Human-Inspired Robots

Silvia Coradeschi, *Orebro University*

Robots are just now becoming part of our everyday life and being used by ordinary people. Future robots will work in hospitals, elder care centers, schools, and homes. Similarity with humans can facilitate interaction with a variety of users who don't have robotics expertise, so it makes sense to take inspiration from humans when developing robots. However, humanlike appearance can also be deceiving, convincing users that robots can understand and do much more than they actually can. Developing a humanlike appearance must go hand in hand with increasing robots' cognitive, social, and perceptive capabilities.

In this installment of Trends & Controversies, Hiroshi Ishiguro and Minoru Asada present the ultimate imitation of a human's appearance: an android. Androids are robots that not only look like humans but also move like us and have artificial skin that feels like ours.

Humanlike robots must communicate clearly using natural language. Stuart Shapiro outlines natural language's importance both as a practical way to communicate and a means of encouraging social interaction.

Cognitive abilities are important for robots interacting with humans and acting in complex environments. Michael Thielscher presents the field of cognitive robotics, which investigates how robots can achieve greater autonomy and flexibility in such environments.

Social competence is also essential in natural human-robot communication. Cynthia Breazeal has been a pioneer in developing socially interactive robots, and she presents the challenges and progress in this field.

The growing field of socially assistive robotics is becoming an important testbed for human-inspired robots. Maja J. Matarić outlines the challenges of aiding people by means of social interaction rather than through physical contact alone.

Finally, Hiroshi Ishida presents olfaction as a new sense for robots. The ability to recognize smells will bring robots closer to humans. This sensor offers great new possibilities but also new challenges for robotics.

—Silvia Coradeschi

Humanoid and Android Science

Hiroshi Ishiguro, *Osaka University*
Minoru Asada, *JST ERATO Asada Synergistic Intelligence Project and Osaka University*

Why are we attracted to humanoids and androids? The answer is simple: we tend to anthropomorphize nonhuman things. Humans always anthropomorphize targets of communication and interaction, so we expect much from humanoids. In other words, we find a human in the humanoid. Recently, robotics researchers have begun shifting from traditional studies on navigation and manipulation to studying interaction with robots.

The study of human-robot interaction has been neglecting an issue—appearance and behavior. The interactive robots that have been developed thus far are nonandroid types, so the researchers who developed them didn't focus on the robots' appearances. Evidently, a robot's appearance influences subjects' impressions, and it's an important factor in evaluating the interaction. Although many technical reports compare robots with different behaviors, they haven't focused on the robots' appearances. There are many empirical discussions on simplified robots, such as dolls. However, designing a robot's appearance, particularly to make it appear humanoid, has always been the industrial designers' role rather than that of researchers and engineers. This is a serious problem for developing and evaluating interactive robots. Appearance and behavior are tightly coupled.

Bridging science and engineering

One way to tackle the issue is to use a humanlike robot—an android—to study human-robot interaction. Figure 1a is a humanoid developed by Mitsubishi Heavy Industry, and figure 1b shows an android that Hiroshi Ishiguro and his colleagues developed in cooperation with KOKORO. The android has 42 air actuators for the upper torso, excluding fingers. During development, we determined the actuators' positions by analyzing a human's movements using a precise 3D motion-capture system. The actuators can represent unconscious movements such as chest movements due to breathing, in addition to large, conscious movements of the head and arms. Furthermore, the android can generate facial expressions that are important for interaction with humans. When we publicized the android through the media, we were anxious about ordinary Japanese people's reactions. However, it wasn't uncanny for them; they just praised the quality and technology.

Developing androids requires contributions from both robotics and cognitive science. To realize a more human-

1541-1672/06/$20.00 © 2006 IEEE
Published by the IEEE Computer Society

like android, knowledge from human science (that is, from studying humans) is necessary. This new interdisciplinary framework between engineering and cognitive science is called *android science*.[1]

In the past, robotics research used knowledge from cognitive science, and cognitive science research utilized robots to verify hypotheses for understanding humans. However, robotics' contribution to cognitive science has been inadequate. Appearance and behavior couldn't be handled separately, and nonandroid type robots weren't sufficient as cognitive science tools. We expect that using androids with a humanlike appearance can solve this problem. Robotics research based on cues from cognitive science faces a similar problem, because it's difficult to recognize whether the cues pertain solely to a robot's behavior, isolated from its appearance, or to the combination of its appearance and behavior. In the framework of android science, androids enable us to directly share knowledge between the development of androids in engineering and the understanding of humans in cognitive science.

So, android science has several major research issues. The issues in robotics are to develop a humanlike appearance using silicon, humanlike movement, and humanlike perception by integrating ubiquitous sensor systems. In cognitive science, the issue is conscious and subconscious recognition. When we observe objects, various modules are activated in our brain. Each of these matches the input sensory data with our brains' models of human faces, voices, and so on, which affect our reactions. So, even if we recognize that a robot is an android, we react to it as if it were human. Android science's goal is to find the essential factors of humanlikeness and realize a humanlike robot.

How can we define humanlikeness? Furthermore, how do we perceive it? That humans have both conscious and unconscious recognition is fundamental for both the engineering and scientific approaches. It will be an evaluation criterion in androids' development, and it provides us cues for understanding the human brain mechanism of recognition.

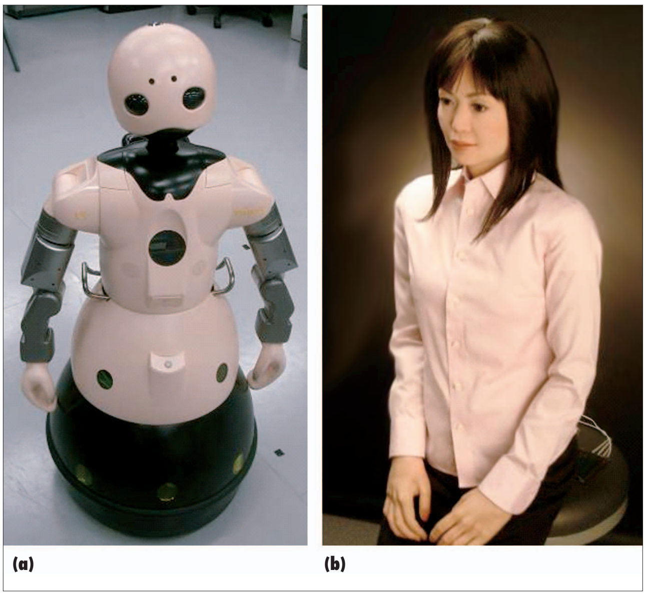

Figure 1. Humanoid and android robots. (a) Humanoid Eveliee P1 is based on WAKAMURU, developed by Mitsubishi Heavy Industry. (b) Android Repliee Q2 was developed in cooperation with KOKORO (www.kokoro-dreams.co.jp).

