

Robots of the Future

V K Wadhawan

Robots are a subset of “smart structures” – engineered constructs that can “think” and adapt to the environment. Robots can work even in environments not conducive to humans (e.g. areas with high levels of nuclear radiation). Intelligent robots will find a large range of applications in defence and nuclear industries, as also other scientific, technological, and commercial applications. Robotics is therefore a strategically important field of research. The question is: How smart or intelligent can the robots become? The answer is that perhaps there is no readily-conceivable upper limit. In fact, there is a strong section of opinion that, within the present century itself, robots will overtake humans in practically all aspects of mental and physical capability, and they will then evolve further at a rapid rate, with or without our help. Expert opinion is bound to be divided when we look too far into the future. Some scientists think that there is nothing much that we can do to control or prevent the *inevitable* and fast evolution of machine intelligence.

Introduction

Human beings try to develop machines which can make their own lives easier and richer. Robots are an example of this. There are two main types of robots: *industrial* robots, and *autonomous* robots. Industrial robots do useful work in a structured or pre-determined environment. They do repetitive jobs like fabricating cars, stitching shirts, or making computer chips, all according to a set of instructions programmed into them. The first industrial robot was employed by General Motors in 1961.

Autonomous or smart robots, by contrast, are expected to work in an *unstructured* environment. They *move around* in an environment that has not been specifically engineered for them (although



V K Wadhawan is a Raja Ramanna Fellow at the Bhabha Atomic Research Centre, Mumbai. He is the author of a book on ferroic materials (Gordon & Breach, 2000), and another one on smart structures (Oxford University Press, 2007). He is also the Regional Editor for Asia for the journal *PHASE TRANSITIONS* (Taylor & Francis).

Keywords

Autonomous robots, machine intelligence, smart structures, universal computers, artificial evolution, personal robots.



there are serious limitations at present), and do useful and ‘intelligent’ work. The work may involve avoiding obstacles, sensing the environment continuously (even doing image analysis, competently and quickly), and solving problems. They have to interact with a dynamically changing and complex environment, with the help of sensors and actuators and a brain centre. Some supervision by humans is also required, at least for the next few decades.

Learning, smartness, and intelligence are terms we normally associate with living beings. But if our machines are to help us in truly worthwhile ways, we must make them intelligent. There has already been some progress in that direction, but what we see at present is just the tip of the proverbial iceberg of the future. In this article we focus on autonomous robots and discuss the evolution of robotic intelligence in the decades to come.

Progress in the Development of Autonomous Robots

There has been a steady trend to move towards *evolutionary or adaptive robotics*. (Boxes 1 and 2 summarize the basics of biological and artificial evolution, respectively.) In adaptive robotics one aims at instilling robots with creative problem-solving capabilities. This means that they must *learn* from real-life experiences so that as time passes, they get better and better at problem-solving. The analogy with how a small child learns and improves as it grows is an apt one. The child comes into the world equipped with certain sensors and actuators (eyes, ears, hands, etc.), as well as that marvel of a learning and memory apparatus, the brain. Through a continuous process of experimentation (unsupervised learning), as well as supervised learning (from parents, teachers, etc.), and reinforced learning (the hard-knocks of life, and rewards for certain kinds of action), the child’s brain performs *evolutionary computation*.

Can we build robots which can do all that? That is a somewhat distant dream at present, but there has been progress. The robot ‘Darwin X’, developed by Krichmar *et al* [1], is a brain-based

In adaptive robotics one aims at instilling robots with creative problem-solving capabilities. This means that they must *learn* from real-life experiences so that as time passes, they get better and better at problem-solving.



Box 1. Biological Evolution

To understand the evolution of machine intelligence, it is important to be familiar with the basics of Darwinian evolution and some related topics. As Richard Dawkins has said, ‘Darwin’s theory of evolution by cumulative natural selection is the only theory we know of that is in principle capable of explaining the evolution of organized complexity from primeval simplicity’.

Darwin proposed that adaptation of populations to the environment was a necessary consequence of the interaction between the organisms and their surroundings. His theory of cumulative natural selection started with the observation that, given enough time, food, space, and safety from predators and disease, the size of the population of any species can increase with each generation. Since this indefinite increase does not actually occur, there must be limiting factors in operation. For example, if food is limited, only a fraction of the population can survive and propagate itself. Since not all individuals in a species are exactly alike and possess intrinsic differences, those which are better suited to cope with a given set of conditions will stand a better chance of survival (*survival of the fittest*). The fitter individuals not only have a better chance of survival, they are also likely to procreate more and transmit the adaptive traits to the next generation. Thus, attributes conducive to survival get naturally selected cumulatively at the expense of less conducive attributes over several generations. This is the process of *cumulative natural selection*.

It is now known that the inherited characteristics of the progeny are specified by *genes*. Genes programme embryos to develop into adults with certain characteristics and these characteristics are not entirely identical among the different individuals in the population. In other words, members of a population exhibit genetic diversity. Genes of individuals with characteristics that enable them to reproduce successfully tend to survive in the *gene pool* at the expense of genes that fail.

