Solutions-Driven Marketing • Agents and Financial Risk Monitoring Systems •
Configurable Development Processes • Mobile Telephony in a Connected Life • Internet Time Out
• Techonomics • Managing International Data Communications

COMMUNICATIONS
of the ACM

March 2002 Volume 45,
Number 3

ROBOTS

Intelligence, Versatility, Adaptivity

Gaurav S. Sukhatme and Maja J. Matarić, Guest Editors

ROBOTS: INTELLIGENCE VERSATILITY, ADAPTIVITY

Expect robots to function on their own with people and each other under whichever environmental conditions they happen to find themselves.

Robotics research has come a long way over the past four decades—from the early days of manipulator control in manufacturing plants to today's mobile robots, many finding application in the unstructured, everyday world around us. Much of this progress has been incremental, but from time to time the research community has also achieved significant advances leading to robots with new capabilities. In addition, Moore's Law, whereby processor speed has consistently doubled every 18 months during the same 40 years, and the availability of smaller and cheaper sensors, actuators, and embedded communication devices, have helped make possible a new dimension in operational autonomy.

A key challenge in robotics is designing algorithms that allow robots to function autonomously in unstructured, dynamic, partially observable, and uncertain environments (such as within a home, a crowded shopping mall, or an unstable building after a disaster). Tantalizingly, nature offers a design model by offering examples of solutions to this problem; robust autonomous survival in the real world is demonstrated by the myriad of life-forms around us. Thus there is no shortage of existence proofs that it can be done; the challenge for robotics researchers is to achieve the same flexibility and autonomy with artificial systems.

Today, for the first time, robots are capable of functioning in novel situations, reasoning about and acting in relatively complex domains. Some are beginning to interact with humans, and in general there is a strong thrust in the field toward

The USC Autonomous Vehicle Aerial Tracking and Reconnaissance (AVITAR) helicopter coming in for an autonomous landing; the craft is not under the radio control of a human pilot.

research directly addressing robot autonomy in just such unstructured, dynamic, and uncertain environments, including those in social as well as physical settings. Indeed, Rodney Brooks, one of the authors in this section, predicts in his article "Humanoid Robots" that robots will be common in people's lives by the middle of the century if not significantly earlier.

Along with Brooks, we offer a number of articles surveying the state of the art in five emerging areas in robotics research pursuing autonomous operation. In this context, Brooks also discusses the fascination humans have had with humanoids through the ages, surveying recent progress in the field and emphasizing the research being conducted by his group at MIT. That group has focused on, for example, aspects of embodied social interaction between humanoid robots and people, leading to two of the best-known machines—Cog, a humanoid torso, and Kismet, a humanoid face. Brooks also makes the interesting point that the jury is still out on whether future socially useful robots will have human form or even need to present themselves in human form in the interests of application performance.

A key challenge is designing algorithms that allow robots to function autonomously in unstructured, dynamic, partially observable, and uncertain environments.

In "Self-Reconfiguring Robots," Daniela Rus et al. discuss the promise and flexibility of shape-adaptation to suit the environment, pointing out that adaptation is particularly important in unstructured environments. Their proposed solution involves constructing robots from elementary building blocks capable of autonomously organizing and reorganizing themselves to best fit the given environmental conditions and mission.

In "Robotics and Interactive Simulation," Oussama Khatib et al. describe recent developments allowing robots to work and cooperate with humans via haptics, or the sense of touch. These interactions are based on physical models with sufficient sophistication to recreate a complicated, physically consistent world, allowing natural and intuitive user interaction. Such interaction relies on dynamic simulation that also provides insight into the behavior of physical systems. They also discuss applications in virtual prototyping, animation, surgery, and more.

In "Probabilistic Robotics," Sebastian Thru describes a successful approach to robot control base on the theory of probability. Probabilistic representations are used for reasoning with the uncertain inherent in models and sensed information. Amon the more notable successes in mobile robotics are series of tour-guide robots, one of which has bee used in the crowded Smithsonian Museum in Washington, D.C.

Finally, in "Entertainment Robotics," Manuela M Veloso discusses the positive effects robot soccer ha had on robotics research, forcing researchers to explicitly address the design of cooperative behavior fo robots in dynamic, adversarial environments. Robo soccer has also shown the viability of robot games a entertainment.

Robotics today, however, covers an even greate scope than this eclectic collection of articles migh suggest. We have limited ourselves to a number representative areas of functionally autonomou robotics within the large and constantly expandin field. For example, we've avoided some other emerging areas, including distributed robotics, field robotics, and sensor networks. The field of distribute robotics studies methods for collective and cooperative robot team coordination for a vast variety tasks. Sensor-actuator networks represent an area research closely related to distributed robotics (as we as to the growing field of embedded systems) an serve as a bridge to the computer network and architecture disciplines. Field robotics deals with applications of robots in construction, space exploration, an agriculture. These areas are worth researchin and applying, in addition to the ones described i the articles in this section.

Though we provide a selective view of the emerging research, we want to impart a sense of today's cu ting edge. Written by some of the leading experts i the respective areas, they represent critical parts of th rapidly expanding field of robotics. Following th lead and serving as an inspiration for future desig ideas, the best is yet to come. ■

Gaurav S. Sukhatme (gaurav@cs.usc.edu) is an assistant professor in the Department of Computer Science at the University o Southern California, Los Angeles, and the associate director of the USC Robotics Research Laboratory; robotics.usc.edu.

Maja J. Matarić (mataric@cs.usc.edu) is an associate professor in the Department of Computer Science at the University of Southern California, Los Angeles, and the director of the USC Robotics Research Laboratory; robotics.usc.edu.

Rodney Brooks

The many moods of Kismet: anger, surprise, happy, aslee

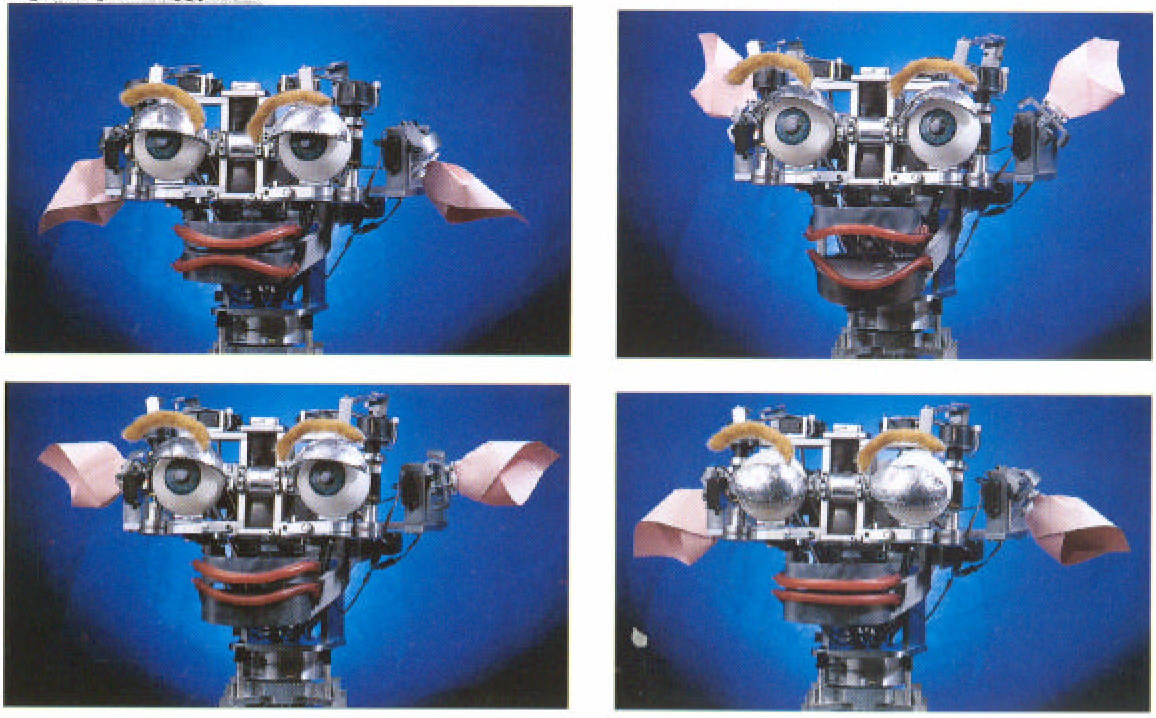

ETER MENZEL/ROBO SAPIENS/WWW.MENZELPHOTO.COM

HUMANOID ROBOTS

The future promises lots of robots in our everyday lives; some, perhaps many, of them could look and behave like people but only if being humanoid represents a technological advantage over their relatively utilitarian counterparts.

People have been interested in building "robots" in the form of humans for thousands of years. There were baked clay figures of humans in both Europe and China 7,000 years ago. At the height of Egyptian civilization 3,000 years ago, articulated statues could be controlled by hidden operators. At Thebes, the new king was chosen by an articulated statue of Ammon, one of the chief Egyptian gods (depicted as a human with a ram's head). Priests secretly controlled it as male members of the royal family paraded before it.

Leonardo da Vinci, the leading student of human anatomy in his time, designed a mechanical equivalent of a human—a humanoid robot—early in the 16th century; unfortunately, the design has still not been constructed.