From android science to humanoid science

The history of intelligent robotics started with Shakey, a robot developed at SRI in 1965. Shakey provided robotics researchers with several important research questions. We focused on the fundamental issues for making Shakey more intelligent, such as AI, computer vision, and language recognition. Since Shakey's development, we've been using these technologies to create new humanoids and androids, which provide us with important new research questions as Shakey did. These questions are in interdisciplinary areas among robotics, cognitive science, neuroscience, and social science.

Android science is also an interdisciplinary framework, but it's rather limited. In addition to humanlike appearance and movement, we must consider the internal mechanisms, such as humanlike dynamic and adaptive mechanisms and complicated sensorimotor mechanisms, for more tightly coupling engineering and science. As we mentioned earlier, what we wish to know is what a human is. We can understand this by developing humanoids comprising humanlike hardware and software. We call this extended framework *humanoid science*.

Our project, JST ERATO (Japan Science and Technology Agency, Exploratory Research for Advanced Technology) Asada Synergistic Intelligence (www.jeap.org), is based on the humanoid-science framework. *Synergistic intelligence* means intelligent behaviors that emerge through interaction with the environment, including humans. Synergistic effects are expected in brain science, neuroscience, cognitive science, and developmental psychology. Synergistic intelligence provides a new way of understanding ourselves and a new design theory of humanoids through mutual feedback between the design of humanlike robots and human-related science.

Synergistic intelligence adopts *cognitive developmental robotics*,[2] a methodology that comprises the design of self-developing structures inside the robot's brain and environmental design (how to set up the environment so that the robots embedded therein can gradually adapt to more complex tasks in more dynamic situations). Here, one of the most formidable issues is nature versus nurture. To what extent should we embed the self-developing structure, and to what

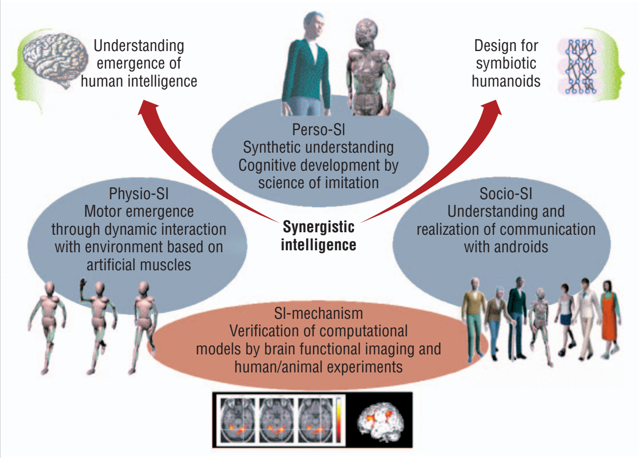

Figure 2. The JST ERATO Asada Synergistic Intelligence project is based on a humanoid-science framework.

extent should we expect the environment to trigger development? We're approaching this issue using the kinds of topics in figure 2.

References

1. H. Ishiguro, "Interactive Humanoids and Androids as Ideal Interfaces for Humans," *Proc. 11th Int'l Conf. Intelligent User Interfaces*, ACM Press, 2006, pp. 2–9; www.ed.ams.eng.osaka-u.ac.jp/research/Android/android_eng.htm.

2. M. Asada et al., "Cognitive Developmental Robotics as a New Paradigm for the Design of Humanoid Robots," *Robotics and Autonomous Systems*, vol. 37, 2001, pp. 185–193.

Natural-Language-Competent Robots

Stuart C. Shapiro, *State University of New York at Buffalo*

Why use natural language to communicate with computer systems? I've always thought the answer was obvious. Because NL is the most natural method of communication for people,[1] computer systems would be easiest to use if they understood and spoke NL. I often get resistance to this opinion, but the resistance almost always takes the form of comparing current NL interfaces to current GUIs. Current NL interfaces don't work as well as current GUIs because computers don't fully understand NL. But that just makes NL understanding an important research topic.

Indeed, not only is NLU in general an active research topic, so is the specific topic of NL interfaces to robots, as evidenced by the recent ACM Conference on Human-Robot Interaction (HRI 2006, www.hri2006.org). Moreover, not only are people who describe themselves as NLU researchers working to make NL understandable to computers, the field of programming languages can be seen as a bottom-up attack on the same problem. Many programming language advances, as far back as the development of assembly languages, have been presented as letting programmers express problems in a high-level language that's natural for humans.

Definitions and motivations

By NL, we could be referring to text or speech. Many NL researchers are working on text-based systems because they're interested in processing information from the Web or other documents. Others are using text as an expedient until speech recognition becomes easier to use. (I'm in this group.) So, by an *NL-competent robot*, I generally mean one that understands and generates speech.

What counts as a robot? By robot, I mean an embodied computer system that has sensor and effector organs. I'm not limiting the term to hardware-implemented robots that operate in the real world. I also include robots with simulated bodies and sensor and effector organs that operate in simulated or virtual reality worlds. I also include teleoperated as well as autonomous robots. A teleoperated robot is controlled by a console, joystick, or other device. I don't, however, want to argue that users of prosthetic limbs or master-slave manipulators, such as those that manipulate objects in toxic environments or perform microsurgery, would necessarily benefit from NL communication.

R&D of NL-competent robots has at least two motivations. One is science. Because NL is the natural communication medium for people, researchers attempting to achieve a computational understanding of human intelligent behavior and create devices that exhibit such behavior must include NL-competent robots in their research agendas (although maybe not in their personal agendas). The other motivation is pragmatic—the belief that robots will be more useful if they are NL-competent. Because I take the scientific motivation to be inarguable, I'll concentrate on the pragmatics.

Speech as a command language

It might seem that controllers of teleoperated robots wouldn't need to use NL, given their more direct control devices. Experiments have shown, however, that using single-word speech commands is beneficial.[2] This is because of both the naturalness of speech and the ease of using it when the operator's hands and eyes are occupied with the other controls. However, Manuel Ferre and his colleagues didn't find a more complete speech-based sentence-understanding system to be beneficial because it required the operator to pause between words, causing a considerable loss in fluency.[2] They found that operators don't work comfortably if the system's recognition rate is under 90 percent. This is another indication that when NL interfaces don't understand fluent speech, they don't compete well with other interfaces.

If speech is useful for robot controllers whose hands and eyes are occupied, it's even more useful for people who can't use their hands or eyes because of long-term disability or short-term injury. The development of speech interfaces for assistive devices is therefore an active research area, and several researchers presented papers on the topic at HRI 2006.

NL as a programming language

Although using speech as a command language is useful, using NL as a robot programming language is even more useful. Any robot will have a repertoire of actions and recognizable domain features designed or programmed into it. A robust understanding system that lets a human controller express these actions and features in NL would provide a useful speech command language. But consider situations in which a human wants to ask the robot to perform a sequence of actions, to do one of several actions depending on what it perceives in its world, or to repeat some action until it has accomplished some effect. As we know, providing NL constructs to express these control structures will make NL a full-fledged programming language. Some of these more complex procedures might be ad hoc procedures the controller wants done only once, but others might be more common. If the human can conceptualize these possibly parameterized procedures and express them in a way acceptable to the NL understanding system, and if the NLU system provides a way for the human to say something like "The way to do X is to do Y," then the NLU robot command language will have all the constructs of a procedural programming language. Stanislao Lauria and his colleagues achieved this to some extent in the context of a hardware robot that can take directions for navigating in a small model city.[3]

To use NL to instruct and command robots, the NLU system must translate NL inputs into some kind of robot acting program. That is, the meaning-representation language that the NLU system uses must be an autonomous-agent acting language. Unfortunately, researchers designing agent acting languages and those designing representation languages for the (contextually grounded) semantics of NL mostly operate independently of each other.