Can the environmental conditions in which the parents live directly affect the genetic characteristics of the offspring? The answer is ‘No’ according to the Darwinian theory of evolution, and ‘Yes’ according to the theory of Lamarck. The Lamarckian viewpoint of *inheritance of acquired characteristics* runs counter to the “central dogma” of modern molecular genetics, according to which information can flow from nucleic acids (genes) to proteins (responsible for phenotype or traits), but not from proteins back to nucleic acids (or from phenotype back to genes). The issue of Lamarckism in natural evolution continues to be debated. However, in *artificial evolution* (say, inside a computer), Lamarckism offers perfectly acceptable and viable scenarios (see *Box 2*).

Variations on original Darwinism have been suggested from time to time. Lamarckism also keeps asserting itself in interesting ways. For example, it is notable that, contrary to the usually long evolutionary time scales, there has been an exceptionally rapid expansion of the brain capacity of our recent ancestors, leading to the evolution of *Homo sapiens*, i.e. ourselves. The evolution of language, speech, and culture are the consequences of this rapid evolution of the human brain under strong selective pressure. Similar to the gene which is the basic unit of biological inheritance, Dawkins introduced the notion of the *meme*,

Continued...



which is the unit of cultural inheritance in his neo-Darwinistic theory. A meme may be a good idea, a soul-stirring tune, a logical piece of reasoning, or a great philosophical concept.

In Dawkins' scheme of things, two different evolutionary processes must have operated in tandem, one of them being the classical Darwinian evolution and the other centred around language, intelligence, and culture. Apes, which have been left far behind by humans in the evolutionary race, lack substantial speech centres in their brains. Emergence of language and speech in one of their species provided a *qualitative* evolutionary advantage to that species, with far-reaching consequences. How did this come about? The answer has to do, not only with gene pools, but also *meme pools*. The totality of genes of a population comprising a species constitutes its gene pool. The genes that exist in many copies in the population are those that are good at surviving and replicating. Through a reinforcement effect, genes in the population that are good at cooperating with one another stand a better chance of surviving. Similarly, memes in a population may jump from one brain to another, and the fittest set of memes has a better chance of surviving to form the meme pool of the population. They replicate themselves by imitation or copying. Cultural evolution and progress occurs through a selective propagation of those memes which are good at cooperating with one another.

Memes can evolve like genes. In fact, any entity that can replicate and can have a variation both in its specific features and reproductive success is a candidate for Darwinian selection. Mastery of a particular skill (say language) required a slight increase in brain size. Having developed this larger size, the brain could have got triggered to launch entirely new activities, some of which had evolutionary advantages. This resulted in a relatively rapid increase in brain capacity, because even Lamarckian evolution could contribute to the meme-related part of the coevolution of the gene pool and the meme pool.

device (BBD). It is an intelligent machine, modelled on the human brain. Experience with an unstructured environment makes it learn and remember, just as a child does. But there is still a long way to go.

Another such developmental robot is 'SAIL' [2]. Its goal is to autonomously generate representations and architectures which can be scaled up to more complex capabilities in unconstrained environments. Like babysitters for human children, it needs human *robot-sitters* for supervised learning. Such robots can 'live' in our company, and become smarter with time, in an autonomous manner (albeit under some human supervision).

In the traditional artificial intelligence (AI) approach to robotics, the computational work for robot-control is decomposed into a chain of information-processing modules, proceeding from



Box 2. Artificial Evolution

Darwin's theory of evolution of new species by cumulative natural selection continues to be challenged occasionally because of the seeming unlikelihood of evolution of complex organisms from very simple beginnings. Among other things, the theory has received very convincing support from a variety of AL (artificial life) simulations. These are designed, not just to simulate evolution, but to actually *be* evolutionary processes, albeit inside a computer. One of the latest in this class, and still continuing, is a series of computer experiments carried out with the computer code called **Avida** (Lenski *et al.* 1999, 2003, Zimmer 2005)*. In it, digital organisms (strings of commands, akin to computer viruses) are created and then allowed to evolve. Each individual has a ('circular') sequence of instructions that are executed sequentially (except when the execution of one instruction causes the instruction pointer to jump to another position); it also has a virtual CPU with two stacks and three registers that hold 32-bit strings. Each item in a sequence is one of 26 possible instructions.

Experiments were begun in Avida with an ancestor that could replicate (by producing tens of thousands of copies of itself in a matter of minutes, thus providing a highly speeded up *and fully recorded* version of natural evolution). The replicated digital bits can undergo mutations (imperfect replication) the way DNA does. They not only replicate and mutate, they also compete for resources, thus meeting all the three requirements for Darwinian evolution. The resources in this case are a supply of numbers. Most often the individual strings of self-replicating code are not able to do anything (perform any logic function) to a presented number. But, once in a while, a mutation (change of command line) in a replicated individual may give it the ability to, say, read the number and then produce an identical output. Or, in a more advanced (more *evolved*) case, the individual may be able to add two numbers correctly. There is a reward system for such evolved behaviour; the more complex the evolved behaviour, the higher is the reward.