Frenchman Jacques de Vaucanson early in the 18th century built three clockwork humanoids. One was a mandolin player that sang and tapped its foot as it played. Another was a piano player that simulated breathing and moved its head. A third

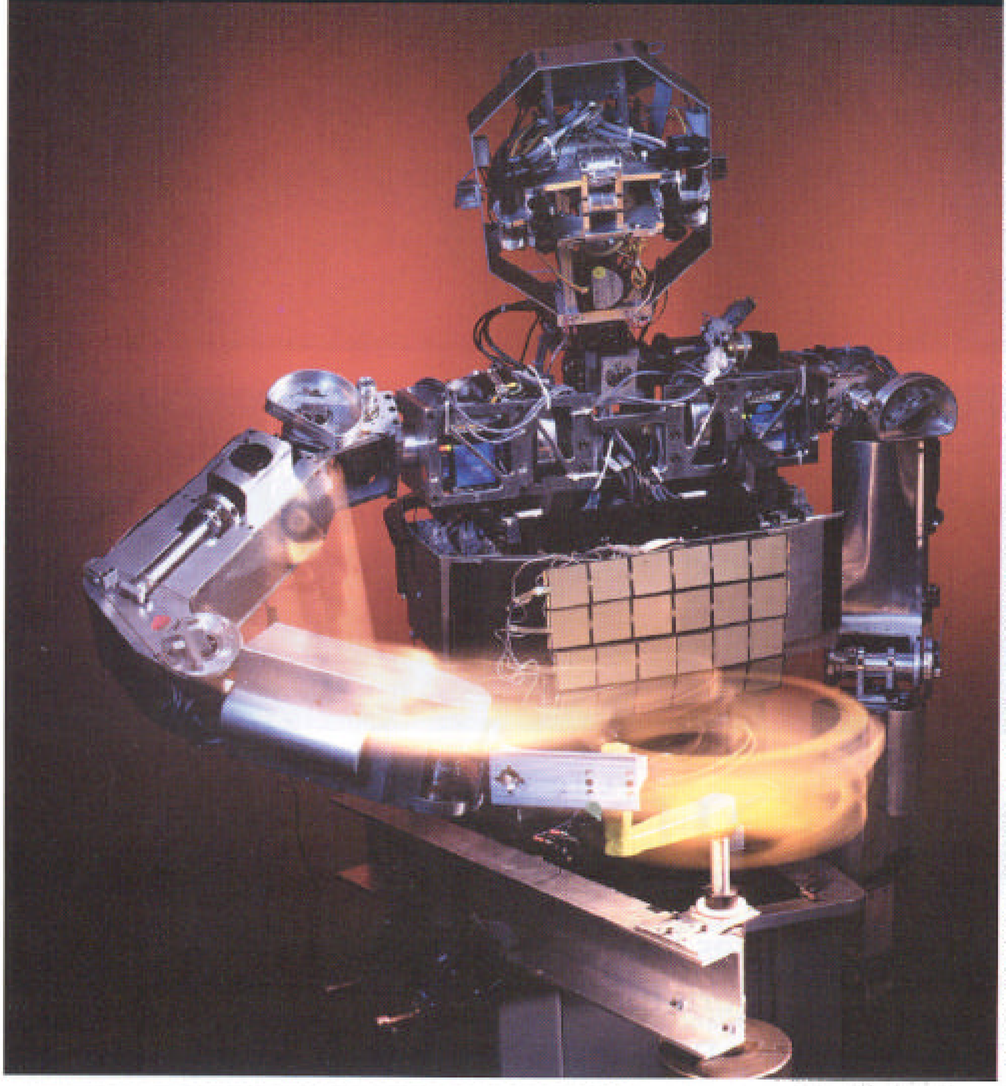

Figure 1. Cog is an upper-torso robot with two force-controlled arms, simple hands, and an active-vision head. It has undergone many revisions since 1993; different versions have appeared in the literature with different heads, arms, and hands.

SAM OGDEN

was a flute player. All were reported to be very lifelike, though none could sense the environment; all were simple playback mechanisms.

Similar humanoids soon followed. In the 18th century Pierre Jacquet-Droz, a Swiss watchmaker, and his son Henri-Louis built a number of humanoids, including a female organ player that simulated breathing and gaze direction, looking at the audience, her hands, and the music. Henri Maillardet, also a Swiss watchmaker, built a boy robot in 1815 that could write script in both French and English and draw a variety of landscapes.

Modern Humanoids

The modern era of humanoid robots was ushered in during the early 1970s by Hirokazu Kato, a professor at Waseda University in Tokyo; he oversaw the building of Wabot-1, a robot that could walk a few steps on two legs, grasp simple objects with its two hands, and carry out some primitive speech interaction with people. But, as with the early humanoids, Wabot-1 was still essentially a playback mechanism.

Kato's next robot, Wabot-2, built in 1984, was much more than a playback mechanism. Like Wabot-1 it had two legs and two arms. Unlike Wabot-1, it could not stand but rather sat on a piano bench. Its feet were used to press the pedals of an organ, and its arms and hands were restricted to playing the organ's keyboard. It had five fingers on each hand and could move its arms from side to side when playing the keys. Its head was a large TV camera; when sheet music was placed on the music stand above the keyboard, it would read the music and play the piece. In some sense it, too, was a playback mechanism, but it played back standard musical notation, perceiving such notation through its vision system and responding appropriately.

By the mid-1990s many humanoid robot projects were under way, most notably in Japan, Germany, and the U.S. Today, more than 100 researchers work in humanoid robotics at Waseda University alone and a similar number at Honda Corp. just outside Tokyo. There are also large humanoid projects at Tokyo University, the Electro-Technical Laboratory (ETL) in Tsukuba, Advanced Telecommunications Research (ATR) in Kyoto, and at other Japanese locations. Germany's Bundeswehr University of Munich and the Technical University of Munich have hosted humanoid robot projects. The major projects in the U.S. have been at the University of Utah, Vanderbilt University, NASA-Houston, and MIT.

Figure 2. Kismet is an active-vision head with a neck and facial features. It has four cameras (two in the steerable eyes and two wide-angle ones embedded in its face) and active eyebrows, ears, lips, and a jaw. Altogether, it includes 17 motors. A new-generation Kismet is under construction.

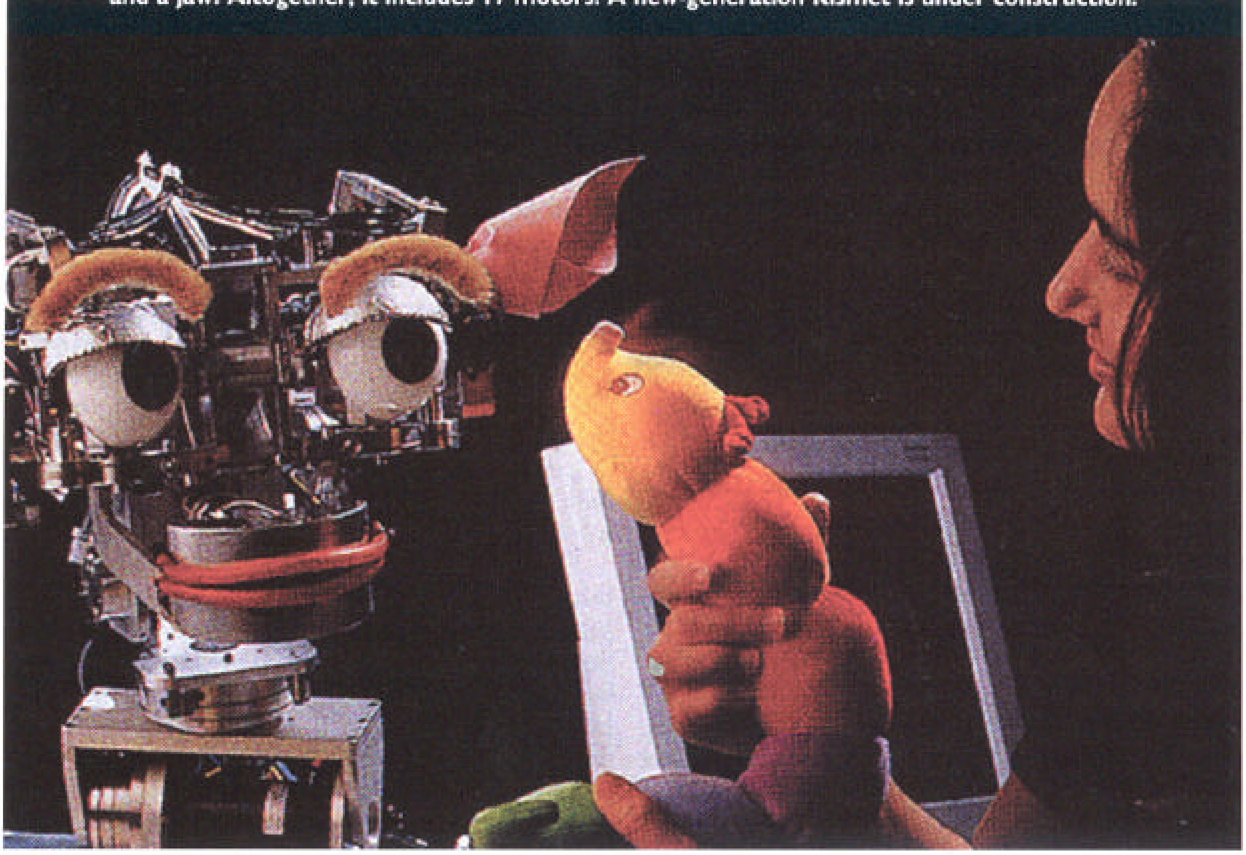

There have been many different motivations for building humanoid robots. Some formally announced ones include: investigating bipedal locomotion; building teleoperated robots to directly take the place of people (such as in spacewalks outside the International Space Station); building robots to maneuver in houses built to be convenient for people; investigating hand-eye coordination for tasks usually done by people; entertaining people; and functioning as a tool to study how people do what they do in the world.

MIT Humanoids

The humanoid robotics group at MIT (one of two groups in the Artificial Intelligence Laboratory working on humanoid robotics, the other concentrating on bipedal locomotion) started out developing humanoid robots as a tool for understanding humans' use of representations of the world around them [6]. Early plans were based on the work of the philosophers George Lakoff and Mark Johnson (best summarized in [8]) who posited that all of our understanding of the world builds upon the embodied experiences we have when we are young. For instance, they argued that the concept of affection uses warmth as a metaphor because children are exposed to the warmth of their parents' bodies when shown affection. Thus we might say, "They greeted me warmly." Likewise, we tend to use bigness as a metaphor for importance, as in "tomorrow is a big day," because parents are important, big, and indeed dominate our visual experience when we are young. Higher-level concepts are built as metaphors less direct than these primary ones but nevertheless rely on our bodily experience in the world. For instance, for time, we use the metaphor of moving forward, walking or running in a straight line. Thus the future is ahead of us, the present is where we are, and the past is behind us.

As the first humanoid robot, called Cog, was being developed in the mid-1990s, many aspects of perception and motor control had yet to be solved [5] (see Figure 1). Its developers realized there were important precursors to explicit representations of metaphors, as had been argued in earlier work on situated and embodied robots [4]. In the case of robots with humanoid form, intended to act in the world as people do, these precursors are social interactions [2], which are themselves based on emotional systems [7], facial expressions, and eye movements. The eye movements are driven by perceptual demands imposed by the underlying architecture of the eye [10]; in turn, they have been hijacked by evolution as significant components of human social interactions.

This realization prompted development of the robot Kismet in the late 1990s to study how social cues can be elicited from people by robots (see Figure 2). Today, both robots are used for researching aspects of social interaction.

Active Vision

Vision systems with steerable cameras that move in purposeful ways as part of the perception process are called active vision systems [1]. A humanoid vision system with the same basic mechanical structure as humans and other mammals and that follows the same motion primitives used by humans appears to be animate and lifelike.

> The robot's coherence of behavior is not determined by some internal locking mechanism but by its direction of gaze out into the world.

