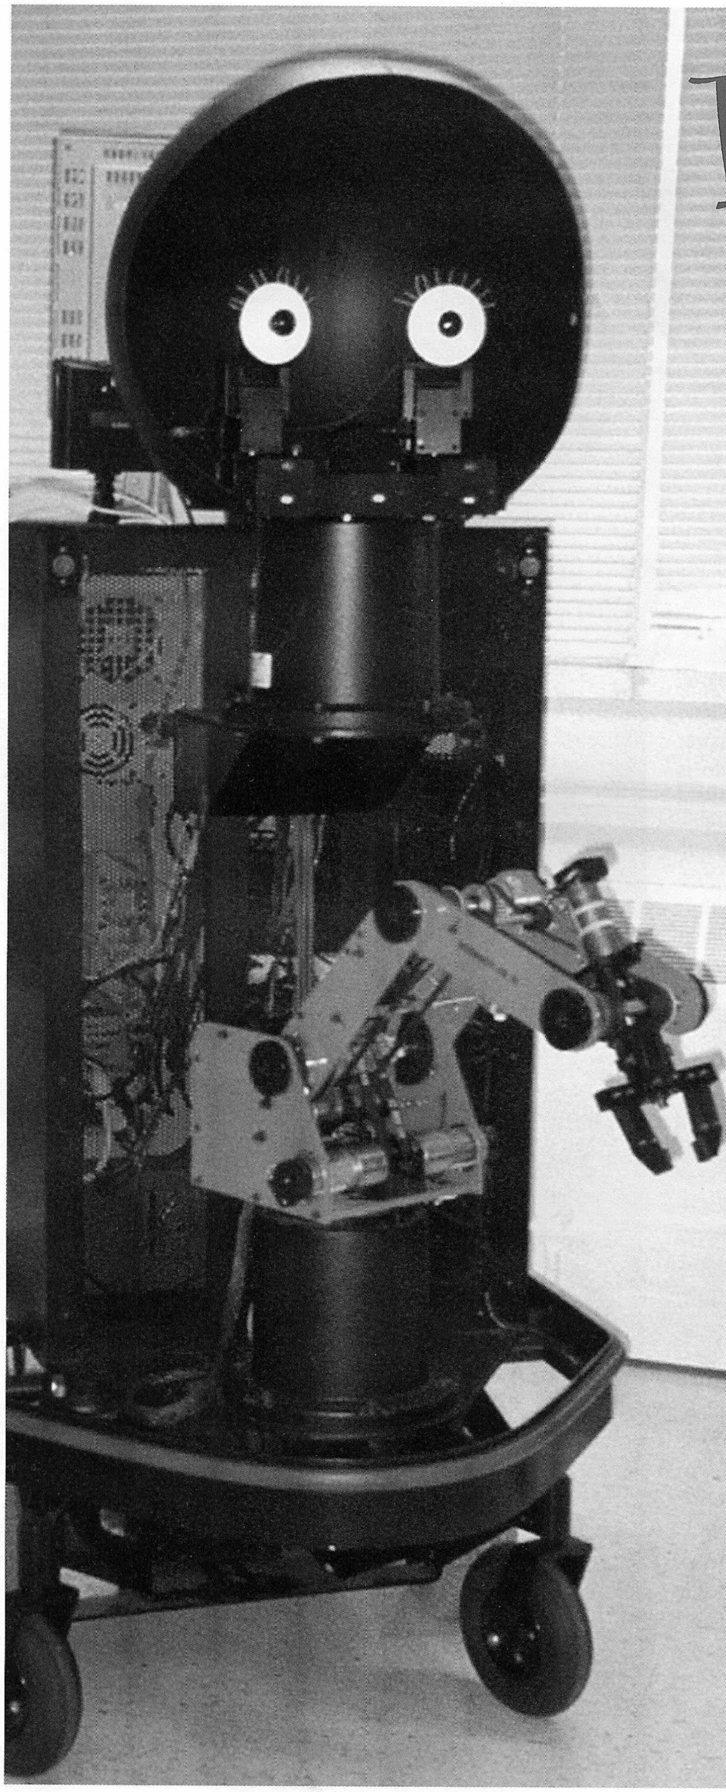

Bringing Up Baby

by Erica Klarreich

"Please don't use that kind of language in front of the robot!"

