Building a Baby Brain in a Robot

Linking artificial intelligence and robotics, researchers are constructing an infantile automaton to learn how the body influences cognition

CAMBRIDGE, MASSACHUSETTS—Excitedly talking about plans for her next child, Lynn Andrea Stein doesn't even pause when a robot the size of a hat box scurries by the open door of her office. Stein, a roboticist who works at the Massachusetts Institute of Technology's celebrated Artificial Intelligence (AI) lab, is already the mother of two human infants, and she seems much too busy to become a parent for a third time.

Stein's next offspring, however, will not be flesh and blood but a jumble of steel rods, motors, video cameras, microphones, and computer chips—the most ambitious humanoid robot ever attempted, she says. It's a machine she hopes will eventually approach, and in some areas perhaps surpass, the mental and physical prowess of a six-month-old infant—a leap ahead in robotics, where many machines are struggling to navigate rooms or butter toast. Complete realization of that goal could take more than a decade of work and millions of dollars, if it's even possible. "When people ask me about this project, I describe it as building a baby the hard way—and not because the other way is easy," says Stein.

She is collaborating in this grand venture with one of the most controversial figures in AI: fellow MIT roboticist Rodney Brooks. The goal of their infantile automaton is to provide a physical system that can test explanations of human cognition. Using mechanical eyes, ears, arms, and a computer brain inspired by human neuroanatomy, the robot will tackle questions such as how eye-hand coordination develops in the brain and how infants learn to interact with others. From observations of live babies and adult humans, psychologists have offered dozens of conjectures for the genesis of sensory and social capacities, some of which have even been successfully simulated on the computer screen.

The baby Stein and Brooks are planning could finally let researchers determine how realistic these ideas are by seeing if they reproduce the expected behavior in a body that interacts on a human time scale with the real world. "We should be able to put to the test a huge bunch of hunches, controversial hunches that have just been lying around," says team member Daniel Dennett, a Tufts University philosopher and AI investigator. The robot might even yield insights into "embodiment", the controversial idea that the ill-defined attribute of "intelligence" is a relatively small step up from the swirl of cognitive processes the brain uses to guide the body in its day-to-day physical interactions with the outside world.

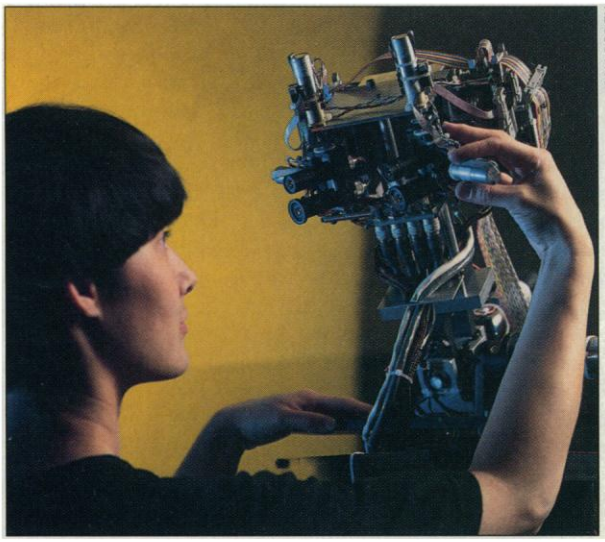
SAM OGDEN

Raising a robot. Cog, the MIT android designed to probe cognition, peers at grad student Cynthia Ferrell.

The name of this mechanical offspring is Cog, a play on the ideas of grinding gears and cognitive processes. Though about the only thing the far-from-complete robot has done to date is track objects with its eyes, Cog's ambitious agenda is drawing applause mixed with guffaws as the research community slowly hears of it. "People are just flabbergasted at the chutzpah of the project," says Dennett. "It's not standard run-of-the-mill robotics by any means," says a delighted Rod Grupen, co-director of the Lab for Perceptual Robotics at the University of Massachusetts. But others already label Cog a ridiculous excursion and say the research community is far too ignorant about the details of cognition to waste time trying to construct a robot.

"Bad Boy of Robotics"

For Brooks, whose graduate students wear T-shirts proclaiming him the "Bad Boy of Robotics," controversy is nothing new. In the 1980s, his lab built more than two dozen small, fast, insect-like robots that challenged the conventional wisdom about mobile, autonomous machines (*Science*, 21 May 1990, p. 959). Traditional AI holds that moving robots must extract a symbolic representation of the world—a map, for instance—from their sensory data before they could plan a course of action within that symbolic structure.

Brooks dismissed that approach as computationally draining and workable only in safe, unchanging environments. His insect robots, in contrast, function through a much more direct coupling with the world, using simple sensors to continually modify and incite physical actions. When a leg of one his insect robots hits an obstacle, for instance, a contact sensor quickly invokes a reflex action like "lift leg" or "back up." In Brooks' robots, dozens of these stimulus-response reflexes are linked together in a rich web, and the resulting behavior is surprisingly complex, robust, and, most of all, quick.