The human eye has a central fovea spanning about 5 degrees vertically and horizontally of the full 160 degrees the eye can see. The brightness and color receptors are much more densely packed in this area; more than half of the region of the brain that first processes signals from the eye is dedicated to the central 2% of the field of view. Humans move their eyes around rapidly, up to four times per second, to aim this high-resolution part of their eyes at whatever it is they are interested in. These rapid motions are called saccades and occur ballistically without feedback about their accuracy during their motion. They are under voluntary control, in that a person can consciously choose to saccade to a particular location, though most saccades are made completely involuntarily by some sort of attention mechanism. Something interesting is often in the low-resolution periphery of human perception, and the eye saccades to that target to see it with higher resolution.

Humans can also scan their eyes to follow something moving in their field of view. Called smooth pursuit, such scanning cannot be done voluntarily. People cannot scan their eyes smoothly from, say, left to right, unless there is a moving object they can lock onto and follow. Lastly, humans use their inner ears to detect head motion, feeding the signal forward to compensate with eye motion much more quickly than the vision system could track how the world appears to be slipping and compensate. This is known as the vestibular-ocular reflex.

These three capabilities—saccades, smooth pursuit, and the vestibular-ocular reflex—have been implemented repeatedly in both Cog and Kismet [5], operating with performance comparable to that of humans, though their cameras have much lower resolution overall than the human eye.

Humans also verge their eyes toward a target and estimate the gross depth by how far off parallel their eyes have to move to see the same point in space. Comparisons are then made between the images in the eyes to get a local relative depth map—the process of stereo vision. Cog and Kismet also have these capabilities and so are able to perceive 3D aspects of the world.

Cog and Kismet are able to detect human faces through a variety of methods [9] and estimate the gaze direction of a person by determining the direction their eyes are pointing. The robots are not able to do as good a job as the human visual system, however, but estimates with 3 to 5 degree accuracy are useful for social interactions.

Cog and Kismet each have their perception and control systems running on more than a dozen computers. There is no central executive and indeed no central locus of control for the robots. Nevertheless, they appear to be operating in a coherent manner. The low-level trick that allows this coherence to happen is the visual attention mechanism (see Figure 3), which determines where the robot is looking; where it is looking determines what all the low-level perceptual processes will see. That in turn determines which of the robot's behaviors are active. The robot's coherence of behavior is not determined by some internal locking mechanism but by its direction of gaze out into the world.

Social Interaction

The Cog and Kismet visual systems are the bases for their social interactions. Even a naive human observer can understand what the robots are paying attention to by the direction of their gaze. Likewise, the robots can understand what a person is paying attention to by the direction of the person's gaze [9].

The visual attention system makes it completely intuitive for naive users to direct the robot's visual attention system to some particular objects. Cynthia Breazeal, now at the MIT Media Laboratory, describes a series of experiments in which subjects were asked to get the robot to pay attention to different objects [2]. Typically, they would bring the object into the field of view of the robot, then shake it and move it to the desired position, with the robot now smoothly pursuing it, paying attention to what the human subject wanted. The experimental subjects had no knowledge of how the robot's visual system operated but were able to use the same strategies they would use with a child, and they worked.

By manipulating the weighting the visual system applies to different visual cues, Kismet's high-level behaviors, such as dialogue turn-taking, can make it make or break eye contact, so indirectly, these high-level behaviors regulate social interaction. Moreover, Kismet expresses its internal emotional state through facial expressions and prosody in its voice. So, for instance, when someone comes very close to Kismet or waves something very quickly near its face, Kismet becomes more fearful. That emotional state is reflected in its posture; it draws back. This reaction triggers a complementary reaction in naive human subjects who also tend to draw back. Thus, Kismet, indirectly through its emotional system and its expression in the world, is able to manipulate people in social settings, just as humans unconsciously manipulate each other.

Figure 3. The visual attention system used in Kismet and Cog. The robots pay attention to objects with skin color, bright colors, or fast motion. Higher-level behaviors determine how these factors are weighted together; the eye motor system then saccades toward the most interesting part of the image. A habituation signal makes any interesting feature eventually appear less interesting, allowing the robot to pay attention to something new.

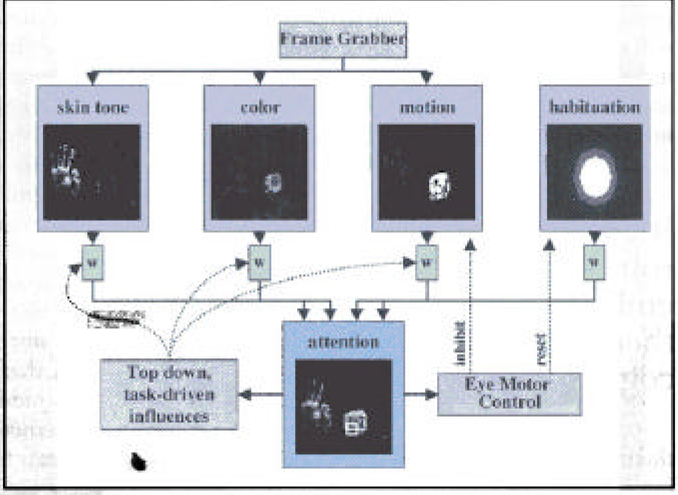

Kismet is also able to detect basic prosody in the voices of people and classify their speech as "praising," "prohibiting," "bidding for attention," or "soothing," four basic prosodic signals used in almost all human cultures by mothers with their babies. Kismet's detection of these cues changes its internal emotional state in appropriate ways; its outward demeanor changes, coupling in people who then intuitively react in appropriate ways.

Breazeal has reported on a number of experiments with human subjects [2]. Naive subjects in one set of experiments sat in front of the robot and were instructed to "talk to the robot." The robot could understand only their prosody and not the actual words they said. The robot generated speech with prosody, though it was always random strings of English phonemes with no intrinsic meaning.

Most subjects were able to determine when it was their turn to speak, but some did not know what to say. Others engaged in long conversations with the robot—even though there was no conventional linguistic transfer. The more basic social interactions often masked the lack of actual language. For instance, in one session a human subject said, "I want you to take a look at my watch," and Kismet looked right at the person's watch. The person had drawn up his left wrist to be in Kismet's field of view, then tapped his right index finger to the face of the watch. That was a sufficient cue to attract Kismet's attention system, and Kismet saccaded to the watch. Just as in human-to-human communication, layers of social interaction smoothed the process.

Because Kismet's processing system made it a little slower at turn-taking than a human, careful examination of the video record showed frequent turn-taking errors (where the robot or the person interrupted the other) at the start of each session, but also that people soon adapted (the robot did not adapt), and that after a few minutes the errors were significantly less frequent. Video clips of many of these experiments are available at www.ai.mit.edu/projects/humanoid-robotics-group.

Humanoids Everywhere?

The first few domestic robots are already on the market, including lawnmowing robots and home floor-cleaning robots. All are easy to use, which will be very important as the functionality of domestic robots is developed further. We can compare robot ease of use with computer ease of use. There are two sorts of computers in people's homes: One is embedded processors in television sets, coffee machines, and practically any tool or appliance powered by electricity; they are trivial to interact with and induce almost no cognitive load on the human user. The other is home PCs with thousands of options that can be quite difficult to understand; they produce high cognitive loads on users. It would be desirable for robots to follow the path of embedded processors, rather then PCs, and produce little cognitive load. However, unlike today's embedded processors, robots will be highly visible because they will move around in home environments. Therefore, it will be desirable for them to understand human social conventions, so they can be unobtrusive; meanwhile, humans should be able to interact with them in the same kind of noncognitive ways they interact with other humans. For instance, it will be useful for a large mobile appliance and a person to be able to negotiate who goes first in a tight corridor with the same natural head, eye, and hand gestures all people understand already.

> It would be desirable for robots to follow the path of embedded processors, rather then PCs, and produce little cognitive load on users.

Should we expect these sociable robots to have humanoid form and be as commonplace in our lives as a number of Hollywood fantasies have portrayed? It is difficult to know today, but there are two compelling, and competing, arguments on opposite sides of this question:

The current infatuation with humanoid robots is a necessary but passing phase. It allows researchers to get at the essence of human-robot interactions, but the lessons learned will ultimately be applicable to robots with much more functional forms. For instance, we can expect driverless trucks in our residential neighborhoods. When human drivers stop at an intersection as other vehicles pull up on the cross street, they often engage in informal social interactions through eye contact, head nodding, and finger motions—social interactions ignored in the formal driving rules but that form a negotiation as to which driver should proceed first.

When another vehicle is a driverless truck, similar sorts of social negotiations should be possible to lubricate the safe flow of traffic. However, current experiences with humanoid sociable robots may well lead to development of social signals for the robot truck requiring no human form but rather signals that can tap into the same subconscious cues used by humans and by humans and humanoid robots.

It may be that the large number of humanoid robot projects under way today, especially in Japan, may produce enough successful prototype robots that people will find them naturally acceptable and expect them to have human form. It has become well understood over the past 20 years that the technologically superior solution may not be the one that wins out in the marketplace (in the same way the VHS video format won out over the Beta format). Rather, it depends on early market share and the little-understood dynamics of adoption. For this reason, humanoid robots might become common by accident. Or it may turn out there will be a discovery (not yet made) that they have some significant advantage over all other forms, and they will be common precisely because they are technologically superior.

The weight of progress in so many forms of robots for unstructured environments leads to the conclusion that robots will be common in people's lives by the middle of the century if not significantly earlier. Whether significant numbers of them will have human form is an open question.

REFERENCES

1. Blake, A. and Yuille, A., Eds. *Active Vision.* MIT Press, Cambridge, MA, 1992.
2. Breazeal, C. *Designing Sociable Robots.* MIT Press, Cambridge, MA, 2001.
3. Brooks, R. *Cambrian Intelligence.* MIT Press, Cambridge, MA, 1999.
4. Brooks, R. *Flesh and Machines: How Robots Will Change Us.* Pantheon, New York, 2002.
5. Brooks, R., Breazeal, C., Marjanovic, M., Scassellati, B., and Williamson, M. The Cog Project: Building a humanoid robot. In *Computation for Metaphors, Analogy, and Agents,* C. Nehaniv, Ed. *Lecture Notes in Artificial Intelligence 1562.* Springer Verlag, New York, 1999, 52–87.
6. Brooks, R. and Stein, L. Building brains for bodies. *Autonom. Robo. 1,* 1 (1994), 7–25.
7. Damasio, A. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness.* Harcourt Brace, New York, 1999.
8. Lakoff, G. and Johnson, M. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Philosophy.* Basic Books, New York, 1999.
9. Scassellati, B. *Foundations for a Theory of Mind for a Humanoid Robot.* MIT Artificial Intelligence Laboratory, Ph.D. dissertation, 2001.
10. Yarbus, A. *Eye Movements and Vision.* Plenum Press, New York, 1967.