In these simulations, *emergent behaviour* is also observed. Avida is not only able to confirm the soundness of Darwin's theory; it also provides answers to a variety of other evolutionary puzzles as well. For example: Why does a forest have more than one kind of plant? Why is sexual reproduction an important aid to the evolution of advanced species? Why should altruism (i.e. regard for other members of the species, as a principle of action) exist at all? More details about this computer-based approach to evolution can be found on the websites <http://dllib.caltech.edu/avida>, or <http://myxo.css.msu.edu/papers/nature2003>.

The current interest in computer-mediated evolution, working on self-replicating filaments of code, brings to the fore Freeman Dyson's assertion that metabolism and replication are logically separable. Avida works on replication (software), without involving metabolism (hardware).

* R E Lenski, C Ofria, T C Collier and C Adami, 'Genome complexity, robustness and genetic interactions in digital organisms', *Nature*, Vol.400, p.661, 1999; R E Lenski, C Ofria, R T Pennock and C Adami, 'The evolutionary origin of complex features', *Nature*, Vol.423, p.139, 2003; C Zimmer, 'Testing Darwin', *Discover*, Vol.26, p.28, 2005.



sensing to action. In a radically different approach adopted by the well-known robotist Brooks [3], a parallel is drawn from how coherent intelligence (swarm intelligence) emerges in a beehive or an ant colony (see *Box 3* on swarm intelligence). In such a ‘vivisystem’, each agent is a simple device interacting with the world with sensors, actuators, and a very simple brain. In Brooks’ *subsumption architecture*, the decomposition of the robot-control process is in terms of *behaviour-generating modules*, each of which connects sensing to action. Like an individual bee in a beehive, each behaviour-generating module directly generates

Box 3. Swarm Intelligence

What is common among a beehive, an ant colony, a shoal of fish, an evolving population, a democracy, a stock market, and the world economy? They are all spontaneously self-organizing, interacting, distributed, large systems. They are *adaptive* (i.e., whatever happens to them, they try to turn it to their collective advantage). There is no central command. The individuals (or *agents*) are autonomous, but what they do is influenced strongly by what they ‘see’ others doing. Here, the term ‘autonomous’ means that each member reacts individually according to internal rules and the state of its local environment.

The autonomous or semi-autonomous members are highly connected to one another through communication. There is a *network*, but not necessarily connected to a central hub. The population is large, and no one is in control. This means, among other things, that even if a section of the population were to be decimated, the system will adjust and recover quickly, and carry on. Kelly (1994)* calls such networked or ‘webby’ systems *vivisystems*.

To get a feel for the nature of vivisystems, let us see how a beehive functions as a single organism, even when there is no central command. How does it, say, select a new site for setting up a hive?

It has been known for long that, at the time of the year when the sources of honey are aplenty, a large colony of honeybees splits into two by a process called *swarming*: A daughter queen and about half of the population in a beehive stays behind, and the rest (including the queen bee) leave the old hive to start a new one *at a carefully selected site*. The sequence of events is roughly as follows:

1. Typically, out of a total of ~10,000 bees, a few hundred worker bees go scouting for possible sites. The rest stay bivouacked on a nearby tree branch, conserving energy, till the best new site has been decided upon.

* K Kelly, *Out of Control: The New Biology of Machines, Social Systems, and the Economic World*, Cambridge: Perseus Books, 1994.

continued...



2. The ‘best’ site for nesting is typically a cavity in a tree with a volume greater than 20 litres, and an entrance hole smaller than 30 cm².

3. The scout bees come back and report to the swarm about possible nesting sites by doing a *waggle dance* in particular ways. Typically there are about a dozen sites competing for attention. During the report, the more vigorously a scout dances, the better must be the site being championed.

4. Deputy bees then go to check out the competing sites according to the intensity of the dances. They concur with the scouts whose sites are good by joining the dances of those scouts.

5. That induces more followers to check out the lead prospects. They return and join the show by leaping into the performance of *their* choice.

6. By compounding emphasis (*positive feedback*), the favourite site gets more visitors, thus increasing further the number of visitors. Finally, the swarm takes the daughter queen, and flies in the direction indicated by mob vote.

As Kelly puts it: ‘It’s an election hall of idiots, for idiots, and by idiots, and it works marvellously’. The remarkable thing about distributed complex systems like beehives is *emergent behaviour*. For them, two plus two is not necessarily four; it can be, say, apples. *More is different*. In a network of bees or ants, each member is programmed in a very simple way. Sheer large numbers and effective communication and interaction result in intelligence (*swarm intelligence*); no single member is capable of achieving this alone.

Four distinct features of distributed existence make vivisystems what they are: absence of imposed centralized control, autonomous nature of subunits, high connectivity between the subunits, and webby nonlinear causality of peers influencing peers. Emergent behaviour in vivisystems is very common in nature and has far-reaching consequences in every case. Human intelligence has also been interpreted as emergent behaviour arising from the interaction and connectivity of the vast number of neurons (or ‘agents’), even though each agent, taken individually, is as unintelligent as can be. Thoughts, feelings, and purpose result from the interactions among these basic components.

some part of the behaviour of the robot. The tight (proximity) coupling of sensing to action produces an intelligent network of simple computational elements that are broad rather than deep in perception and action.