A promising approach to building NL-competent robots is to combine these researchers' efforts. For example, the representation the robot uses to understand a verb phrase could be a construct in the robot's acting language that it can perform. In this way, the robot performs according to an imperative sentence by executing the structure representing its understanding of the command. By appropriately appending a representation of itself and of the current time to the representation of an action it has just performed, the robot believes it has performed the action. Giving that belief to its NL generation routine produces an NL report of what it has just done. My colleagues, students, and I have been pursuing this approach.[4]

Given an NL programming language, having an NL development environment would also be important, and it would be most convenient if the robot participated in the development process. For example, Lauria and his colleagues' navigation robot uses NL generation to ask for explanations of directions it doesn't understand.[3]

Clifford Nass and Scott Brave found that people can't help but interpret speech in a social context and attribute human characteristics to the speaker, even if they know the speaker is a machine.[1] So, taking the social aspects of speech into account is important so that "users will not simply talk *at* and listen *to* computers, nor will computers simply talk *at* and listen *to* users. Instead, people and computers will cooperatively *speak with* one another"[1] [italics in original].

> If speech is useful for robot controllers whose hands and eyes are occupied, it's even more useful for people who can't use their hands or eyes because of disability or injury.

Future research directions

NL-competent robots aren't here yet, but we're making progress, and speech interfaces are becoming common at call centers and in automobiles.[5] Robotics researchers should be encouraged to include NL competence among their robots' abilities, and NL researchers should be encouraged to consider robots as platforms for their work.

References

1. C. Nass and S. Brave, *Wired for Speech: How Voice Activates and Advances the Human-Computer Relationship*, MIT Press, 2005.
2. M. Ferre et al., "Voice Command Generation for Teleoperated Robot Systems," *Proc. 7th IEEE Int'l Workshop Robot and Human Communication* (ROMAN 98), vol. 2, IEEE Press, 1998, pp. 679–685.
3. S. Lauria et al., "Training Personal Robots Using Natural Language Instruction," *IEEE Intelligent Systems*, vol. 16, no. 5, 2001, pp. 38–45.
4. S.C. Shapiro and H.O. Ismail, "Anchoring in a Grounded Layered Architecture with Integrated Reasoning," *Robotics and Autonomous Systems*, vol. 43, nos. 2–3, 2003, pp. 97–108.
5. E. Millard, "Talking Computers Just around the Corner," *Newsfactor Magazine*, 14 Mar. 2006, www.newsfactor.com/story.xhtml?story_id=03200000ST9C.

Cognitive Robotics

Michael Thielscher, *Dresden University of Technology*

Cognitive robotics aims at modeling human-level cognitive faculties in robots. Thinking in abstract terms is humans' most distinct cognitive capability. Specifically, we benefit from this ability when we mentally entertain the effects of different courses of action and then choose the one that will likely satisfy our needs and achieve our goals. This requires us to reflect on the actions we could take and to know what effects they normally have. In other words, coming up with a plan that's suitable under the circumstances results from our ability to draw the right conclusions from our knowledge of how the world around us functions. This ability lets us exhibit intelligent behavior in situations we haven't encountered before and that no preprogrammed routine deals with adequately.

In this regard, robots should greatly profit from taking inspiration from our cognitive ability to reflect on our own acts. This doesn't mean that cognitive robotics must necessarily take inspiration from how human brains implement these functions, considering that brains are networks of neurons. Instead, the focus is on the high-level cognitive functions themselves—the abilities to represent and reason about knowledge to make rational decisions, devise complex plans, react sensibly to unforeseen circumstances and problems, and adapt to new environments. All these capabilities are characteristics of human-level intelligence. Today's robots, however, can't reflect on what they can do or decide on their own whether a preprogrammed behavior is sensible under given circumstances. Most AI researchers expect that future cognitive robots will be much more autonomous and flexible than those of today.

Formalisms and applications

Cognitive-robotics research goes back to AI's beginning. One of the earliest knowledge representation formalisms, developed in the early 1960s, was designed specifically to provide systems with knowledge of actions. This was accompanied by a mechanism for solving tasks by fully automated reasoning. In this way, an autonomous system would be able to compose a suitable plan entirely on its own on the basis of predicted outcomes of actions. However, the first implementations of this knowledge representation formalism—known as the *situation calculus*—didn't scale up beyond domains with a small state space and few actions. That is, they suffered from the *frame problem*, which turned out to be one of the most interesting and challenging formal problems in logic-based AI. While simple to understand, it proved difficult to solve and over the years has led to an abundance of insights and results. The frame problem was, for instance, the driving force behind the emergence of *nonmonotonic reasoning*, now a vibrant, active AI subfield in itself.

The first satisfactory solution to the frame problem didn't appear until the early '90s, still using the classical situation calculus as the underlying knowledge representation language.[1] This marked the beginning of a thriving era for the theory of cognitive robotics. Other solutions to the frame problem emerged in quick succession using different representation techniques, such as the *event calculus*[2] or the *fluent calculus*.[3] In recent years, each of these basic solutions has been generalized in view of the additional challenges that autonomous systems face in real-world environments. Examples are hybrid environments with both discrete and continuous change or environments where actions frequently fail because of unknown or uncontrollable circumstances. These developments show that using logic-based representation formalisms can let us deal with large state spaces, thanks to a high level of abstraction, and can provide simple, computationally efficient ways to deal with incomplete knowledge and nondeterminism.

In addition to the emergence of increasingly sophisticated representation formalisms, theoretical approaches to cognitive robotics have evolved into actual programming languages and implementations that let you write high-level control programs for autonomous systems. Examples of such general languages are GOLOG (Algol in Logic), based on the situation calculus; FLUX (Fluent Calculus Executor), based on the fluent calculus; and a reasoning system for cognitive agents known as SNePS (Semantic Network Processing System). Thanks to a high level of abstraction paired with efficient reasoning mechanisms, these languages and their implementations let you write and run control programs for complex tasks in environments that give rise to large state spaces.

Cognitive robotics systems have proved to be competitive in a variety of areas. For example, planning itself has turned into an important AI subfield with a variety of applications and its own conferences and organized competitions for implemented systems. Symbolic reasoning is also an integral part of software agents that must interact and collaborate in heterogeneous multiagent systems. A third example of an application of logical reasoning about actions is *general game playing*, which is concerned with the development of systems that can understand descriptions of arbitrary games and play them well without human intervention.