RODNEY BROOKS (brooks@ai.mit.edu) is Fujitsu Professor of Computer Science and Engineering and Director of the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology, Cambridge, MA.

Daniela Rus, Zack Butler, Keith Kotay,
and Marsette Vona

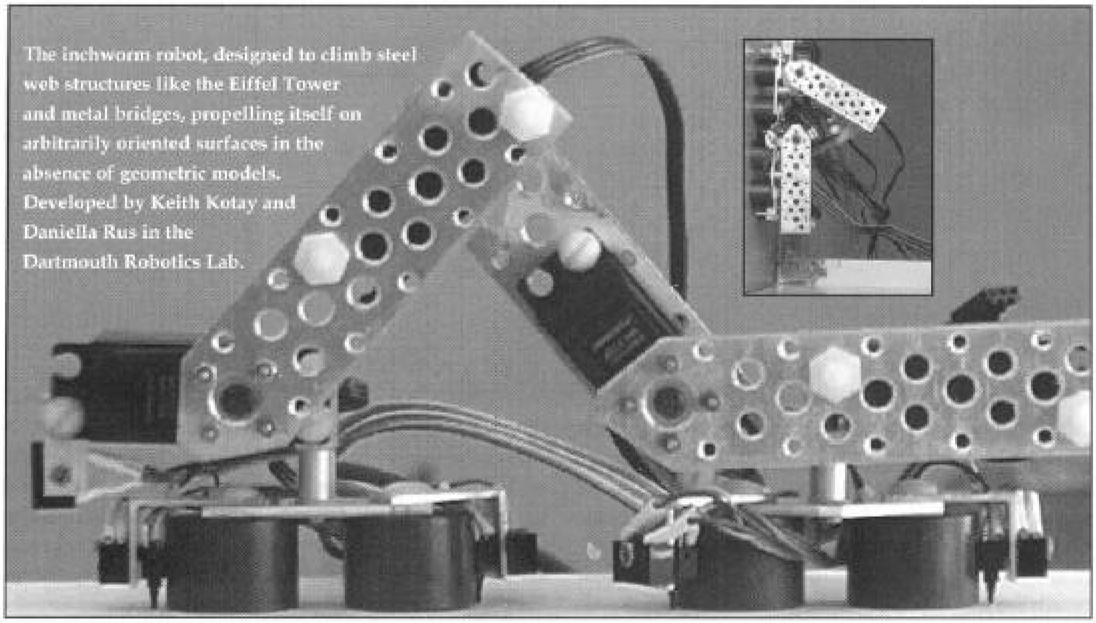

The inchworm robot, designed to climb steel web structures like the Eiffel Tower and metal bridges, propelling itself on arbitrarily oriented surfaces in the absence of geometric models. Developed by Keith Kotay and Daniella Rus in the Dartmouth Robotics Lab.

SELF-RECONFIGURING ROBOTS

Mimicking the adaptability of living biological cells, robot modules will reconfigure themselves toward a common purpose within the limits imposed by the local environment.

A robot designed for a single purpose can perform some specific task well but poorly on some other task, especially in a different environment. A fixed-architecture machine is acceptable if the environment is structured, but for tasks in unknown environments, a robot with the ability to change shape to suit the new environment and the required functionality is more likely to succeed than is its fixed-architecture counterpart.

Our vision for ultimate robotic design and functionality is to create vastly more versatile robots by building them out of simple modules with general self-reconfiguration capabilities. Hundreds of small, relatively limited, modules would be able to autonomously organize and reorganize as geometric structures to best fit the terrain on which the robot has to move, the shape of the object the robot has to manipulate, or the sensing needs for the given task. This function mimics the organizational principles of biological systems; a simple living cell does not do too much by itself, but a collection of cells can be organized to form a climbing inchworm, a crawling crab, or a running child. Large collections of small robots

Figure 1. Two different Molecule robots, each consisting of four modules, each composed of two atoms connected by a right-angle rigid bond. The connectors are implemented with electromagnets and a gripper mechanism. The Molecule has two rotational degrees of freedom about the bond and one rotational degree of freedom per atom about a single inter-Molecule connector.

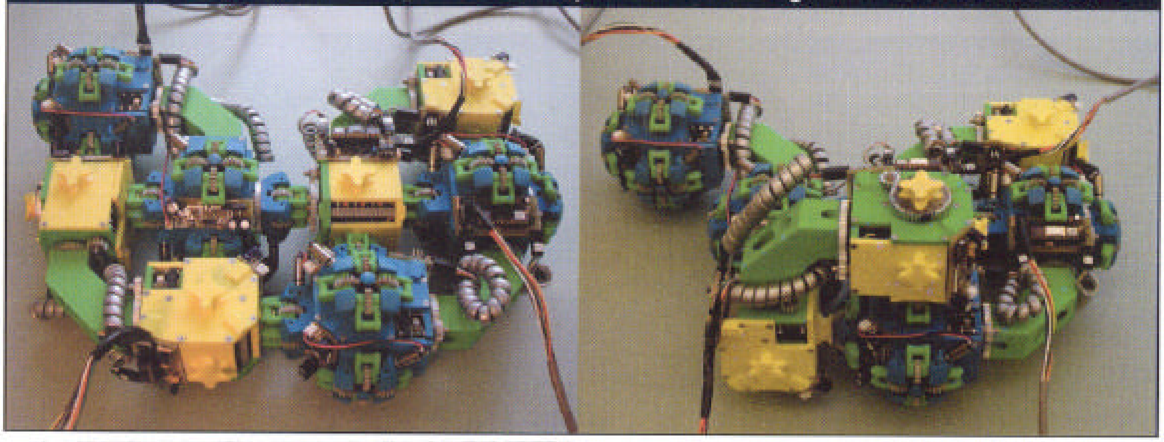

will someday actively organize themselves as the most optimal geometric structure under the local environmental conditions to perform coordinated and useful work, such as repairing a damaged structure, transporting artifacts, or manipulating surface properties.

Modular reconfigurable robots have a number of other advantages over their more traditional, fixed-architecture counterparts. First, they support multiple modalities of locomotion and manipulation. For example, if a particular robot system needs to climb stairs, it would reconfigure itself so it can crawl up stairs. If it needs to cross a gap or reach a window, just-in-time bridges and towers might be created. This adaptability can be achieved by requiring the robot to metamorphose from one shape into another to best match the shape of the terrain in a statically stable gait. Second, these robots are more fault tolerant than their fixed-structure counterparts. For example, if one module fails, a modular reconfigurable robot carrying some additional modules might excise the failed part and replace it with a spare unit. Third, they can be used in tasks requiring self-assembly, such as a structure in outer space or on the floor of the deep ocean. Fourth, they can provide a means for physically modeling 3D data—a real breakthrough in visualization. It is now common to use software to create computer models for 3D data. We envision creating physical models that can be manipulated directly by using modular self-reconfiguring robots. The robot would morph itself into the geometric shape dictated by the data. The user would be able to touch and manipulate the physical model provided by the robot. Ultimately, when the technology is available at a very small scale (even down to nano-scale), ordinary objects, say, a lamppost, would be able to reorganize themselves on demand into other ordinary objects, say, a bench or barricade.

Self-reconfiguration represents a paradigm shift for studying the fundamental principles of organization and reorganization in physical systems. In robotics, self-reconfiguration defines a rich class of questions about designing, controlling, and using massively distributed systems of robots. Self-reconfiguration also offers fertile ground for applying existing concepts in novel ways. For example, the domain involves the need for two kinds of planning algorithm: one to achieve a desired geometric shape and one to globally move the resulting shape in any physical direction. To explore the value of these algorithms and their applications, the following sections cover some of the hardware and algorithmic challenges of achieving task versatility with self-reconfiguring robots.

Hardware Systems

Self-reconfiguring robot systems employ physical mechanisms allowing modules to dynamically and automatically configure and reconfigure themselves into more complex forms. Creating self-reconfiguring robot systems poses many engineering challenges centered on designing the basic self-reconfiguring module and inter-module connections, as well as on how to aggregate distributed systems from the modules.

We want the basic module to be as small and simple as possible in terms of physical size and numbers of components, linkages, and functions, because the smaller the module, the greater the range of shapes that can be built from it. The modules should also be able to function independently of one another. Although simple modules are easier to design and build, simplicity may constrain functionality. Simplicity is also a key consideration in designing the inter-module connection mechanism. Because the modules make and break connections frequently, the connections should be simple and reliably independent of human intervention.

Other important design issues include communication between modules, actuator power, and the method used to supply electrical power to the system. Inter-module communication is necessary for module cooperation and distributed control of the various components in a system. A good connection mechanism can also be used to transmit messages between modules. Actuator power is the amount of force actuators have to exert to move the modules around; minimally, modules must be able to move their own weight. In a 3D system, modules must be able to move their own weight against the force of gravity. But supplying electrical power to modules is difficult and costly in terms of design and fabrication. Moreover, if the modules supply their own power using batteries, their weight and size increase, thus requiring more power to move them around. One possibility is to use the connection mechanism to simultaneously transmit power to all modules.

Several research groups have developed solutions to designing self-reconfiguring robots. For example, Toshio Fukuda et al. of Nagoya University first proposed in 1990 the idea of a cellular robotic system in which a set of specialized modules called "cells" are coordinated to accomplish a complex task [1]. The first hardware systems were developed a few years later at the Mechanical Engineering Laboratory (MEL) in Japan, as well as at Johns Hopkins University and Xerox PARC; the sophistication of the devices has increased ever since. Meanwhile, several groups have contributed groundbreaking ideas, such as the expansion/contraction actuation mechanism of [6], the connection mechanism of [3], and the deformation-based actuation of [4]. A good overview of the state of the field is presented in [5].

Most proposed unit-reconfigurable robot systems are actuated by rotating a module relative to the rest of the robot or expanding and/or contracting a module; their connection mechanisms are magnetic or mechanical. Modular robots are characterized as either homogeneous (all modules are identical) or heterogeneous (the modules are different). For example, Mark Yim of Xerox PARC defined a reconfiguring system to be "unit modular" if it is homogeneous. Most existing self-reconfiguring robot systems are based on the unit-modular approach.

Our own work has focused on the principle of mechanical simplicity, or the simplest design with the fewest components to accomplish the job. We've also characterized the properties that would confer a unit-modular robot system with self-reconfiguration [6]. Guided by these results, we've developed two unit-modular systems: the Molecule robot and the Crystal robot. The main goal of the former is self-reconfiguration in 3D. The latter uses a novel actuation mechanism, called scaling, that allows an individual module to double in size by expanding or halve its size by contracting, thus providing more robust motion than the previous rotation-based actuation systems. Instead of moving individual modules by rotating them, we change the overall robot shape by expanding and contracting the modules.