Setting a good example for an impressionable child is hardly a new idea. After all, human society depends on the careful rearing of infants into responsible, thoughtful adults. But when a youngster is a hulking five-foot-nine assemblage of wheels, gears, and video cameras, giving it a good upbringing may not seem like a high priority.

That is, unless you're computer scientist John Weng. Weng is working to build robots that are as quick to pick up our bad habits as our good ones. And somewhere down the road, he says, they may even achieve the Holy Grail of artificial intelligence: self-awareness.

Experience Is the Teacher

With colleagues at Michigan State University, Weng has built a rudimentary robot named SAIL (Self-organizing Autonomous Incremental Learner) that learns through its experiences with the world, just as a child does. Looking like a cross between a person and an upright vacuum cleaner, SAIL has already learned several of an infant's first activities, such as paying attention to a toy and reaching out to it. It has acquired a vocabulary of about twenty words and has learned to follow simple spoken commands.

From these humble beginnings, Weng hopes to spawn robot progeny that could someday be smarter than he is. "In traditional artificial intelligence, a machine will do reasonable things because the designer is smart, not because the machine is," Weng says. "But a child can learn a language its parents don't understand. That's the power of animal learning."

Robots are taking their time going from the realm of fantasy to the reality of daily life. The reason, Weng says, lies in the way scientists have been trying to design robots. In traditional artificial intelligence, like the kind that went into the chess whiz Deep Blue, a computer scientist writes a program specifically geared towards the task to be executed—a method that works well for something like chess, which has clean rules.

But if you want a robot to perform complicated real-world tasks, Weng argues, the key—paradoxically enough—is not to program the tasks into the robot. Instead, start the robot off with a mind that can structure itself in response to its environment, and then let the robot develop by interacting with the world autonomously, in real time. While many computer scientists have written programs that can learn from experience to perform a particular task, for Weng, it's essential that SAIL should be able to learn tasks that he himself hasn't imagined.

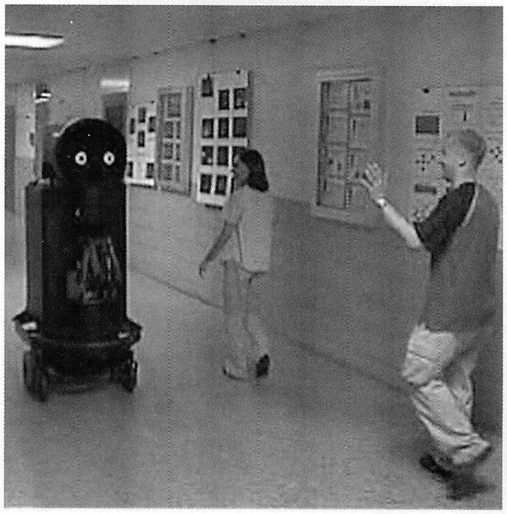

SAIL, which stands for Self-organizing Autonomous Incremental Learner, is Michigan State University's version of a robot "child." SAIL interacts independently with its surroundings and learns from its mistakes—and successes. Actions such as navigating the hallways of the university's engineering building (*above*) and reaching out to grab a toy (*right*) are the result of SAIL's learning by trial and error rather than by preprogrammed instructions.

Weng has endowed SAIL with a developmental program that allows a "robot-sitter" to teach it in much the same way a parent teaches a child. To teach SAIL to reach out to a moving object, for example, the researchers started by moving SAIL's arm to grab the object, the way a parent guides a child's hand when teaching the child to write. After many repetitions, SAIL practiced on its own, receiving encouragement or discouragement from trainers via "good" and "bad" buttons on its back.

Eventually Weng plans to use language to let the robot know how it's doing. "Good job" may be the first lesson. In any case, positive or negative feedback from the trainers becomes part of the information that the robot stores about each experience and considers when it's deciding on a new action. Over time, this will enable the robot to develop long-term values. Humans build up complex values out of such simple beginnings, Weng says.