There are two further concepts in this approach: *situatedness*, and *embodiment*. Situatedness means the incorporation of the fact that the robot is situated in the real world, which directly influences its sensing, actuation, and learning processes. Embodiment means that the robot is not some abstraction inside a computer,



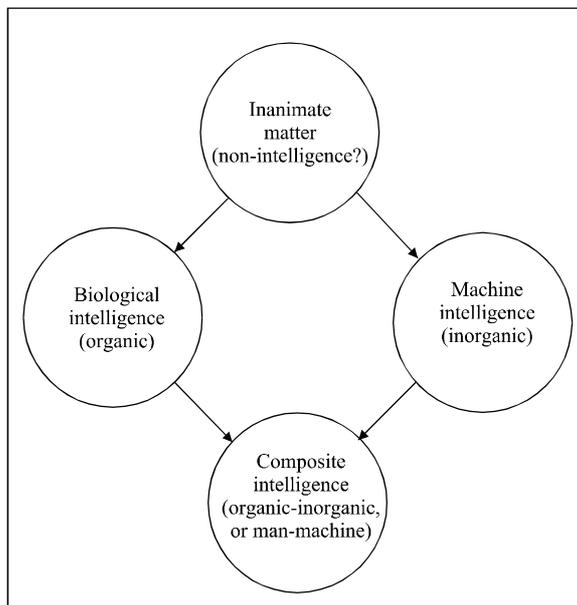


Figure 1. Evolution of intelligence. Biological or organic intelligence, and machine or inorganic intelligence, each has its strengths and limitations. The future course of evolution may be towards a composite (organic-inorganic, or man-machine) kind of intelligence (Wadhawan [4]).

millions of years, gave our body and large brain a highly developed proficiency for locomotion, navigation, talking, manipulation, recognition, and commonsense reasoning. These were essential for the battle for survival against competitors and predators. Only those of our early ancestors survived and evolved who were good enough at doing certain things repeatedly: procure food, outsmart predators, procreate, and protect the offspring. For achieving this, the human brain evolved into a *special* kind of computer, rather than a general-purpose or *universal* computer. Skills like number crunching are of recent origin, and therefore our brains are no match for present-day universal computers. Mathematical skills were not necessary for the survival of our early ancestors. For our machines, on the other hand, calculating is easier than reasoning, and reasoning is easier than perceiving and acting.

The principles of Darwinian evolution are now being used for evolving robot brains for specific tasks like obstacle avoidance. The robot adapts to the surroundings under the control of an algorithm which is not provided beforehand, but must evolve on an ongoing basis. The important thing about adaptive robots is that, although the various candidate algorithms (e.g. for obstacle

but has a body which must respond dynamically to the signals impinging on it, using immediate feedback.

Although adaptive robots need not be *humanoid robots*, research and development work is progressing in that direction also. Three important characteristics of such ‘creatures’ are walking, talking, and manipulation. Gait, voice, and optimal gripping are tough problems to tackle in robotics, although they are, quite literally, child’s play in the case of humans. There is an evolutionary underpinning to this situation (Figure 1). The evolutionary natural-selection processes, spread over hundreds of



avoidance, combined with the criteria for fastest movement and farthest foray from a reference point) are still created inside a computer, actual tests for evolutionary 'fitness' are carried out in the real world, *outside the computer*. In one of the approaches, this is done by mapping each of the candidate algorithms onto the robot brain and checking for fitness (in the evolutionary sense).

There have been several distinct or parallel approaches to the development of machine intelligence (*Figure 2*). The classical AI approach attempted to imitate some aspects of rational thought. Cybernetics, on the other hand, tended to adopt the human-nervous-system approach more directly. Evolutionary robotics embodies a convergence of these two approaches.

Moravec [5] has discussed in substantial detail the progress made in the development of autonomous robots. Evolutionary robotics involves a good deal of simulation and generalization work, rather like the creation of 'invariant representations' in the human brain. Now, suppose the robot is able to continuously update, using simulation algorithms, the knowledge about its own configuration, as well as that of the environment. Suppose further that the robot carries out this a little bit faster than the real rate of change in the physical world. The smart robot can then see the consequences of its intended action *before* taking the action! If the simulated consequences are not desirable, the robot would change its 'mind' about what would be a more appropriate course of action under the circumstances. Such a robot can be viewed as having an *inner life* or consciousness.