The Molecule robot. A Molecule robot consists of multiple units called Molecules, each consisting of two atoms the size of an apple linked by a rigid connection called a bond [2] (see Figure 1). Each atom has five inter-Molecule connection points and two degrees of freedom. One degree of freedom allows the atom to rotate 180 degrees relative to its bond connection; the other allows the atom (thus the entire Molecule) to rotate 180 degrees relative to one of the inter-Molecule connectors at a right angle to the bond connection. The current design uses R/C servomotors for the rotational degrees of freedom. A feature of the prototype is the use of a gripper-type connection mechanism for linking the various individual molecules.

Ordinary objects, say, a lamppost, would be able to reorganize themselves on demand into other ordinary objects, say, a bench or barricade.

The rotating connection points on each atom are the only ones required for Molecule motion. The other connection points are used for attachment to other Molecules in order to create stable 3D structures. Each Molecule also contains a microprocessor and the circuitry needed to control the servomotors and connectors. Each atom is four inches in diameter, and the Molecule weighs three pounds. Each individual Molecule has three basic motion capabilities: linear motion in a plane on top of a lattice of identical Molecules, irrespective of the absolute orientation of the plane; convex 90-degree transitions between two planar surfaces composed of Molecules; and concave 90-degree transitions between two planar surfaces composed of Molecules.

These primitive motions for a Molecule relative to a substrate can be combined and sequenced by the robot control system to achieve global motions for the entire robot. We've found that a four-Molecule robot is the smallest one that can move in general ways in the plane [2]; the smallest one that can also climb stairs is an eight-Molecule robot. Figure 1 shows our prototype four-Molecule robot.

The Crystalline atom. The Crystal is our novel self-reconfiguring module, a mechanism with some of the same motive properties as biological muscles that can be closely packed in 3D space and that can attach themselves to similar units (see Figure 2). Each of the 24 Crystal modules we've developed so far is actuated by expansion and contraction of its four faces. By expanding and contracting the neighbors in a connected structure, an individual module can be moved in general ways relative to the entire structure. Crystal atoms never rotate relative to each other; their relative movement is actuated by sliding via expansion/contraction. This basic operation has yielded new algorithms for global self-reconfiguration planning [6]. When fully contracted, each atom is a square measuring two inches on each of its sides. When fully expanded, each atom is a square with a four-inch side. Each atom weighs 12 ounces. Crystal robot systems can realize a wide range of geometries; for example, Figure 2 (top) shows a dog-shaped object transforming itself into a couch-shaped object.

Crystalline robot systems are dynamic structures. They move using sequences of reconfigurations to implement locomotion gaits and undergo shape metamorphosis. The dynamic nature of these systems is supported by the ability of individual modules to move globally relative to the entire structure. Unlike all previously proposed unit modules, in which modules are relocated only by traveling on the surface of a structure, Crystalline atoms can be relocated by traveling throughout the volume of a Crystal. Instead of propagating the module along the surface of the robot structure (requiring a linear number of expansion and contraction operations in the modules on the surface of the cube), the same goal can be achieved using a constant number of internal expansion and compression operations [6].

Algorithmic Issues

The algorithmic challenges involved in achieving self-reconfiguring robotic systems in a distributed fashion concern the metamorphosis of a given structure into a desired structure and how to use self-reconfiguration to implement multiple (adaptive) locomotion and manipulation gaits. These issues can be formulated as motion-planning problems. The key observation for automated planning is that most self-reconfiguring systems consist of identical modules. Since all modules are also interchangeable, it is not necessary to compute goal locations for each element. Thus, self-reconfiguration is different from the related intractable warehouse problem in which modules are assigned unique IDs and have to be placed at desired locations. Several groups have proposed architecture-dependent planners [2–4, 6, 8, 9]. This work can be divided into two categories of approaches: centralized and decentralized. The former is easier to analyze for performance guarantees but is not scalable for large robots. The latter supports parallelism but is generally more difficult to analyze.

Most approaches to planning have two parts: a set of device-level primitives for controlling the motion of one module relative to a structure or substrate; and general planning algorithms built by "composing" these primitives. Two of our centralized planning approaches are described in detail in [2, 6]; other groups [3, 4, 7–9] have pursued similar approaches.

Some of the most interesting future applications for self-reconfiguring robots promise to employ thousands of modules working together. Such systems represent ultra-high-degree-of-freedom systems that might be able to synthesize a robotic pet or a couch at one's request. However, centralized planning algorithms move one module at a time and may be too slow and impractical for controlling robots made of thousands of modules. For this reason, it is important to consider distributed planners that are scalable, support parallelism, and are better suited for operation in unstructured environments.

Distributed planning for Crystal robots. One possible approach is an algorithm called PACMAN distributed control we've developed for unit-compressible systems like our Crystal robots. An overall desired shape, or locomotion gait, is given to the robot's modules, each of which then determines whether or not it needs to move using only local information. If motion is necessary, each module initiates a path search through its fellow modules using only local information at each step. After a path is found, it is instantiated by marking each atom through which the path travels. We therefore developed data structures called "pellets" as a way to mark the path a module should follow to perform its part of the reconfiguration. This process could be viewed as the engineering equivalent of Hansel and Gretel's bread crumbs left through the forest, but since the pellets are permanent, the result is better. Because of the Crystal's unique actuation principle, a single physical module does not follow the entire path. Rather, it exchanges identities with other modules along the path, so it appears to 'follow' the entire path while actually moving only locally. Additionally, by marking each pellet with the identity of the

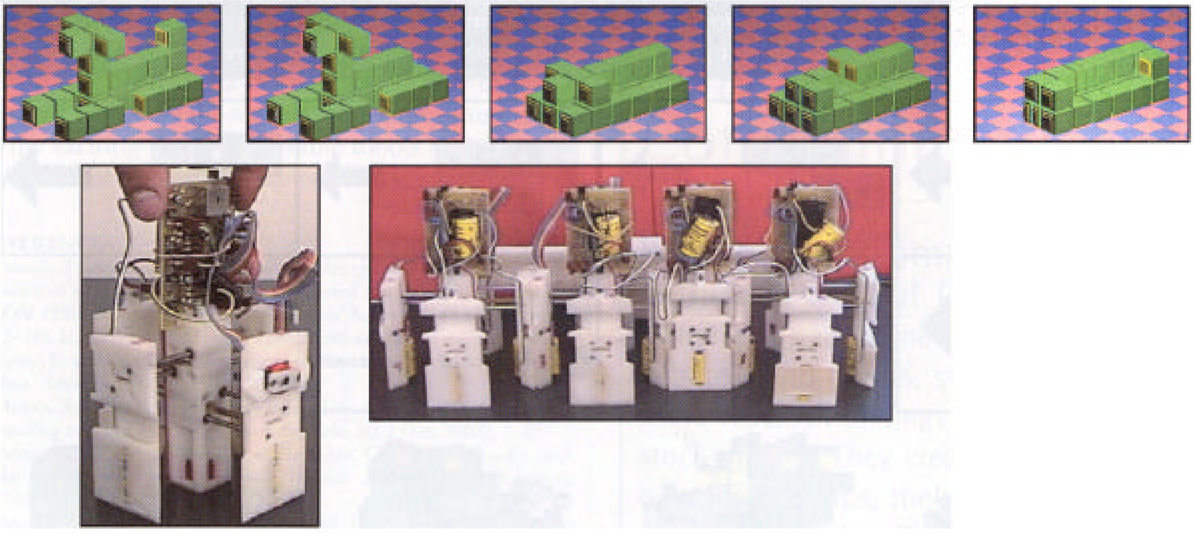

Figure 2. (Top) Five snapshots from a simulation using Crystalline robots showing the intermediate steps in the transformation from dog-shaped object to a couch-shaped object; (bottom left) the physical prototype of the Crystalline atom; and (bottom right) a robot consisting of four Crystals.

module that is to eat it, paths for several modules can coexist in the Crystal robot. This navigation and configuration strategy in turn allows modules to perform simultaneous reconfigurations without relying on a central clock or the actuation of only a single path at a time.

The reconfiguration process involves two main steps. First, a path is planned for each module in a distributed fashion. The result is a set of pellets distributed through the atoms of the robot. Once the pellets are in place, the actuation happens asynchronously, as each atom looks for pellets and "eats" them without adhering to a strict schedule. This strategy means that although the intermediate structure of the crystal is undetermined, the final structure is as specified.

This actuation protocol can direct the active module to move and trade identities with other modules along the path, eventually resulting in a module appearing at a location specified in the goal statement. It provably allows for many paths to be planned and executed simultaneously through the robot, since each active module needs to look only in its immediate neighborhood to discover and actuate the next step on its path.

General approaches to decentralized planning. The current direction in self-reconfiguring robotics focuses on designing and building hardware and developing algorithms coupled to specific hardware. We are at a point where we can step back to examine more general questions about self-reconfiguration planning in an architecture-independent way. It is important to examine architecture-independent algorithms that can be instantiated for many different systems because they have the potential of providing a more general science base for the field. By outlining general principles for reconfiguration planning, we hope to learn how to better design hardware and control algorithms.

The ability of self-reconfiguring systems to change shape can lead to water-flow-like locomotion algorithms allowing the robot to conform to the terrain on which it has to travel. These algorithms have the potential for working well in unstructured environments. In most existing self-reconfiguring robot systems [2, 3, 6, 9], an individual module can move in general ways relative to a structure of modules by traveling on the surface of the structure. Specifically, an individual module is capable of:

- Linear motion on a plane of modules;
- Convex transitions into a different plane; and
- Concave transitions into a different plane.

The details for how to accomplish these goals are architecture-dependent. We can use these motion abstractions as the basis for a general decentralized cellular-automata algorithm that can work for any system capable of such motions. The cellular automata control uses the primitives within local rules to model water-flow-like locomotion.

We have developed several sets of simple rules for cellular automata that produce reliable and provable locomotion; Figure 3 shows snapshots from a simulation of a blob-like robot moving over irregular obstacles using this approach. Each rule requires a set of

Figure 3. (Top) Five rules for rightward locomotion without obstacles. The robot modules are represented as squares; a variety of symbols characterize the local configuration of the active cell, which is represented as a black square with a dot in the middle. The first rule allows, for example, the active module to move up by one unit when its top is free and the surrounding structure consists of at least three units arranged in an L shape. (Bottom) Screenshots of a compliant locomotion simulation using cellular automata rules; the light gray squares model obstacles, or terrain irregularities.