Theoretical approaches to cognitive robotics have evolved into programming languages and implementations that let you write high-level control programs for autonomous systems.

More than 40 years of research have led to significant progress, and cognitive robotics is now a vibrant, active field that plays a major role at all large AI conferences. Surprisingly (and disappointingly) enough, the progress has been mostly restricted to theory and simulation; little progress has occurred with regard to endowing actual robots in real-world environments with cognitive functions. One exception is the control program of a robotic museum guide, which employs GOLOG to compose high-level tour plans according to visitors' interests.[4] Another is a high-level control module for unmanned airborne vehicles (helicopters), which observes traffic scenarios on the ground, predicts what traffic will do in the near future, and plans the helicopters' actions.[5]

However, these real-world systems use only a limited form of symbolic reasoning specifically targeted to the application at hand. Cognitive-robotics research hasn't achieved its main goal, which is to significantly enhance robots' autonomy and flexibility with the help of high-level reasoning capabilities. Why not?

Challenges and future work

In comparison to devising software agents that live in pure virtual environments, controlling actual robots that move in the physical world poses two major additional challenges: First, the symbols that form the basis for the cognitive processes must be grounded in the real world. Second, an autonomous physical system's actions are far less reliable than those of a program in a simulated world.

The *symbol-grounding problem* arises because high-level reasoning is based on abstract terms and concepts, but a robot processes nothing but raw sensory data. How to map these data onto a single symbol that expresses as complex an entity as, say, "Sandra's coffee mug" is a largely unsolved problem. Lacking a sufficiently general solution, the relationships between symbols and sensory data must be predefined in every detail, causing a significant loss in the flexibility and autonomy of high-level reasoning.

With regard to the problem of unreliable actuators and sensors in real-world environments, established solutions are all based on probabilistic approaches. Specifically, robust algorithms for the basic tasks of self-localization and navigation for autonomous robots require a robot to maintain probability distributions as its internal world model, which is updated according to probabilistic knowledge of its actions' effects. Contemporary theories of cognitive robotics, however, still lack a tight integration of logical and probabilistic reasoning.

Recently, researchers have looked into combining probability theory with knowledge representation languages such as the situation or fluent calculus. But so far, no competitive implementation combines the exactness of probabilistic calculations with the power of logical reasoning when it comes to coping with large state spaces containing diverse information.

Significant advances toward cognitive robotics' actual goal will therefore require a new generation of theories and implementations, which will emerge from sufficiently general solutions to the symbol-grounding problem and novel combinations of logical and probabilistic reasoning. This would pave the way toward a new generation of autonomous robots. High-level reasoning is a key aspect of truly intelligent behavior, which emerges from the ability to make rational decisions and to devise plans for complex tasks, as well as to exhibit a great deal of flexibility dealing with unforeseeable circumstances. Probabilistic reasoning alone won't be enough to reach this level of intelligence because it doesn't allow for the high degree of abstraction needed to cope with a large body of diverse information.

References

1. R. Reiter, *Knowledge in Action*, MIT Press, 2001.
2. M. Shanahan, *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*, MIT Press, 1997.
3. M. Thielscher, *Reasoning Robots: The Art and Science of Programming Robotic Agents,* Springer, 2005.
4. W. Burgard et al., "Experiences with an Interactive Museum Tour-Guide Robot," *Artificial Intelligence*, vol. 114, nos. 1–2, 1999, pp. 3–55.
5. P. Doherty et al., "The WITAS Unmanned Aerial Vehicle Project," *Proc. 14th European Conf. Artificial Intelligence*, IOS Press, 2000, pp. 747–755.

Human-Robot Partnership

Cynthia Breazeal, *MIT Media Lab*

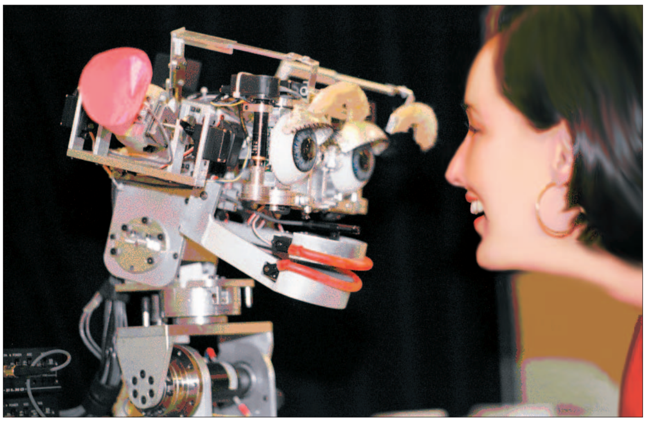

Figure 3. Cynthia Breazeal and her collaborators developed Kismet at the MIT Artificial Intelligence Lab in the 1990s.

What kind of future do you envision with robots? The science fiction stories and movies that fascinated me as a child shape my dream. There are the mechanical droids R2-D2 and C-3PO from the movie *Star Wars*. There are many wonderful examples in the short stories of Isaac Asimov, such as *Robbie*. And more recently, there's Teddy from the movie *Artificial Intelligence*.

These sci-fi robots differ dramatically from today's autonomous mobile robots. We generally view modern autonomous robots as tools that human specialists use to perform hazardous tasks in remote environments (such as sweeping minefields, inspecting oil wells, mapping mines, and exploring other planets). In contrast, my favorite science fiction robots aren't tools that do things for us. Rather, they're socially savvy partners that do things with us.

Recent commercial applications are motivating research into building robots that play a beneficial role in ordinary people's daily lives (see Maja J. Matarić's following essay on socially assistive robots). This new breed of social (or sociable) robot must be natural and intuitive enough for the average consumer to interact and communicate with, work with as partners, and teach new capabilities.

One of the best known social (or sociable) robots is Kismet—the first autonomous robot explicitly designed to explore face-to-face interactions with people.[1] Kismet is considered humanoid because of its expressive anthropomorphic face, even though it doesn't have a humanoid body (see figure 3). The research with Kismet focused on exploring the origins of social interaction and communication in people, namely that which occurs between caregiver and infant, through extensive computational modeling guided by insights from psychology and ethology.