Low-Cost Universal Computers and Robots

Although progress in the development of smart robots has been slow because of a lack of adequate economic incentive for the

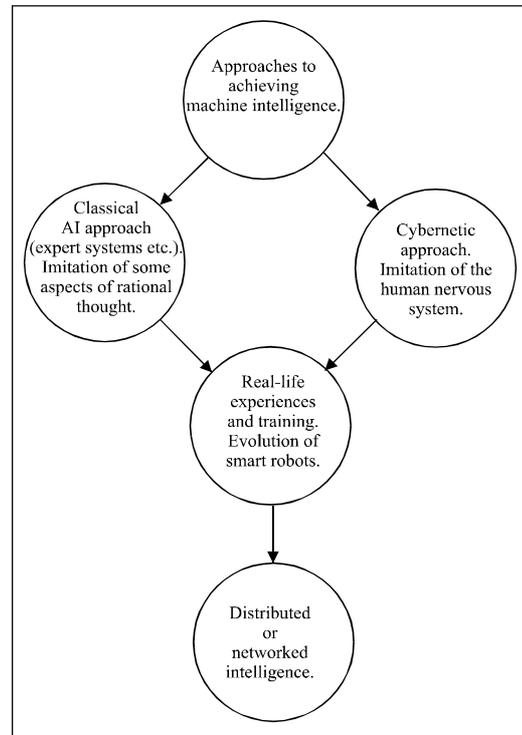


Figure 2. Approaches to machine intelligence. The classical Artificial Intelligence approach and the cybernetics approach get merged when we expose our autonomous robots to real-life experiences, and make them learn and evolve as intelligent machines. When there are several such interacting intelligences, distributed over space, and communicating and sharing information and inferences, a superintelligence can emerge ([4]).

There are three parameters to consider for the computational aspect of robotic action: The processing power (speed) of a computer, its memory size, and the price for a given combination of processing power and memory size.

robotics industry, there has been a major upswing in the recent past. There are three parameters to consider for the computational aspect of robotic action: The processing power (speed) of a computer, its memory size, and the price for a given combination of processing power and memory size. The processing power or computing power can be quantified in terms of ‘million instructions per second’ or MIPS. The size of the memory is specified in megabytes. By and large, the MIPS and the megabytes for a computer cannot be chosen independently: Barring some special applications, they should have a certain degree of compatibility (per unit cost), for reasons of optimal performance *in a variety of applications*. (See Box 4)

Box 4. Speed vs Memory

An analysis by Moravec [5] reveals that, for general-purpose or *universal computers*, the ratio of memory (the megabytes) to speed (the MIPS) has remained remarkably constant during the entire history of computers. A ‘time constant’ can be defined as roughly the time it takes for a computer to scan its own memory once. One megabyte per MIPS gives one second as the value of this time constant. This value has remained amazingly constant as progressively better universal computing machines have been developed over the decades. Machines having too much memory for their speed are too slow (for their price), even though they can handle large programs. Similarly, lower-memory higher-speed computers are not able to handle large programs, in spite of being fast.

Special jobs require *special computers* (rather than universal computers), entailing higher costs, as also a departure from the above universal ‘time constant’. For example, IBM’s Deep Blue computer, developed for competing with the chess legend Garry Kasparov (in 1996-97) had more speed than memory (~3 million MIPS and ~1000 megabytes, instead of the ‘universally’ optimal combination of, say, 1000 MIPS and 1000 megabytes). Similarly, the MIPS-to-megabytes ratio for running certain aircraft is also skewed in favour of MIPS. Examples of the other kind, namely slow machines (less MIPS than megabytes), include time-lapse security cameras and automatic data libraries.

Recently, Nvidia has introduced a chip, GeForce 8800, capable of a million MIPS speed, and low-cost, enough for use in commonplace applications like displaying high-resolution videos. It has 128 processors (on a single chip) for specific functions including high-resolution video display. In a multicore processor, two or more processors on a chip process data in tandem. For example, one core may handle a calculation, a second one may input data, while a third one sends instructions to an operating system. Such load-sharing and parallel functioning improves speed and performance, and reduces energy consumption and heat generation.



Moravec estimated in 1999 that the most advanced supercomputers available at that time were within a factor of 100 of having the power to mimic the human brain. But then such supercomputers come at a prohibitive cost. Costs must fall if machine intelligence is to make much headway. This has indeed been happening for a whole century. What about the future? How long can this go on? The answer is: For quite a while, provided technological breakthroughs or new ideas for exploiting existing technologies keep coming. An example of the latter is the use of *multicore processors*. Multicore chips, even for PCs, are already in the market.

Nanotechnology also holds great promise as the next-generation solution to faster and cheaper computation. DNA computing is another approach being investigated; this technique has the potential for massive parallelism.

Another factor hindering rapid progress in robotics has been the high costs incurred on sensors and actuators. Progress in nanotechnology (e.g. the development of MEMS) is resulting in continuously falling costs of sensors and actuators. It is now far less expensive to incorporate GPS (Global Positioning System) chips, video cameras and array microphones, etc., into robots.

Bill Gates has recently announced the development of universally applicable software packages by his company Microsoft that would further facilitate the use of ordinary PCs for controlling and developing robots of ever-increasing sophistication (see [6]). Many robots, especially research robots, already have PC-based controllers. It is anticipated that a large-scale move of robotics towards universally applicable PC-based architecture will cut costs and reduce the time needed for developing new configurations of autonomous robots. In the 1970s the development of Microsoft BASIC provided a common foundation that made it possible to use software written for one set of hardware to run on another set. Something similar is now going to happen in robotics.