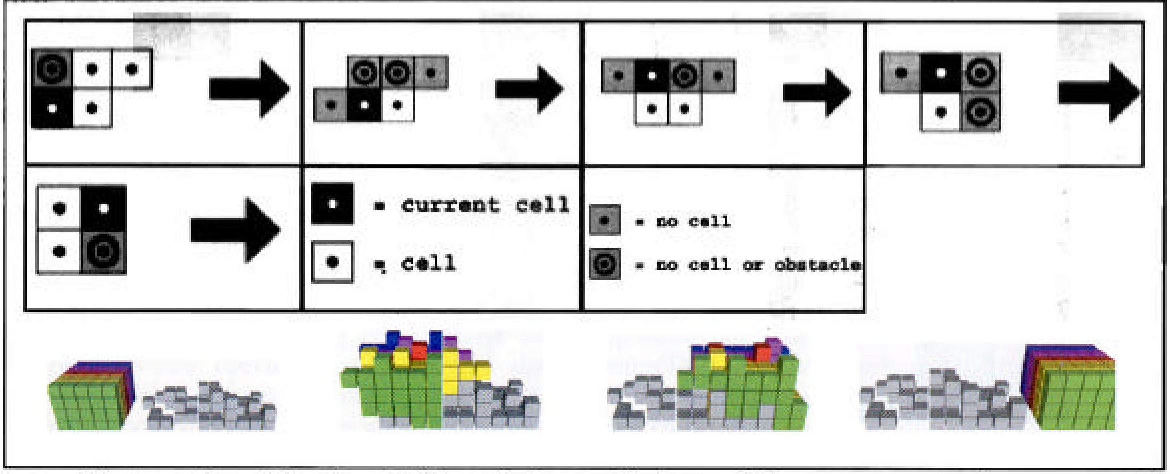

preconditions on the neighborhood of the cell that, when activated, causes the cell to move to an adjacent location on the surface of the system. These correct abstract algorithms can be instantiated on a variety of robot systems with their correctness properties intact.[1]

In addition to the locomotion task, we have investigated other reconfiguration tasks, including: self-replicating robots (dividing a large modular system into a collection of smaller systems, an operation useful for distributed search and rescue applications); and self-merging robots (regrouping two smaller robots into one robot). These tasks are interesting by themselves and useful for applications involving distributed monitoring and surveillance. Moreover, insights into how to approach these problems will improve our understanding of more general problems, such as how the number of modules in a robot affects the robot's task repertoire and whether a large number of small simple modules is more powerful than a single complex robot.

Conclusion

Self-reconfiguring robots are able to adapt to the operating environment and required functionality by changing shape. They consist of a set of identical robotic modules that can autonomously and dynamically change their aggregate geometric structure to suit different locomotion, manipulation, and sensing tasks. However, creating robots with self-reconfiguration capabilities is a serious challenge now being met through new designs for reconfigurable systems and new ideas about algorithmic planning and control that confer autonomous reconfigurability. We've discussed hardware design issues and presented two solutions developed in our laboratory. We also discussed planning issues and illustrated a hardware-specific distributed planner and a generic distributed planner that can be instantiated to many different designs.

These results are encouraging first steps toward creating self-reconfiguring robotics applications. However, we have a way to go before we can engineer modular self-reconfiguring robot systems that can be embedded into the physical world and respond in real time to requests for self-assembly. Because these robot systems will constitute long-lived distributed systems, all the supporting hardware and software will have to be robust, long-lasting, fault-tolerant, scalable, and self-healing. The hardware will have to rely on simple and robust actuation. The units will have to be powered for long periods of time. Adding and removing units into the system will have to be incremental, in that these changes should affect the overall system only locally. When units break, the system should be able to repair itself without altering overall global functionality. The units will have to be networked with a reliable wireless ad-hoc communication infrastructure. And control will have to be highly parallel, scalable, and distributed.

---

[1] We instantiated them to the Molecule robot [2], the MEL Fracta robot, and the robot in [3]. The proof is developed from three statements: At any point in time, at least one of the five rules in Figure 3 applies to one cell in the system; any sequence of rules produces a net rightward motion; and no sequence of rules causes the system to disconnect its components from one another. The correctness result for decentralized locomotion is encouraging, as it is generally difficult to prove the correctness of distributed algorithms specified in a bottom-up fashion.

To develop such systems, we have to improve ou understanding of the general properties that confe modular robots with self-reconfiguration capabilities as well as more generic (rather than architecture-specific) solutions to control and planning. Programming and giving commands to these systems should be a least as easy as writing a HTML page. These issues are being addressed by the research community, motivated by the exciting vision of versatile robots achieving the same level of flexibility as biological systems of cells. ■

References

1. Fukuda, T. and Kawauchi, Y. Cellular robotic system (CEBOT) as one realization of a self-organizing intelligent universal manipulator. In *Proceedings of the 1990 IEEE Conference on Robotics and Automation* (Cincinnati, May 13–18). IEEE Computer Society Press, 1990, 662–667.
2. Kotay, K. and Rus, D. Locomotion versatility through self-reconfiguration. *Robo. Autonom. Syst. 26* (1999), 217–232.
3. Murata, S., Yoshida, E., Kurokawa, H., Tomita, K., and Kokaji, S. Self-repairing mechanical systems. *Autonom. Robo. 10,* 1 (Jan. 2001), 7–21.
4. Pamecha, A., Ebert-Uphoff, I., and Chirikjian, G. Useful metrics for modular robot motion planning. *IEEE Transact. Robo. Autom. 13,* 4 (Aug. 1997), 531–545.
5. Rus, D. and Chirikjian, G., Eds. Special Issue on Self-Reconfiguring Robots. *Autonom. Robo. 10,* 1 (Jan. 2001), 1–124.
6. Rus, D. and Vona, M. Crystalline robots: Self-reconfiguration with uni-compressible modules. *Autonom. Robo. 10,* 1 (Jan. 2001), 107–124.
7. Shen, W.-M., Will, P., and Castano, A. Robot modularity for self-reconfiguration. In *Proceedings of the International Society for Optical Enginee (SPIE) Conference on Sensor Fusion and Decentralized Control in Robotic Systems* (Boston 1999), 236–245.
8. Unsal, C. and Khosla, P. A multi-layered planner for self-reconfiguration o a uniform group of I-cube modules. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (Maui, HI). IEEE Press Piscataway, NJ, 2001, 598–606.
9. Yim, M., Zhang, Y., Lamping, J., and Mao, E. Distributed control for 3D shape metamorphosis. *Autonom. Robo. 10,* 1 (Jan. 2001), 41–56.

Daniela Rus (rus@cs.dartmouth.edu) is an associate professor in th Computer Science Department of Dartmouth College, Hanover, NH.
Zack Butler (zackb@cs.dartmouth.edu) is a research associate in the Computer Science Department of Dartmouth College, Hanover, NH.
Keith Kotay (kotay@cs.dartmouth.edu) is a Ph.D. candidate in th Computer Science Department of Dartmouth College, Hanover, NH.
Marsette Vona (vona@helios.jpl.nasa.gov) is a member of the technical staff at NASA's Jet Propulsion Laboratory.

This work was done in the Dartmouth Robotics Lab. Support was provided through a U. National Science Foundation CAREER award IRI-9624286 and NSF awards IR 9714332, EIA-9901589, IIS-9818299, and IIS-9912193, the Sloan foundation, and NASA Spacegrant award.

Manuela M. Veloso

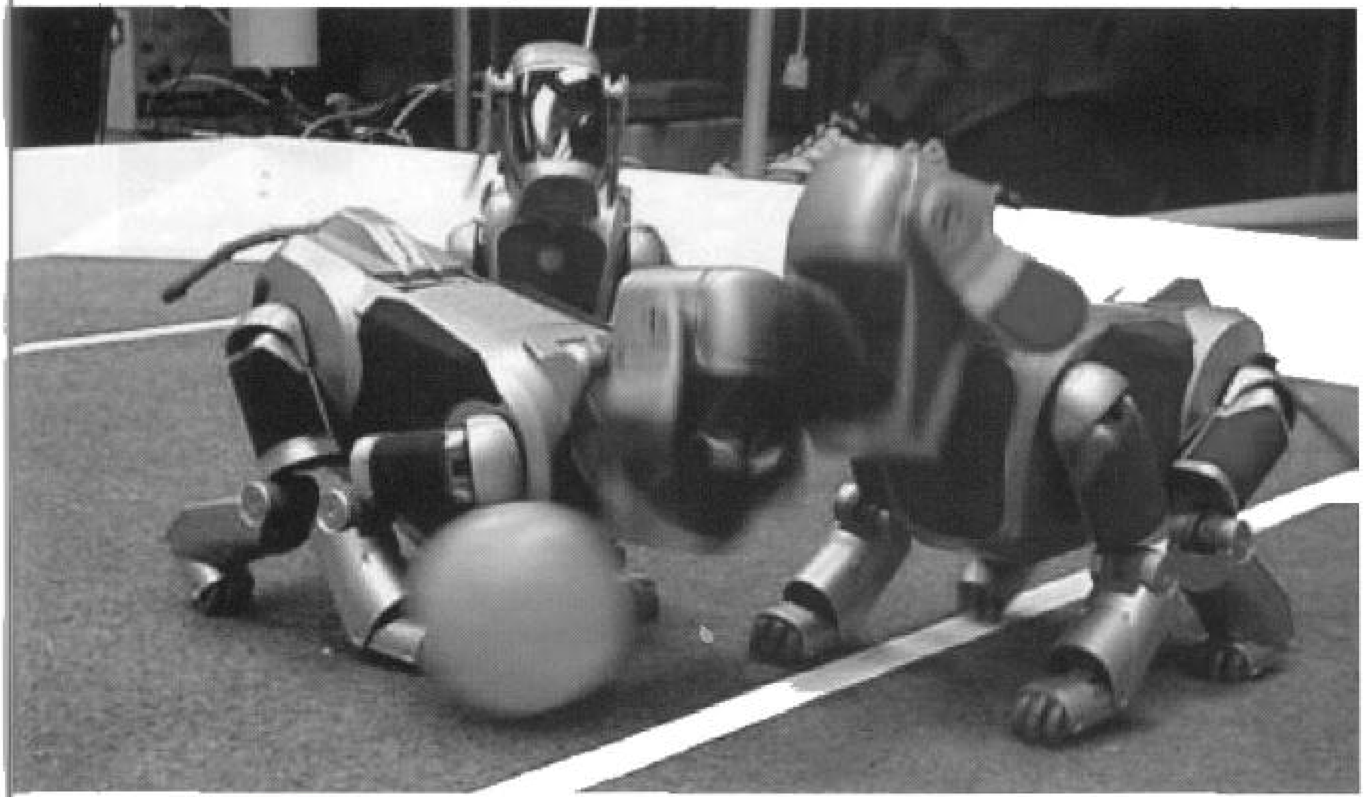

ENTERTAINMENT ROBOTICS

Competing teams of autonomous robot soccer players illustrate the challenges, pleasures, and promise of developing collaborative multi-robot applications.