In particular, Kismet was one of the earliest works to explore socio-emotive interaction between humans and robots. Inspired by our earliest infant-caregiver exchanges, which are heavily grounded in the regulation of emotion and its expression, Kismet used its emotion system and corresponding expressive modalities to communicate internal affective states reflecting the degree to which its drives and goals were being met. Internally, Kismet's emotion models interacted intimately with its cognitive systems to influence behavior and goal arbitration. In effect, Kismet socially negotiated its interaction with people via its emotive responses so that humans would help it achieve its goals and satiate its drives. The scientific literature on emotion theory, social cognition, and its development through infancy and childhood has inspired and guided the development of several socially interactive robots.[2]

Challenges and progress in social robotics

As with any emerging area of inquiry, the fundamental scientific questions that social robotics pursues distinguish it from other areas.[1,2] A few challenges of social robotics that I'm particularly interested in include these three:

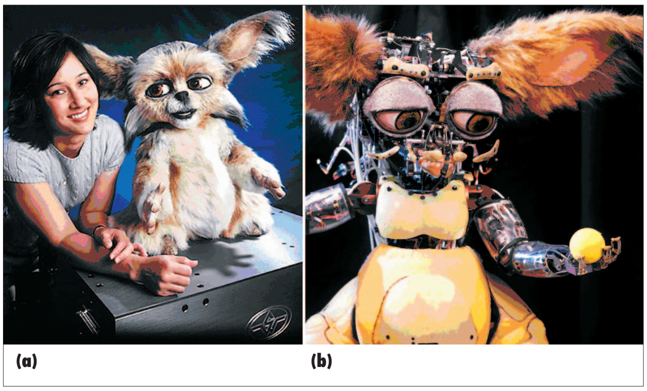

Figure 4. Leonardo, (a) with and (b) without cosmetic finishing, is being developed at the MIT Media Lab (aesthetic design by Stan Winston Studio). (images courtesy of Sam Ogden Photography)

- *Robots with social-cognitive skills*. Whereas research with modern autonomous robots has largely focused on their ability to interact with inanimate objects, socially intelligent robots must understand and interact with people. They must understand people in social-psychological terms to appreciate the goals, beliefs, feelings, motives, and other mental states underlying our behavior.
- *Robots that collaborate with people as full-fledged partners*. Socially intelligent robots work alongside humans as peers, communicate with us in human-centric terms, and participate in successful relationships with people that provide benefits over an extended period of time.
- *Robots as social learners that learn from people*. Socially intelligent robots engage in rich forms of social learning with people, such as imitation and tutelage, to learn new skills, knowledge, and tasks while on the job.

Promisingly, there have been initial and ongoing strides in all of these areas.[2–4]

This article briefly surveys my own group's research at the MIT Media Lab and our ongoing progress on these three challenges (see http://robotic.media.mit.edu). Leonardo is the expressive full-body humanoid robot that serves as our experimental platform (see figure 4). We study and evaluate the robot's performance from a technical perspective to assess its ability to communicate with, collaborate with, and learn from people in situated task scenarios. Additionally, we conduct human subject studies to assess whether the robot's performance is cognitively and behaviorally synergistic for people and to confirm that the human-robot team is more effective as a whole.

Robots with social-cognitive skills

Our computational models are deeply informed by scientific theories of human and animal social cognition to endow our robots with useful social-cognitive skills that are also human compatible. In particular, recent experimental data from neuroscience and psychology provide empirical evidence that human social cognition is deeply embodied.[5] *Simulation theory*, for instance, posits that we use our own body and cognitive and affective processes to simulate the mental processes of others—to take their mental perspective to infer their intents, beliefs, affect, motives, and so on.[6] Inspired by this work, we've been developing a simulation theoretic framework to endow our robot with mental perspective-taking skills. For instance, we've demonstrated the robot's ability to learn affective appraisals of novel objects via social referencing or to reason about and infer its human partner's goals, attention, beliefs, and affect from their observable behavior in various initial tests, such as the *false-belief task*. (The false-belief task, which determines whether a person can appreciate that others can have beliefs that differ from his or her own, is used as a diagnostic tool in developmental psychology.) These skills play an important role in robots that collaborate with humans as teammates and learn from human teachers.

Robots and people as full-fledged partners

Collaborative discourse theory has provided a rich theoretical framework to understand and model human-style collaboration for human-robot teams. Researchers have developed sophisticated discourse models to support human-computer collaboration where the human and computer establish and maintain shared beliefs, goals, and plans to coordinate joint action through natural language. Human-robot teamwork introduces an important nonverbal dimension where physical actions communicate mental states (for example, gaze direction communicates attention, facial expression communicates emotion, and actions communicate intentions). These mental and body states must be coordinated between human and robot to successfully collaborate on a physical task in a shared space.

We've demonstrated our robot's ability to apply its social-cognitive skills to compare and reason about how its human partner's goal and belief states (as communicated through verbal and nonverbal behavior) relate to its own "mental" states to provide the person with informational and instrumental support. For example, in the case of informational support, the robot can relate its own beliefs about the state of the shared workspace to those of the human on the basis of the visual perspective of each. If a visual occlusion prevents the human from knowing important information about that region of the workspace, the robot knows to direct the human's attention to bring that information into common ground. On the basis of the principles of the *joint intention* and *speech act* theories, the robot uses a versatile range of nonverbal behaviors to coordinate effective, efficient teamwork. In the case of instrumental support (that is, performing an action to help someone achieve a task-oriented goal), the robot can infer the human's intent (that is, a desired effect on the workspace) from observing the person's behavior. If the human fails to achieve his or her goal, the robot can reason about how it might help the human by either performing a direct action that achieves that goal or by executing a complementary action that enables the human to achieve the goal (for example, the robot might disengage the lock on a device to help the human activate that device).

Robots as social learners

The ability to learn new tasks and skills quickly and effectively from real-time interaction with everyday people is critical for socially situated robots. I believe this entails understanding how people naturally want to teach robots and then using these constraints to design learning algorithms and architectures to support the teaching-learning interaction.

We've carried out human subject studies to explicitly study how people teach robotic agents. Our results empirically support the popular assumption that people want robots to engage in social forms of learning (such as by demonstration, imitation, verbal instruction, or animal training). In particular, people want to guide exploration, motivate the learner, and adapt their instruction to be appropriate for the learner.

Given our findings, we're developing a learning framework called *socially guided machine learning*. In SG-ML, the robot is intrinsically motivated to learn and explore on its own and is socially motivated to leverage its interactions with people to learn new skills, goals, and tasks. Importantly, we view the teaching-learning interaction as a collaboration in which the human teacher guides the robot's exploration process and the robot communicates its internal state (such as what it understands and where it's confused) back to the teacher to help guide teaching. By making its learning process transparent to the human, the robot actively improves its own learning environment by helping the human better tune their instruction for the robot. We've demonstrated this bidirectional process whereby the quality of the human's guidance improves and the robot becomes a more efficient and effective learner.

Furthermore, we've found that the robot's teamwork and perspective-taking skills also play an important role in making teaching-learning interactions more robust to the miscommunications or misunderstandings that inevitably arise even in human-human tutelage. As the robot observes the human's demonstrations, it internally simulates "what might I be trying to achieve were I performing these demonstrations in the human's context?" The robot therefore interprets and hypothesizes the intended concept being taught not only from its own perspective but also from the human teacher's visual perspective. Through this process, the robot successfully identifies ambiguous or confusing demonstrations given by the human instructor and clarifies the human's intent. After disambiguating these problematic demonstrations, the robot correctly learns the intended task.

My goal

Unlike AI's original goal—to create a technological system with human-equivalent intelligence—my goal is to create robots that are synergistic and compatible with humans. Specifically, they bring value to us and are valued by us because they're different from us in ways that enhance and complement our strengths and abilities. The goal of creating robots that can engage us as full-fledged partners is as challenging and deep as AI's original goal because it requires scientists of the natural and artificial to deeply understand human intelligence, behavior, and nature across multiple dimensions (that is, cognitive, affective, physical, and social). It also requires scientists to understand the dynamics of the human-with-robot system. This fascinating area of technological and multidisciplinary scientific inquiry has the potential to transition the robots of today into our sci-fi robot sidekicks of tomorrow.