For now, though, Weng is content with more modest aspirations. Recently, his team has been training SAIL to navigate the intricate corridors of the university's engineering building. "It'll take ten or twenty tries to learn the first corner you show it," says Jason Massey, one of Weng's graduate students. "But the next one takes only half as many tries."

Two video monitors on SAIL's back show its teacher what the robot is seeing at every moment. But the teacher can't see into SAIL's brain—which, Weng says, is the point. Once SAIL's developmental program was put into place, Weng locked off its brain from programmers who might feel the itch to tamper with it to make SAIL perform a particular task.

Weng is "taking a huge step towards something much more brainlike than what is normally being done," says James McClelland, a psychologist and computer scientist at Carnegie Mellon University. "What John is fundamentally trying to do is say, 'Look, the human brain doesn't have a programmer. It programs itself in the course of experience. What we need to do is find ways of creating machines that program themselves.' "

A growing body of evidence supports the idea that our own brains are molded by our surroundings. "Although there is certainly some predisposition in the brain," McClelland says, "experience does a lot more in shaping what we come out with than people have argued in the past."

Making Memories

Weng has engineered SAIL to base its actions on the experiences it has already laid down in its memory. To teach SAIL a command like "turn left," for instance, researchers repeatedly turn the robot while saying the words. On each go, SAIL stores a wealth of information, including what it sees as it turns, the sound of the trainer's voice saying "turn left," and the direction and speed of its wheels. If the trainer gives SAIL a command and doesn't impose an action, SAIL sifts through its memory to find the past experiences closest to the current one, and uses those to determine its current action. So if SAIL is instructed to turn left, it will see that previous successful responses all involved turning its wheels in a particular way and act accordingly.

SAIL can't store every experience it has, of course. Imagine if our brains tried to recall every detail of everything that happened to us—we'd soon be overwhelmed. "Each image SAIL takes in has a quarter of a million pixels," Weng says. "It can't record everything, only what's essential." To allow SAIL to respond in real time, Weng wrote a clumping mechanism into the developmental program. When SAIL has similar experiences, instead of storing each one separately, it gathers them into what Weng calls a prototype memory, in the same way a young child might look at many flat surfaces with four legs and finally decide they are all specimens of something called

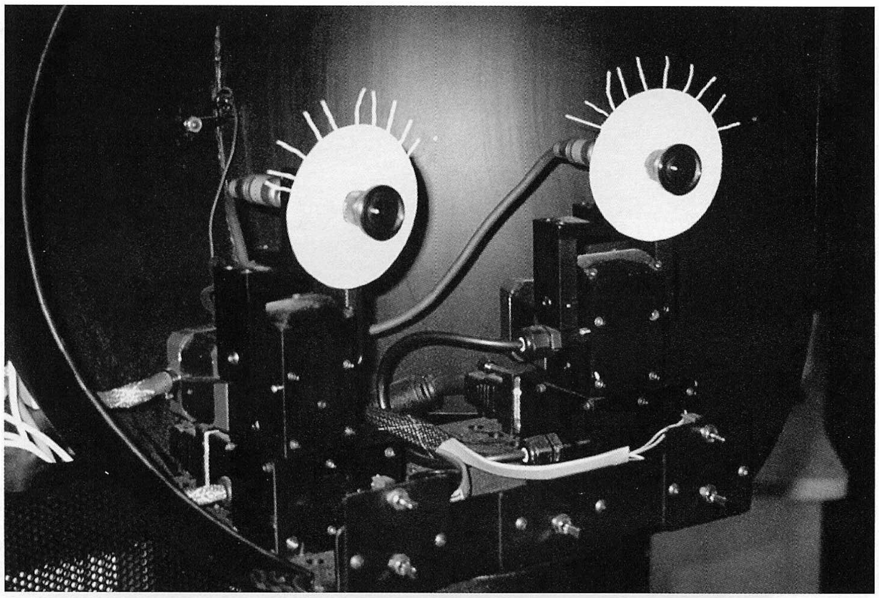

a "table." Thus, for example, SAIL eventually gathers together the memories associated with the words "turn left" into a prototype, which it refers to and updates as needed. By grouping and regrouping its memories, SAIL constantly restructures its brain in response to the environment.