Artificial intelligence (AI) focuses on using computers to manipulate symbolic and numerical information to perform a wide variety of intelligent reasoning similar to humans. Robotics includes investigating the feasibility of creating mechanical creatures—robots—to perform like humans in real-world physical environments. The fields of robotics and AI converge in pursuing this goal. The automated preprogrammed action of robotic artifacts has been developed extensively and successfully for industrial, inaccessible, and hazardous environments, including volcanos, the arctic, outer space, and the deep ocean floor. But

Sony Aibo robots (2001 research version) competing in the RoboCup-2001 soccer tournament.

DEBRA TOBIN, CARNEGIE MELLON UNIVERSITY

the accelerating advances in computer power now appear to add enormous credibility to the notion that robots with humanlike AI can be developed fully. This expectation offers many new opportunities for people to interact and coexist with robots.

Recent efforts have sought to bring robots into our daily lives, including in the form of autonomous vehicles, museum tour guides, helpmates for the elderly, and robot competitions, particularly robot soccer. My remarks here about entertainment robotics are based on my involvement with my students in robot soccer research over the past six years.

> We learned that robots playing a game cannot afford to be stopped and not act; the game goes on at a speed that requires rather proactive behaviors.

The late Herbert A. Simon, a professor of computer science and psychology at Carnegie Mellon University and a founder of the AI field, concluded his lecture "Forecasting the Future or Shaping It?" at the October 2000 Earthware Symposium (see video at www.ul.cs.cmu.edu/) by saying: "Here around CMU, we have been amazed, amused, gratified, and instructed by the developments in robot soccer. For four years, and with rapidly increasing skill, computers have been playing a human game requiring skillful coordination of all the senses and motor capabilities of each player, as well as communication and coordination between players on each team, and strategic responses to the moves of the opposing team. We have seen in the soccer games an entire social drama played out with far less skill (thus far) than professional human soccer, but with all the important components of the latter clearly visible.

"Here we see, in a single example, a complex web of all the elements of intelligence and learning—interaction with the environment and social interaction, use of language—that AI has been exploring for half a century and a harbinger of its promise for continuing rapid development. Almost all of our hopes and concerns for the future can be examined in miniature in this setting, including our own role in relation to computers." The lecture went on to forecast our interactions with computers and robots. But his impressions and assessment of robot soccer are the best introduction to entertainment robotics I know.

Robot soccer pioneered multi-robot research and entertainment. Until its development began in 1996, most robotics research focused on single-robot issues. Robot soccer presented a new horizon; teams of autonomous robots have to respond to a highly dynamic environment, including other teams of robots, to accomplish specific goals, like get the ball in the opponents goal. Moreover, choosing a popular game like soccer to explore such a rich research objective has apparently made a big difference. The entertainment component of robot soccer has been significant in attracting researchers, as well as crowds of spectators. For example, the fifth annual RoboCup International Competition held for the first time in the U.S. last August in Seattle included more than 500 participants, 200 robots, and thousands of spectators (see Figure 1). The ambitious official RoboCup motto is: "By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champions." I now briefly illustrate the technical challenges faced by teams of autonomous robots dealing with such real-time dynamic tasks as robot soccer.

Autonomous Robots

To robot researchers, an autonomous robot is capable of handling problems without the help of an outside source, particularly a human. We humans are in general autonomous in our everyday lives, capable of surviving in our relatively unstructured environments. Autonomy includes three main capabilities:

Perception. The ability to recognize the surrounding environment, including the five human senses: vision, hearing, taste, smell, and touch.

Action. The ability to respond to perceived sensations, enabling one to change one's own state or the state of the environment; many actions are available to autonomous creatures, possibly in an infinite number in some continuous space; common actions include all sorts of motion and manipulation.

Cognition. The ability to reason, including selecting from among the actions that are possible in response to sensations; reasoning is a complex process that can include the ability to experiment and learn from feedback from the effects of the actions selected.

Research in robotics has a very long way to go to actually achieve the level of perception, action, and cognition we humans demonstrate in our everyday lives. But the research is advancing in that direction. Inherent in this advancement is the fact that robots will be part of our lives and in particular will be able

Figure 1. Researchers and robots participating in RoboCup-2001, Seattle, August 2001.

to coexist with us in entertainment tasks. Indeed, the fact that scientific and technology advances are contributing to the development of autonomous robots with perception, action, and cognition similar to our own motivates us to use our discoveries well and learn to coexist with robots.

Teams of Robots

In robot soccer, the robots face a highly dynamic and uncertain environment in which they have to achieve clear goals like advancing the ball toward the opponent's goal. Robot soccer teams need to effectively integrate perception, action, and cognition in real time. Each team of robots needs to continuously live in a cycle, perceiving the world, deciding what to do, and performing actions. One of the main challenges in developing such integration is how to provide the robots with the ability to close this autonomy cycle, so they perceive the environment, make decisions about which actions to take, actually take actions in the world, and continuously perceive the environment, making decisions and acting.

My students and I have been developing several different teams of autonomous robot soccer players, each reflecting a variety of challenges of perception, action, and cognition. I briefly describe two of them in the following paragraphs to illustrate the concrete challenges of multi-robot entertainment.

Multi-robot soccer teams. Robot perception is one of the main bottlenecks. Robots need to be equipped with sensors from which they can accurately and reliably infer the state of the world. Figure 2 shows small-wheeled soccer-playing robots [3, 4]. Each team designs and builds its own robots under specific size constraints. The robots play with an orange golf ball on a field approximately the size of a ping-pong table. Each robot team is allowed to hang a vision camera over the playing field to provide a global view of the field of play. Processing images globally in real time is

Figure 2. Carnegie Mellon soccer robots (designed by Brett Browning).

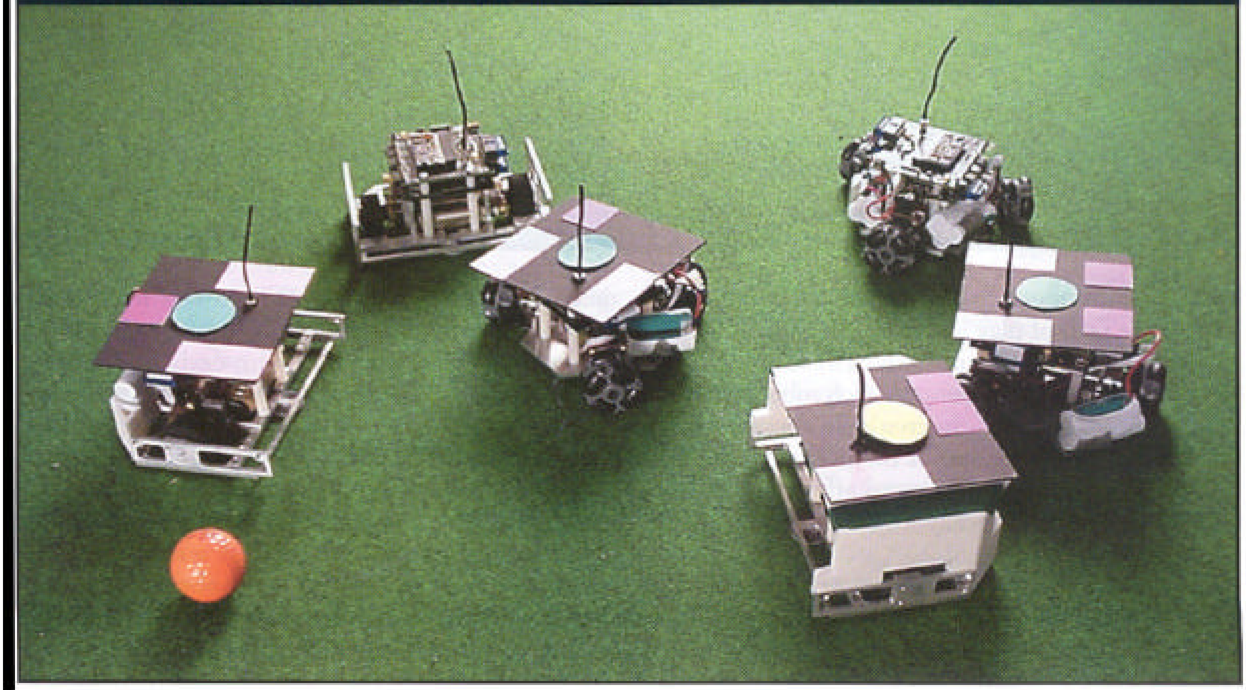

Figure 3. Sony-built Aibos are programmed by Veloso's students to play soccer in teams of three fully autonomous robots.