> This fascinating area of technological and multidisciplinary scientific inquiry has the potential to transition the robots of today into our sci-fi robot sidekicks of tomorrow.

References

1. C. Breazeal, *Designing Sociable Robots*, MIT Press, 2002.
2. T. Fong, I. Nourbakshsh, and K. Dautenhahn, "A Survey of Socially Interactive Robots," *Robotics and Autonomous Systems,* vol. 42, nos. 3–4, 2003, pp. 143–166.
3. R.W. Picard, *Affective Computing*, MIT Press, 1997.
4. S. Schaal, "Is Imitation Learning the Route to Humanoid Robots?" *Trends in Cognitive Science*, vol. 3, no. 6, 1999, pp. 233–242.
5. L.W. Barsalou et al., "Social Embodiment," *The Psychology of Learning and Motivation*, B. Ross, ed., Academic Press, vol. 43, 2003, pp. 43–92.
6. A. Meltzoff and J. Decety, "What Imitation Tells Us about Social Cognition: A Rapprochement between Developmental Psychology and Cognitive Neuroscience," *Neuroscience of Social Interactions: Decoding, Influencing, and Imitating the Actions of Others*, D. Wolpert and C.D. Frith, eds., Oxford Univ. Press, 2003, pp. 491–500.

Socially Assistive Robotics

Maja J. Matarić, *University of Southern California*

Five decades into its development, robotics is finally on the brink of entering our daily lives. Having transitioned from success in factory automation (auto assembly) and high-precision laboratory work (aiding DNA sequencing and medication delivery), it is now poised to move into much messier, less structured, and more uncertain human-inhabited environments. While simple robot vacuum cleaners are already at work in over a million homes, robotics researchers are eagerly developing much more complex and intelligent robots that will work with and for people in hospitals, elder care centers, schools, and homes.

What happens when intelligent robots and people share an environment and even goals? This question is at the heart of human-robot interaction, an inherently interdisciplinary research endeavor. HRI brings together engineering and social science, cognitive science, neuroscience, ethics, and the disciplines directly related to the application domains the robot technology would operate in, such as health-care and education. HRI is as much a study of people as it is of robots; it requires us to gain insight into a new and dynamically changing relationship between humans and intelligent machines.

Within the larger context of HRI, the

research in my Interaction Lab focuses on *socially assistive robotics*, which aims to aid people by means of social interaction rather than physical contact. SAR brings together HRI and assistive robotics, which has demonstrated successes in hands-on rehabilitation through the use of high-precision but low-autonomy and low-interactivity systems. Consequently, SAR presents an entirely new set of research challenges, as we strive to understand how robots can interact with people to effectively and measurably help in the difficult processes of convalescence, rehabilitation, socialization, training, and education.

Using physical robots for social assistance might seem counterintuitive. After all, why not use software agents or devices such as PDAs endowed with intelligent interfaces? The answer rests on the fundamentally human tendency to socially engage with and attribute lifelike properties to machines, especially embodied ones such as robots, which exhibit sufficient (often quite simple) attributes of biological-like movement or appearance. Researchers have shown the human projection of intentions, goals, and emotions to such embodied devices to be inescapable, culture independent (although featuring culture-specific responses), uninhibited by sophistication of knowledge and familiarity with the details of the technology, and resistant to change over time or with repeated encounters. This means that people cannot help but respond to engaging physical machines.

Our research builds on this human characteristic to develop a novel robot-assisted research paradigm with two distinct but interrelated goals:

- to study human social and performance-related behavior to better understand those phenomena and their underlying mechanisms and
- to elicit sustained productive goal-driven behavior as part of the diagnostic, therapeutic, or educational process.

We're motivated by the vision of using assistive technology as a tool for both scientific discovery and societally relevant application of assistive robots.

From the basic-science perspective, SAR presents a unique new tool for scientific inquiry. Machines have long been used in research, but usually as passive (if sophisticated) tools. Through socially aware robots, we can introduce a fundamentally new type of machine into the research realm: one that can detect, monitor, and respond to user behavior in a physically shared context. As tools of scientific inquiry, robots can have both their form and behavior manipulated in a fully controllable manner. We can carefully design the information that we share with participants about the robot to attempt to control their beliefs and biases. Using that experimental paradigm, we can measure participants' responses, both voluntary and involuntary, ranging from observable behavior (such as gaze, posture, movement patterns, and linguistic interactions) to physiological responses (such as heart rate, body temperature, and galvanic skin response). We can also combine robot experiments with functional brain imaging, aiming toward structural insights about relevant social and learning mechanisms. As a result, we can ask scientific questions that we could not previously ask or experimentally answer. By perfectly tuning how lifelike, humanlike, believable, socially engaging, and intelligent robots are, we can probe into human social behavior and use the insights to design social interaction toward measurably improved therapeutic, educational, and training outcomes.

Through socially aware robots, we can introduce a fundamentally new type of machine: one that can detect, monitor, and respond to user behavior in a physically shared context.

The robot's physical embodiment is a critical means of eliciting the richness of the user's response, both in the context of scientific discovery and in real-world applications. Properties and implications of embodiment present rich avenues for research into both human nature and assistive human-robot interaction. Such properties include

- how the robot looks (Is humanlike or animal-like appearance important and desirable? Should the robot match the user in approximate size and, if humanoid, apparent age?),
- how the robot behaves (How expressive should the robot be relative to its user? What means of expression are most effective? How far do physical embodiment and embodied cues go in conveying information, and what is the appropriate balance between nonlinguistic and linguistic interaction? Should the robot have a detectable personality? If so, what should it be?), and
- how the robot relates to its niche (How does the hospital environment differ from school or home? How should the robot handle different types of users in an environment, such as doctors, nurses, therapists, patients, and family at the hospital?).

Situated at the intersection of science and technology, SAR research aims to help solve major societal challenges, focusing specifically on the growing aging population, rising costs of medical care and education, and shortages of social programs for education and training of persons with disabilities. To aid people in need, robots will have to perform truly challenging tasks in both structured environments (hospitals, elder care facilities, and schools) and unstructured settings (homes, cities, roadways, and rubble piles).

An outstanding challenge for SAR is the robots' need to be both useful and engaging. The most popular human nurses are not known to be the most effective; it is hard indeed to compel a person to do something painful, difficult, or dull in a sustained manner, such as exercising six hours daily post-stroke, breathing into a spirometer 10 times an hour after cardiac surgery, or performing months of vocabulary and pronunciation drills. The socially assistive robot must remain appealing as well as effective over a long period, whether months in stroke rehabilitation, years in special education, or, potentially, a lifetime. This presents a novel, compelling challenge for both robotics and AI. Effective human teachers, coaches, nurses, and therapists are all too rare yet too important to do without. Those shortages serve as motivators for SAR research, whose goal is not to replace humans but to fill in large gaps where attentive and individ-

Figure 5. A stroke patient interacting with a socially assistive robot from Maja Matarić's lab that is monitoring, coaching, and motivating prescribed exercise therapy.

ualized human care is diminishing or entirely unavailable.