One of the challenging problems faced at present in robotics is

Progress in nanotechnology (e.g. the development of MEMS) is resulting in continuously falling costs of sensors and actuators. It is now far less expensive to incorporate GPS (Global Positioning System) chips, video cameras and array microphones, etc., into robots.



One of the challenging problems faced in robotics is that of *concurrency*, namely how to process *simultaneously* the large amount of data coming in from a variety of sensors.

that of *concurrency*, namely how to process *simultaneously* the large amount of data coming in from a variety of sensors, and send suitable commands to the actuators of the robot. The approach adopted till recently was to write a ‘single-threaded’ program that first processes all the input data and then decides on the course of action, before starting the long loop all over again. This was not a happy situation because the action taken on the basis of one set of input data may be too late for safety etc., even though subsequent input data indicates a drastically different course of action.

To solve this problem, one must write ‘multi-threaded’ programs that can allow data to travel along many paths. This tough problem has been tackled by Microsoft by developing what is called CCR (‘concurrence and coordination run-time’). Although CCR was originally meant to exploit the advantages of multicore and multiprocessor systems, it may well be just the right thing needed for robots also. The CCR is basically a sequence of library functions that can perform specific tasks. It helps in developing multithread applications quickly, for coordinating a number of simultaneous activities of a robot.

Of course, competing approaches already exist, not only to CCR, but also to the so-called ‘decentralized software services’ (DSS). The latter enable the creation of applications in which the various services operate as *separate* processes that can be orchestrated just like we can aggregate images, text, and information from different servers on a Web page. The DSS need not reside entirely on the robot. It can be distributed over many PCs.

Although low-cost universal robots will be run by universal computers, their proliferation will have more profound consequences than those engendered by low-cost universal computers alone. Computers only manipulate symbols; i.e. they basically do ‘paperwork’ only, although the end results of such paperwork can indeed be used for, say, automation. A sophisticated universal robot goes far beyond such paperwork. It goes into perception and action in real-life situations. There is a far greater diversity of situations in the real world than just paperwork, and in far greater



numbers. There would thus be a much larger number of universal robots in action than universal computers. This, of course, will happen only when the cost per unit capability falls to low levels.

Networked Intelligent Systems

Vivisystems like beehives (see *Box 3*) continue to inspire more and more inventions in various branches of human activity and research. Distributed ‘perceptive’ networks are an example of this. Smart sensor systems that are small and cheap enough to be produced in large numbers and distributed over the area to be monitored have been developed (*Box 5*). They may or may not involve robots in the conventional sense at present, but robots are bound to come in very soon.

It is not mandatory that the brain and the sensors of a smart or intelligent structure be housed in one monolithic edifice. Let us see what can be the implications of such *pervasive computing* for applications concerning terrestrial phenomena like weather forecasting, animal and human migrations, etc. Suppose we let loose weather sensors with a certain degree of artificial intelligence all over the globe, all communicating with local brain centres, and indirectly with a centralized brain. Since these are intelligent artificial brains, they will form a world-view with the passage of time, as more and more sensory data are received, processed, and generalized. Since the sensors need not be like those of humans, the pattern formation in the artificial superbrain will be different from that in our brains. The machine-brain will recognize local patterns and global patterns about winds, etc., and there will be a perspective about weather patterns at different time scales: hours, days, months, years, decades. The crux of the matter is that the intelligent machine-brain will see patterns we humans have not and cannot.

With the ready availability of technologies like the Microsoft DSS (or other competing technologies) for writing in an easy manner the programs needed for distributed robotic applications, it is now possible for a network of wireless robots to tap into the

‘Capability, numbers sold, engineering and manufacturing quality, and cost-effectiveness will increase in a mutually reinforcing spiral’ (Moravec [5]).



Box 5. Motes – A Perceptive Network

Very simple computers and operating systems are linked by radio transceivers and sensors to form small autonomous nodes called *motes*. Running on a specially designed operating system called TinyOS (cf. Ananthaswamy; Culler and Mulder)* each mote links up with its neighbours. Although each unit by itself has only limited capabilities, a system comprising hundreds of them can spontaneously emerge as a *perceptive network*. Such networks, like other vivisystems, work on the principle that '*people learn to make sense of the world by talking with other people about it*' (Kennedy)*.

A network of motes is *not* just a network of miniaturized PCs, although each member does have some simple computing capability. The overall approach is based on the following considerations: Cut costs; conserve power; conserve space (miniaturize); have wireless communication (networking) among the nodes or agents; let collective intelligence emerge, like in a beehive; ensure robust, efficient, and reliable programming of a large and distributed network of motes; incorporate regrouping and reprogramming; include redundancy of sensor action to increase the reliability of the motes, keeping in view the fact that they may have to operate in hostile environmental conditions.