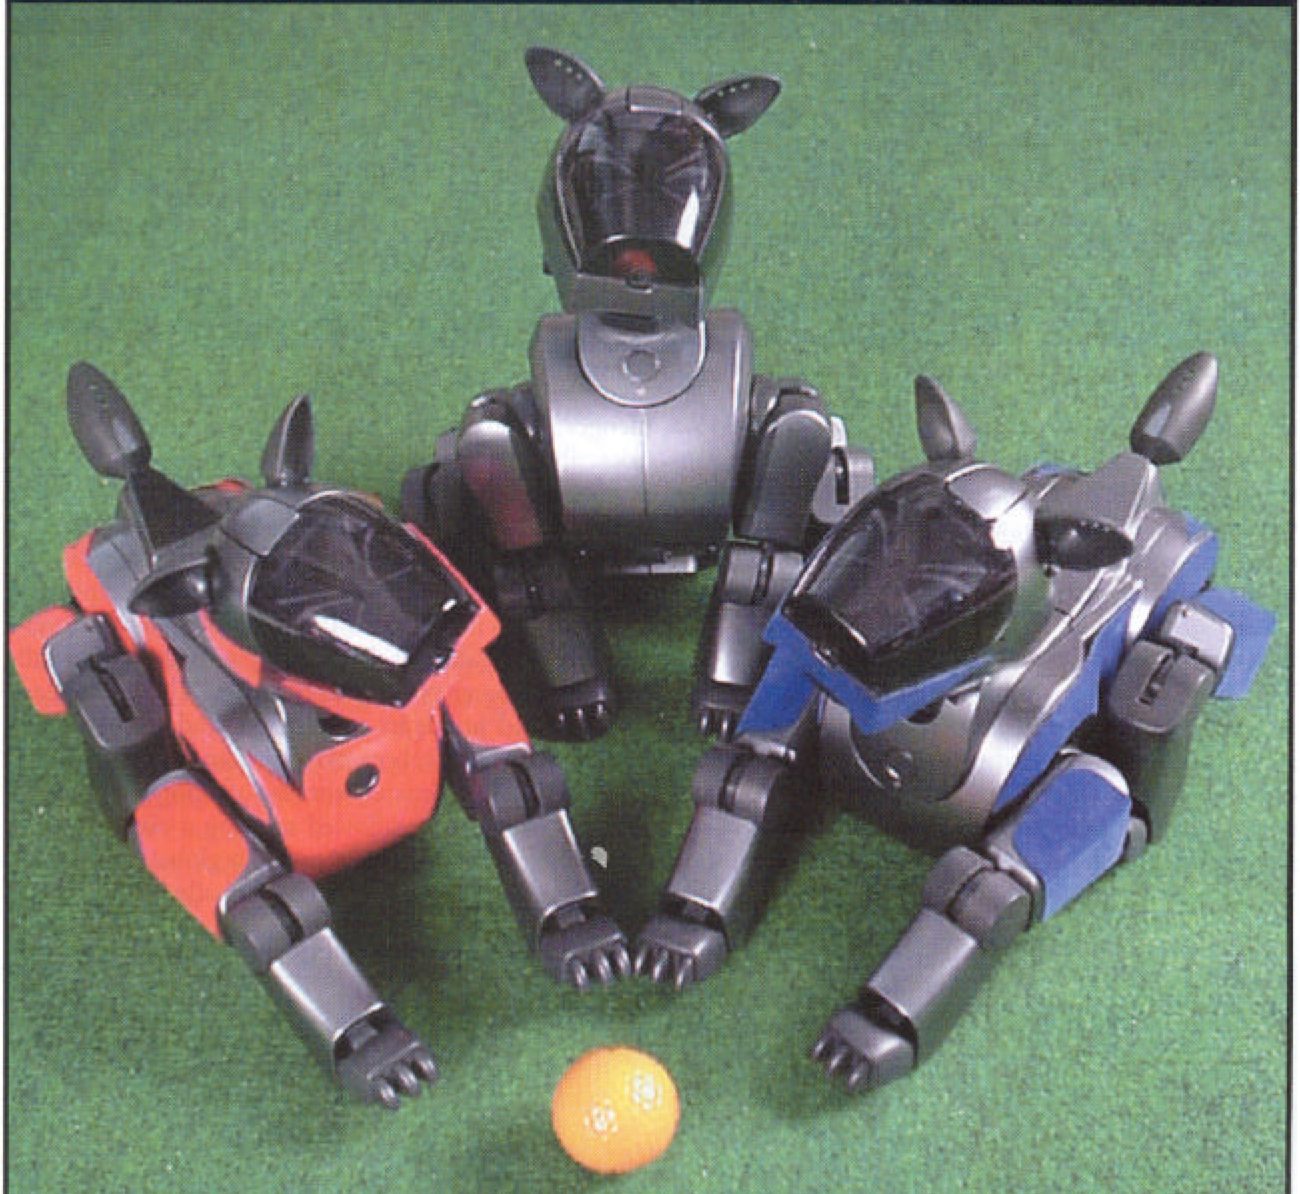

itself a significant perception challenge, but the global view of the field of play is provided to each robot on each team. The image can also be sent to an offboard computer that remotely controls each robot's motion, usually through radio. Interestingly, because each robot has a complete view of the positions of all its teammates and opponents, it can effectively use this information to strategically collaborate with other team members.

This scenario may seem far removed from human sport, as humans cannot see in all directions. But humans can share information through communication and collaboratively develop complete information of the relevant world. Within this framework, these artificial robot artifacts cannot control a ball as humans can and indeed cannot devise compelling strategic teamwork when compared with humans.

We have also developed teams of fully autonomous legged robots with onboard vision and computational power. Figure 3 shows the legged robots we use—programmable versions of the Aibo designed and built by Sony Corp. We have used them as a hardware platform since Sony introduced its first version in 1998 [5, 6]. For each one's onboard processor, we have developed algorithms to provide image processing, localization, and control. None of them is remotely

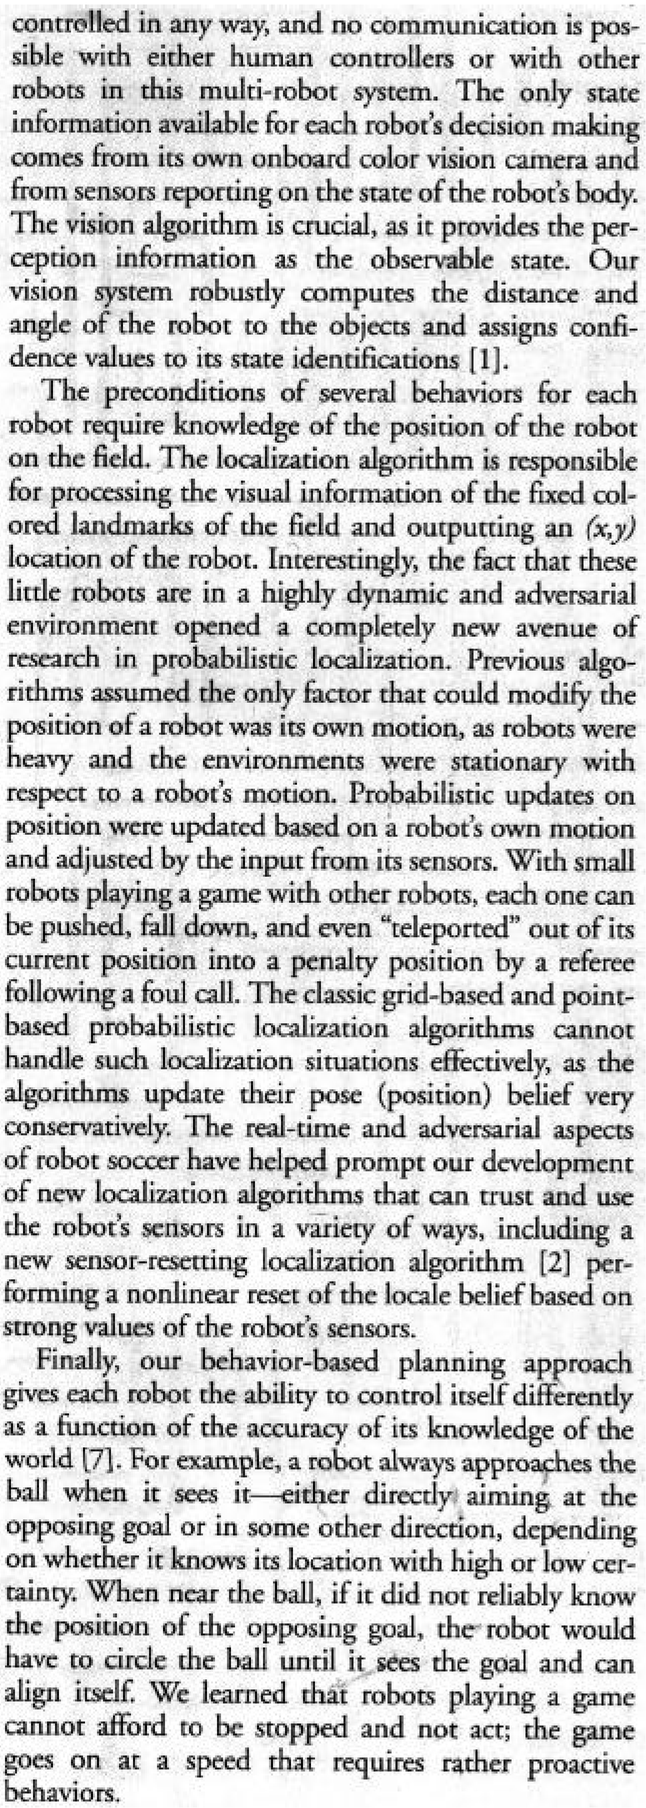

controlled in any way, and no communication is possible with either human controllers or with other robots in this multi-robot system. The only state information available for each robot's decision making comes from its own onboard color vision camera and from sensors reporting on the state of the robot's body. The vision algorithm is crucial, as it provides the perception information as the observable state. Our vision system robustly computes the distance and angle of the robot to the objects and assigns confidence values to its state identifications [1].

The preconditions of several behaviors for each robot require knowledge of the position of the robot on the field. The localization algorithm is responsible for processing the visual information of the fixed colored landmarks of the field and outputting an *(x,y)* location of the robot. Interestingly, the fact that these little robots are in a highly dynamic and adversarial environment opened a completely new avenue of research in probabilistic localization. Previous algorithms assumed the only factor that could modify the position of a robot was its own motion, as robots were heavy and the environments were stationary with respect to a robot's motion. Probabilistic updates on position were updated based on a robot's own motion and adjusted by the input from its sensors. With small robots playing a game with other robots, each one can be pushed, fall down, and even "teleported" out of its current position into a penalty position by a referee following a foul call. The classic grid-based and point-based probabilistic localization algorithms cannot handle such localization situations effectively, as the algorithms update their pose (position) belief very conservatively. The real-time and adversarial aspects of robot soccer have helped prompt our development of new localization algorithms that can trust and use the robot's sensors in a variety of ways, including a new sensor-resetting localization algorithm [2] performing a nonlinear reset of the locale belief based on strong values of the robot's sensors.

Finally, our behavior-based planning approach gives each robot the ability to control itself differently as a function of the accuracy of its knowledge of the world [7]. For example, a robot always approaches the ball when it sees it—either directly aiming at the opposing goal or in some other direction, depending on whether it knows its location with high or low certainty. When near the ball, if it did not reliably know the position of the opposing goal, the robot would have to circle the ball until it sees the goal and can align itself. We learned that robots playing a game cannot afford to be stopped and not act; the game goes on at a speed that requires rather proactive behaviors.

All the teams in the RoboCup legged-robot league use the same Sony hardware platform, creating a very interesting research AI problem, as all the robots have in principle the same low-level perception and motion capabilities. Therefore, their eventually different performance should mainly reflex their cognition. However, this is indeed not the case. Although they do differ as to cognition, it remains a challenge to program them to use their similar hardware. The result is that some robots move faster or see better than other robots. Robotics researchers focus on different research directions, leading to robots that vary by performance, even though they have the same physical components. This variation is similar to how we all handle the limits of our physical and cognitive abilities in different ways, achieving different results for similar tasks.

Conclusion

Robot soccer illustrates the challenges of building complete autonomous robots able to perform active perception and sensor-based planning while playing a multi-robot game. The games are not only a source of entertainment but a great source of advances in robotics research. I am confident in extrapolating that further advances in entertainment robotics will continue to serve this twofold goal.

ì