Embarking on assistive robotics involves being not only curiosity-driven but also problem-driven, and facing complex scientific and technological challenges in messy, noisy, and sensitive real-world domains that involve interactions with vulnerable users. We have placed simple robots with stroke patients (see figure 5), cardiac patients, and children (see figure 6) in special education classrooms. Our research has demonstrated positive and promising results,[1–5] as well as pointed toward a plethora of fascinating research questions. SAR's rich potential for gaining novel insights into human cognition and social behavior, and for improving human quality of life for populations that most need it, represents one of the most exciting and uniquely compelling topics in modern robotics.

Figure 6. A child-sized humanoid robot, which Maja J. Matarić and her colleagues developed, interacts with a child. The robot works from a desktop or on a mobile platform and is an example of the research toward robots that can help in the education and socialization of children with special needs.

References

1. J. Eriksson, M.J. Matarić, and C. Winstein, "Hands-Off Assistive Robotics for Post-Stroke Arm Rehabilitation," *Proc. IEEE Int'l Conf. Rehabilitation Robotics* (ICORR 05), IEEE Press, 2005, pp. 21–24.
2. D. Feil-Seifer and M.J. Matarić, "Socially Assistive Robotics," *Proc. IEEE Int'l Conf. Rehabilitation Robotics* (ICORR 05), IEEE Press, 2005, pp. 465–468.
3. R. Gockley and M.J. Matarić, "Encouraging Physical Therapy Compliance with a Hands-Off Mobile Robot," *Proc. 1st Ann. Conf. Human-Robot Interaction*, (HRI 06), ACM Press, 2006 pp. 150–155.
4. K.I. Kang et al., "A Hands-Off Physical Therapy Assistance Robot for Cardiac Patients," *Proc. IEEE Int'l Conf. Rehabilitation Robotics* (ICORR 05), IEEE Press, 2005, pp. 337–340.
5. A. Tapus and M.J. Matarić, "Introvert vs. Extrovert Hands-Off Robot for Post-Stroke Rehabilitation Therapy," to be published in *Proc. 10th Int'l Symp. Experimental Robotics* (ISER 06), 2006.

Do Androids Dream of the Sweet Smell of Roses?

Hiroshi Ishida, *Tokyo University of Agriculture and Technology*

Robotics and intelligent systems research have historically aimed to make machines that can replace human workers, that can do jobs that humans can't do, or both. Industrial robots, for example, are designed to more quickly and accurately accomplish tasks that human workers formerly did. In this context, robots must do their jobs autonomously with minimal interaction with and supervision from human users. However, human-machine interaction has become significantly more important as the application areas of robots and intelligent systems have expanded. Home care robots for the elderly, for example, must be able to understand users' demands and to tell them what actions the robots are about to take.

This is one reason that robotics researchers are eager to implement olfactory sensing capabilities in robotic and intelligent systems. It's desirable for robots working with humans to share sensory information with human users. If the robot can see the world and hear sound exactly as the human user does, it can perceive the situation similarly, and working cooperatively with the human will become much easier. However, this won't be enough. Although most people don't appreciate the value of olfaction in our daily lives, our noses play essential roles in various scenarios, such as enjoying meals and detecting danger. A robot cook that can't discriminate old milk with a nasty smell, a security robot that can't detect the smell of burning, or a nanny android that can't smell that it's time to change diapers all lack an important functionality.

Taking inspiration from humans or other animals

For robots that work closely with humans, the ultimate goal would be an olfactory capability that reproduces human olfaction. We sense odors when chemical compounds stimulate the olfactory receptor cells at the nasal cavity's roof.[1] What makes olfaction a distinctive sensing modality is the diversity in the stimuli. Several hundreds of thousands of different chemical substances are known to elicit odor sensation. Moreover, what we perceive as a single smell is often the result of a mixture of a number of chemical substances. The vapor emanating from a cup of coffee includes more than 400 chemicals that have odors.

Such complexity and diversity is the result of the olfactory system's huge number of sensors. The human nose has 10 million olfactory receptor cells. Genetic-engineering advances have revealed that those cells are classified by their receptor proteins into approximately 400 types. The levels of interactions between odorant molecules and the receptor proteins determine the cells' response levels. Each class of cells responds to a wide spectrum of chemical substances, but in a slightly different way from the cells of the other classes. So, exposing the cells to a single puff of a smell produces a response pattern unique to that smell. We discriminate odors by recognizing the cells' response patterns.

Human olfaction isn't the only model for building artificial olfactory sensing systems. Many animals have keener olfaction than humans. We might use such "supersenses" as a model to enable our robots to do things that humans can't do. Dogs are famous for their ability to track trace scents. It's said that trained dogs outperform metal detectors in landmine detection. Dog noses are both more sensitive and more selective than human noses. However, the basic olfaction mechanism—that is, recognition of the receptor cells' response patterns—seems to be common not only to all mammals but also to a wide variety of animal species. The main difference between human and dog noses is that a dog nose has more olfactory receptor cells with a greater variety of receptor proteins.

Another famous example of super olfaction is male moths that track airborne sexual pheromones released by conspecific females over long distances. To achieve such feats, the olfactory sensors must be able not only to detect extremely diluted target odors but also to discriminate target odors from any background odors. Moths have solved this problem by developing, through evolution, receptor cells more sensitive than humans'. A male moth's pheromone receptor cell responds even to a single molecule of the sexual pheromone, but shows little response to any other chemical substances. Despite their efforts, researchers don't yet fully understand how the pheromone receptor cell differs from the mammalian olfactory receptor cell.

Electronic noses and their applications in robotics

The first research on an artificial sensing system able to discriminate different odors using mammalian olfaction as a model was published in 1982.[2] Since then, researchers have done extensive research on developing *electronic noses*. Typically, an electronic nose consists of an array of chemical sensors with partial specificity and a pattern-recognition system just like its model's. Unfortunately, no engineered replica of olfactory receptor cells exists. The most popular sensor in electronic noses is a metal-oxide semiconductor gas sensor. The sensing element is generally made of sintered tin dioxide particles, which show variation in conductivity on exposure to reducing gases. Although originally made for flammable gas leak detectors, the sensor responds to a wide range of volatile organic compounds. The sensitivity to different gases can be tuned to some extent in several ways, such as changing the operating temperature and doping the sensor with small amounts of catalytic materials. Sensors that use organic materials are expected to be more compatible with human olfactory receptor cells. Adsorption of gas molecules onto organic materials is detectable—for example, by measuring the mass increase or swelling of the polymer matrix. Typically, an electronic-nose system has only eight to 32 sensors because fabricating numerous sensors with different selectivities into a single array is difficult.