Let us see how motes are able to meet the objectives set out above. Each mote has its own TinyOS (like the tiny brain of a bee). This software runs on microprocessors that require very little memory (as little as eight kilobytes). A strategy for saving power (and for achieving something else with far-reaching consequences) is that of *multi-hop networking*. It amounts to giving each mote an extremely short-ranged radio transmitter, which saves power. A multi-tiered network is established, with motes in a particular tier or layer communicating only with those in the next lower and the next upper layer. Information hops from a mote to another mote only by one layer-step. This hierarchy is reminiscent of the neocortex of the human brain, although it is much simpler in concept and connectivity. It also makes room for parallelism, like in the human brain, and like in a beehive. If, for example, a particular mote stops functioning, there is enough redundancy and parallelism in the network that other motes reconfigure the connectivity to bypass that mote.

To update or replace the software for a network of motes, the method used is similar to how viruses or worms spread in PCs via the Internet. The new software is placed only in the 'root' mote, which then 'infects' the neighbouring motes with it, and so on, up the line.

* A Ananthaswamy, March of the motes, *New Scientist*, p.26, 23 August 2003.

D E Culler, and H Mulder, Smart sensors to network the world, *Sci. Amer.*, p.85, June 2004.

J Kennedy, Swarm intelligence, In Zomaya, A Y (ed.), *Handbook of Nature-Inspired and Innovative Computing: Integrating Classical Models with Emerging Technologies*, New York, Springer, p.187, 2006.

power of desktop PCs for carrying out tasks like recognition and navigation. This seemingly innocuous development will have far-reaching consequences, apart from cutting costs, for building and



operating robots. One can link wireless domestic or *personal robots* (PRs) to PCs. This will make it possible, for example, to keep long-distance tabs of what is being done by the PRs when one is away from home. A variety of surveillance applications can also be envisaged in the battlefield and elsewhere. The robots will of course communicate with one another also.

The PR revolution is round the corner. In a decade from now, we shall have affordable, remote-controlled, mutually-interacting PRs, the way we have PCs at present. And that will just be the beginning of the robots era.

Evolution of Machine Intelligence

Human intelligence continues to be something of an enigma (*Box 6*). But there is room for optimism that we shall be able to evolve machine intelligence of comparable sophistication in the near future. As argued by Moravec, a fact of life is that biological intelligence has indeed evolved from, say, insects to humans. There is a strong parallel between the evolution of robot intelligence and biological intelligence that preceded it. The largest nervous systems doubled in size every fifteen million years or so, since the Cambrian explosion 550 million years ago. Robot controllers double in complexity (processing power) every year or two. They are now barely at the lower range of vertebrate complexity, but should catch up with us within half a century. This will happen so fast because artificial evolution is being assisted by the intelligence of humans, and not just determined by the blind processes of Darwinian evolution (*Box 7*).

Robots of the Future

By the year 2050 or so, intelligent robots would have evolved to such an extent that they would take their further evolution into their own hands. The scenario beyond this crossover stage has been the subject of considerable debate. For example, the books by Moravec [5] and Kurzweil [7,8] continue to invite strong reactions.

Box 6. Human Intelligence

The human brain has $\sim 10^{11}$ neurons of various types. When the axon from one neuron touches dendrites of other neurons, connections called synapses are formed. A synapse is where the nerve impulse from one neuron is transmitted to another neuron. Formation and strengthening of synapses is what causes memories to be stored. A neuron has several thousand synapses. The human neural network is a massive parallel computational system.

According to Jeff Hawkins' recently formulated 'memory-prediction framework' for understanding human intelligence, the brain uses vast amounts of memory to create a *model* of the world. Everything we know or have learnt is stored in this model. The brain uses this memory-based model to make continuous *predictions* of future events. This ability to make predictions is the crux of intelligence.



Box 7. Goal-Directed Evolution

'Our intelligence, *as a tool*, should allow us to follow the path to intelligence, *as a goal*, in bigger strides than those originally taken by the awesomely patient, but blind, processes of Darwinian evolution. By setting up experimental conditions analogous to those encountered by animals in the course of evolution, we hope to retrace the steps by which human intelligence evolved. That animals started with small nervous systems gives confidence that today's small computers can emulate the first steps toward humanlike performance' (Moravec).

Moravec has estimated that artificial smart structures in general, and robots in particular, will evolve millions of times faster than biological creatures, and will surpass the humans in intelligence in the present century itself. At present this evolution in machines is being assisted by us. Unlike 'natural' Darwinian evolution (which is not goal-directed), artificial evolution has goals set by us. And Lamarckian evolution (adaptive change induced by the selective force) is not at all taboo in artificial evolution. Sophisticated machine intelligence, successfully modelled on the human neocortex, should be only a few decades away. This optimism stems partly from the observed trends in computational science. As mentioned above, the megabytes-to-MIPS ratio has remained remarkably constant (~1 second) during the entire history of universal, general-purpose computers. Extrapolation of this trend, as also projections about what lies in store after Moore's law has run its course, tell us that cost per unit computing power in universal computers will continue to fall rapidly. This will have a direct bearing on the rate of progress in the field of computational intelligence.