Learning to Forget

Weng has bestowed another human trait on SAIL that is perhaps undervalued by its possessors: forgetfulness. If one of SAIL's memories fails to be reinforced by succeeding memories, it fades gradually, according to a precise rule. When the memory gets faint enough, it's erased.

Sorting through its brain for faded memories or experiences to merge can take SAIL anywhere from minutes to hours, so when there's a large backlog of memories to be organized, the researchers switch the robot into a sleep mode in which it temporarily stops interacting with the world. So far, the trainers are putting SAIL to sleep manually, but eventually Weng plans to teach SAIL to obey the robot equivalent of "Okay, kids, time for bed."

With time out for sleep, SAIL can keep pace with its environment when it's awake, reacting to changes in real time. And since it's not distracted by activities such as nursing, crying, and diaper changing, it learns some tasks faster than a human baby. Teaching SAIL to reach out and grab a toy, for instance, took the researchers about half an hour. They are now working on building SAIL's vocabulary, trying to teach it more complicated tasks such as recognizing faces, and thinking about ways to improve the developmental program.

The Right Approach?

Of course, the list of things SAIL can't do is much longer than the list of things it can do. And some researchers don't believe that SAIL—or any other robot without preprogrammed cognitive abilities—will be able to develop the traits we think of as strictly human.

Researchers who try to leap past understanding specific tasks run the risk of oversimplifying, warns Patrick Winston, a professor in the artificial intelligence lab at the Massachusetts Institute of Technology. "It's hard to make a program learn if you don't know what, in general terms, it is to learn, and romantic to suppose that you can," he says. "I'm sure we will learn something from their approach, but I do not think it is the answer, in the sense that all else should be abandoned."

Arthur Markman, a psychologist at the University of Texas, agrees that we shouldn't discard traditional approaches lock, stock, and barrel. "But I do think this is an important approach, and it's a good group of people doing it," he says. "We won't know what the limitations are of a particular approach until people actually go out and build the robots, and show what you can or can't do."

But MIT computer scientist Marvin Minsky argues that it's improbable that a purely self-organizing machine will ever be able to generate the complex representations that go into higher-order thought. He points out that cats and mice have much more powerful brains than the SAIL robot, but still they seem to be blocked from learning language. "No matter how much experience they have, there seem to be things these animals simply can't represent," he says.

No one can dispute that, agrees cognitive scientist Mriganka Sur, also at MIT, but no one knows whether a cat or mouse is held back by the structure of its brain or simply by its size. "We have to ask, what is in human brains that is not in cat brains," he says. "One of the amazing things is that we have no clue what the answer is. It would be prematurely limiting ourselves to put arbitrary limits on what robot brains can do."

Weng doesn't expect SAIL to learn complex language any time soon, but the limitations are in the hardware, not the principle, he says. Like human beings, a robot is only as good as the materials it's made of. As more people start building developmental robots, Weng predicts, we'll see dramatic improvements in the quality of the robots' eyes and ears, their body design, and perhaps most crucially, in their processors.

If intelligent robots can someday be built, Weng says, one of their great advantages will be their immortality. Weng likes to tell his students that when Albert Einstein died, his brain was wasted because we couldn't download it into another body. "Biological intelligence is very much tied to the body," Weng says. "But when a robot's hardware gets worn, you can download its developed brain, then upload it into another body."

Developmental robots like SAIL will probably first appear in the toy and entertainment industry, Weng says, maybe within just a few years. Already, a major toy company has been sniffing around his laboratory, looking for ideas for "smart toys." Beyond that, Weng sees developmental robots working in complex, ever-changing human environments. Tasks might include guarding major bridges against potential attacks, or carrying out wartime reconnaissance missions. Unlike robots that have been programmed to learn or perform specific tasks, developmental robots will be able to handle the unpredictability of situations. By sending these robots where danger or extreme environments bar humans, "we can extend our physical limit," he says. "That's a very exciting thing." ☺