

TRUE SOFTWARE ENGINEERING REQUIRES SUITABLE PERSONNEL, I.E. RESEARCHERS FAMILIAR BOTH WITH SOFTWARE ENGINEERING AND THE PROBLEMS OF AUTONOMOUS ROBOTICS.

"Perception-action-learning" systems: Following these explanations about the state of research and technology in the field of autonomous robotics, it is also important to again highlight the foundations of these systems, i.e. integrated perception-action-learning systems. All three components of these systems are linked in a closed loop. Consequently, such systems cannot simply be divided into independent modules, as the quality of one module can easily affect the function of another. The hardware of a robot also being part of this loop, the software has to be adapted to the hardware. It is therefore no surprise that every robotics project quickly becomes unique, with the perception, control and learning algorithms generating very specific solutions that cannot be transferred to other robot systems. Therefore, one of the major goals of research and technology is to understand what building blocks and complex learning algorithms can generally be used to build autonomous perception-action-learning systems, so that the hardware, the environment and submodules autonomously adapt, analogous to the way in which biological systems deal with a constantly growing and ageing body.

SUMMARY

Experts in all technically advanced countries agree that robots, particularly autonomous robots, will become part of everyday life in human societies in the foreseeable future. However, a great deal of R&D work is still required to devise such systems. In all research areas of autonomous robotics, namely perception, control, adaptability, learning capability and mechatronics, there are significant unknown quantities that still do not permit robust autonomous robots to be created for everyday use. It remains a challenge for basic and applied research to develop general building blocks for autonomous robots. Scientific councils and political decision makers also face a challenge in providing suitable support for such highly interdisciplinary and complex research in this field.