> A security robot that can't detect the smell of burning or a nanny android that can't smell that it's time to change diapers lacks an important functionality.

Despite current sensor technologies' limitations, electronic noses have been tested in a variety of contexts. A typical electronic nose setup stores sensor response patterns for a repertoire of odors into the computer memory in the learning phase. The electronic nose recognizes odors by finding the best match between the measured and the stored response patterns. For example, it can estimate the freshness of a piece of fish by comparing its smell with those of meat samples with different levels of rottenness. Researchers have successfully demonstrated electronic noses' odor-discrimination capabilities in applications ranging from quality control of cosmetic products to food process monitoring and medical diagnosis.[1] Some electronic-nose instruments are commercially available.

For these applications, electronic noses are designed to have odor discrimination capabilities similar to human olfaction. However, artificial olfaction can have different selectivities to odors by using different sensors, just as in animals' supersenses. To date, supersensitive gas sensors have been developed without a specific animal as a model because the mechanisms underlying animal olfaction are still open questions. One key to realizing highly sensitive sensors is a built-in signal-amplification mechanism. For example, Nomadics' landmine detector has sensitivity comparable to dog noses in detecting vapor emanating from explosives.[1] The key element is a fluorescent conjugated polymer film that has a high affinity for trinitrotoluene. Adsorption of a single TNT molecule quenches the whole polymer chain's fluorescence, causing significant variation in the sensor output.

In robotic applications, detecting and discriminating a specific odor might be just a part of the whole task. If a security robot with artificial olfaction detects a flammable gas leaking in a factory, an appropriate action would be to find the leak by tracking the gas plume to its source. This task isn't as easy as it seems. Because molecular diffusion is slow, even a slight airflow greatly affects gas dispersal. The flows we encounter in the real world are almost always turbulent. Because the gas plume becomes patchy and meandering, tracking the gas-concentration gradient often fails to locate the gas source. To help solve this problem, researchers have tried to develop mobile robots that mimic moths' instinctive behavior.[3] Male moths predictably steer upwind when they come in contact with a plume of sexual pheromone. The basic feature of plume-tracking behavior was implemented into a robot equipped with gas sensors and airflow sensors. Successful demonstrations have been provided on tracking gas plumes formed in indoor environments.[4]

Outlook

The major problem in the development of artificial olfaction is that no sensor as versatile as odor receptor cells exists. Each sensor has certain advantages and disadvantages, so users must select a set of sensors appropriate for their applications. Because the methodology for systematic sensor selection hasn't yet been established, selection often involves much trial and error. Moreover, all gas sensors show drift in their response over time, and frequent recalibration is laborious. Human olfaction involves both sensors of many different types and many sensors of the same type. The former makes our noses versatile; the latter appears to help alleviate drift. Olfactory receptor cells regenerate in four to eight weeks. Although at least some sensors are always being replaced, we can still recognize the same odors as before the sensor replacement.

Recently, research has been initiated to exploit olfactory receptor proteins themselves as the ultimate sensing materials.[5] Challenges include how to keep receptor activity on a sensor chip, how to prepare receptor proteins for mass production, and how to measure with high sensitivity the adsorption of odor molecules onto the receptors. Once these problems are solved, we'll obtain a large family of artificial sensors with selectivities closely matched to human olfaction. If we could somehow add the self-generating function to those sensors, we should be able to equip our androids with versatile, drift-free artificial olfaction.

On the software side of artificial olfaction, an interesting issue is how to deal with unknown odors. So far, electronic noses merely report which smell in their learned repertoire elicits the sensor response pattern that best matches the one being measured, even if the repertoire doesn't include the measured odor and the match is far from perfect. Some electronic noses can recognize an unlearned odor as "unlearned," but this might not be enough. A robot that works closely with humans should be able to tell human users what smell it's detecting and, if the smell is unfamiliar, how it smells. However, when you try to explain an unfamiliar smell to someone else, you realize the difficulty of describing smells. Our language system seems to be built to describe what we see. Logically describing a perceived smell's features isn't always possible. This is partly because odor perception involves not only conscious odor recognition but also subconscious evocation of memories and emotions. We don't know yet what features in the response pattern of our olfactory receptor cells make us feel sweet when we smell roses. Some researchers are attempting to make robots explain the smells they detect in terms of abstracted descriptions, such as herbal, oily, or fragrant.[6] In theory, if we succeed in providing an array of sensors compatible with human noses, we should be able to let our robots appreciate the sweet smell of roses. ■

References

1. T.C. Pearce et al., eds. *Handbook of Machine Olfaction: Electronic Nose Technology*, John Wiley & Sons, 2003.
2. K. Persaud and G. Dodd, "Analysis of Discrimination Mechanisms in the Mammalian Olfactory System Using a Model Nose," *Nature*, vol. 299, no. 5881, 1982, pp. 352–355.
3. R.A. Russell, *Odour Detection by Mobile Robots*, World Scientific, 1999.
4. H. Ishida et al., "Controlling a Gas/Odor Plume-Tracking Robot Based on Transient Responses of Gas Sensors," *IEEE Sensors J.*, vol. 5, no. 3, 2005, pp. 537–545.
5. G. Gomila et al., "Artificial Nose Integrating Biological Olfactory Receptors and NEMS," *Proc. 11th Int'l Symp. Olfaction and Electronic Nose*, 2005, pp. 14–18.
6. A. Loutfi and S. Coradeschi, "Forming Odour Categories Using an Electronic Nose," *Proc. European Conf. Artificial Intelligence* (ECAI 04), 2004, IOS Press, pp. 119–124.

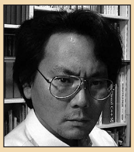

Hiroshi Ishiguro is a professor at Osaka University. Contact him at ishiguro@ams.eng.osaka-u.ac.jp.

Minoru Asada is a research director at the JST ERATO Asada Synergistic Intelligence Project and a professor at Osaka University. Contact him at asada@ams.eng.osaka-u.ac.jp.

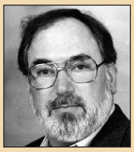

Stuart C. Shapiro is a professor in the Department of Computer Science and Engineering and director of the Center for Cognitive Science at the State University of New York at Buffalo. Contact him at shapiro@cse.buffalo.edu.

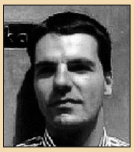

Michael Thielscher is a professor at Dresden University of Technology. Contact him at mit@inf.tu-dresden.de.

Cynthia Breazeal is an associate professor at the MIT Media Lab. Contact her at cynthiab@media.mit.edu.

Maja J. Matarić is a professor of computer science and neuroscience, founding director of the Center for Robotics and Embedded Systems, and senior associate dean for research in the Viterbi School of Engineering at the University of Southern California. Contact her at mataric@usc.edu; http://robotics.usc.edu/~maja.

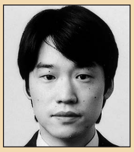

Hiroshi Ishida is an associate professor at the Tokyo University of Agriculture and Technology. Contact him at h_ishida@cc.tuat.ac.jp.