As Moravec [5] points out, already, machines read text, recognize speech, and even translate languages. Robots drive cross-country, crawl across Mars, and trundle down office corridors. Moravec also discusses how the music composition program EMI's classical creations have pleased audiences who rate it above most human composers. The chess program Deep Blue, in a first for machinekind, won the first game of the 1996 match against Gary Kasparov.

We mentioned notes above. Perceptive networks like these will be increasingly used for spying. Such distributed supersensory systems will not only have swarm intelligence, they will also undergo *evolution* with the passage of time. Like in the rapid evolution of the human brain (cf. *Box 1*), both the gene pool and the meme pool will be instrumental in this evolution of *distributed intelligence*. This ever-evolving superintelligence and knowledge-sharing will be available to each agent of the network, leading to a snowballing effect (*Figure 2*).

Moravec, from whose work I have quoted extensively, asks the



question: Is there a way to get our mind out of our brain? The mind is, after all, only a *pattern* of information. His 'bush robot' could possibly remove our mind from the biological brain, a little at a time, and transfer it to a machine. That would be total freedom: physical death, but mental immortality.

Summary and Concluding Remarks

Why is robotics such a hot field of research? We mention just a few economic and military reasons here:

Nations strong in smart-structures research will be able to cut production costs by letting smart robots do most or all of the work (not only manual work, but also design, planning, production, product innovation, sales strategies, mergers, buy-outs). Outer space will be colonized by sending out teams of robots (humans are too fragile for this), who will set up factories there to exploit solar energy and the mineral wealth. Only the fittest robot designs: companies, and nations will survive (evolution again!).

Possible defense applications include precise targeting and destruction of enemy installations. A nation with a highly advanced smart-structures technology will be able to minimize (and even eliminate) its own human casualties; just send smart robots to the battle field.

It appears inevitable that, aided by human beings, an empire of *inorganic life* will evolve, just as biological or organic life has evolved. We are about to enter a *post-biological world*, in which machine intelligence, once it has crossed a certain threshold, will not only undergo Darwinian and Lamarckian evolution on its own, but will do so millions of times faster than the biological evolution we are familiar with. The result will be smart or intelligent structures with a composite, i.e. organic-inorganic or man-machine intelligence (*Figure 1*). An important factor responsible for this rapid evolution will be the *distributed* nature of this networked intelligence.



As of now, we need to develop much more computing power for creating thinking robots, or for designing an artificial superbrain (envisaged, for example, by Isaac Asimov in 1950 in his science-fiction book *I, Robot*). Computing based on three-dimensional semiconductor design, and other kinds like DNA and quantum computing, will hopefully provide the necessary breakthroughs for achieving the computing power needed for creating truly intelligent artificial structures.

Stephen Hawking has commented that the present century will be the century of complexity. Complexity and evolution have a strong link. Evolution of really smart artificial structures will go hand in hand with our steadily improving understanding of complexity in nature. Apart from the economic and military power such developments will afford to nations which acquire a lead in these aspects of science and technology, there will also be a welcome fallout for the earth as a whole. The unprecedented levels of efficiency and economy achievable for manufacturing goods of all kinds with the help of (or entirely by) intelligent robots will not only bring prosperity for all, but will also have highly salutary effects on the ecology of our planet.

Address for Correspondence

V K Wadhawan
Raja Ramanna Fellow (DAE)
Solid State Physics Division
Bhabha Atomic Research
Centre
Mumbai 400 085, India
Email:
wadhawan@barc.gov.in

Suggested Reading

- [1] J L Krichmar, D A Nitz, J A Gally and G M Edelman, Characterizing functional hippocampal pathways in a brain-based device as it solves a spatial memory task, *PNAS, USA*, Vol.102, p.2111, 8 February 2005.
- [2] J Weng, In C W Chen and Y Q Zhang (eds.), *Learning in Computer Vision and Beyond: Development in Visual Communication and Image Processing*, New York: Marcel Dekker, 1998.
- [3] R A Brooks, New approaches to robotics, *Science*, Vol. 253, p.1227, 1991.
- [4] V K Wadhawan, *Smart Structures: Blurring the Distinction Between the Living and the Nonliving*, Oxford University Press, Oxford, 2007.
- [5] H Moravec, *Robot: Mere Machine to Transcendent Mind*, Oxford University Press, Oxford, 1999.
- [6] B Gates, A robot in every home, *Sci. Amer.*, p.48, January 2007.
- [7] R Kurzweil, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, Viking Penguin, New York, 1998.
- [8] R Kurzweil, *The Singularity is Near: When Humans Transcend Biology*, Viking Adult, New York, 2005.
- [9] F Fukuyama, *Our Posthuman Future*, Farrar, Straus and Giroux, New York, 2002.
- [10] M Minsky, *The Society of Mind*, Simon and Schuster, New York, 1986.
- [11] H Moravec, *Mind Children: The Future of Robot and Human Intelligence*, Harvard University Press, Cambridge, 1988.
- [12] S Nolfi and D Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*, MIT Press, Cambridge, Massachusetts, 2000.