Many AI researchers, however, argue that Brooks' strategy holds little interest when it comes to behaviors of higher creatures. And that is a challenge Brooks can't resist. "I think there's an interpretation of neuroscience and cognitive science which says you don't need those central representations," he says, explaining that the brain must have developed thousands of solutions to everyday problems like seeing, hearing, and moving. Higher intelligence may simply be a matter of adapting those solutions in a novel way. "When you're building in the machinery that knows not to bump into things, how to get around them, et cetera, you're also building in a lot of the same machinery that does abstract reasoning," elaborates Stein. So, to pursue these beliefs and answer critics, says Brooks, "I had to go and ratchet things up beyond the insect level to a more sophisticated animal. The question was how far to go in one jump. The logical thing might be to go after an iguana or dog or something like a rat. On the other hand I figure I have one 10-year project left in me and I would hate to go out as 'he built the best artificial cat in the world.'"

Instead, recalls grad student Cynthia Ferrell, who named the robot and leads its construction, her boss decided to go for "the whole enchilada" and build Cog. As for Stein, whose past work focused on traditional AI symbol processing, she decided that constructing an embodied robot would be the best way to ground her more abstract research in reality. "What I'm interested in building is something that connects to the world but still does AI tasks," she says.

Placing Cog on a pedestal

A first glance at Cog doesn't bring to mind Star Trek's Commander Data (although his picture is pinned to what can loosely be termed the robot's chest) or any other well-known fictional android. The robot is only an upper body bolted to a metal base. Its mechanical hips and neck provide a range of motion like the human torso. Though looking a bit like a metallic Venus de Milo for the moment, Cog will soon have arms and two hands, each sporting three fingers and a rigid thumb. Cog won't be running around, however, as legs, wheels, and other forms of locomotion have been rejected as too challenging for the moment.

The robot's head will have sets of microphones and video cameras that mimic human hearing and sight—although Cog will only see in black and white. One crucial element, however, is still missing from Cog's head: the brain. "It doesn't do its computing onboard. It has its huge brain in the other room," explains Ferrell. A thick tangle of colored wires runs down Cog's back and through a hole in the wall, connecting to a sophisticated network of microprocessors roughly equivalent to 64 Mac IIs.

The human brain has often been compared to a massively parallel computer because it handles many diverse jobs simultaneously. Inspired by that notion, Cog's builders plan to divvy up the functions of the robot's sensors and motors among its processors. "Each processor is like a little piece of the brain doing a dedicated task," explains Brooks. Distinct groups of processors are already lined up to handle roughly the same functions as the visual and auditory cortexes; others will be dedicated to the role of the cerebellum, which translates motor commands into smooth physical actions. The Cog team members are not trying to directly model the complexity of a mature human brain, but they are betting that their crude cerebrum is intricate enough to mimic the mental capacity of an infant.

Driving all this wiring is Cog's software, which will enable the robot to gather, interpret, and store the torrent of information expected to flow in. This programming will determine how Cog "learns" from sensory experience by creating associations among sounds, images, and contacts. For this task, Cog's builders will draw on dozens of learning techniques, such as neural networks that simulate the behavior of neurons in the brain by strengthening connections among the elements when they produce a desired behavior. "There isn't a single answer to how Cog will learn," says Stein.

Bringing up baby

Initially, Cog's creators will be preoccupied with merely getting the robot to control its eyes, ears, and arms. When first turned on, for instance, Cog will "stretch" to calibrate the range of motions its motors can perform. Tougher challenges loom ahead for the robot. For example, babies at an early stage of development can bat a static object and later learn to intercept moving ones. Infants also learn to react visually and physically to auditory stimuli, first shifting their eyes and then turning their heads when a sound attracts their attention. Cog, it is hoped, will even reproduce the "cocktail party effect," where an individual can interact with a nearby person, but at the same time pay close attention to a distant conversation. The Cog team can monitor the robot's visual and auditory processing while this is happening, but the clearest measure of success may be if the robot's interactions with people and its environment match those of a living being.

Cog will be "born" with certain base-level behaviors. Dennett says that Cog will have a recognizable "mother" (identifiable through face-recognition programs), and it will be instructed to want to interact with that person. This "bonding," Dennett says, may make the robot try to attract its mother's attention with waving arms or other actions, much as infants try to prolong contact with adults by crying or making noise. And as thought to be the case in babies, cross-cultural cues, such as smiles, frowns, soft encouraging tones, and nodding heads, will also be hard-wired, providing Cog with a crude ability to interpret the reactions of people around it. Those cues may assist Cog in reaching another goal, the development of "shared visual space," which is when two individuals deliberately focus their eyes on the same object—even without verbal cues.

If Cog masters all these ways of interacting with and sensing its world, the robot's builders hope it will also reignite the basic debate about what intelligence is, parents argue. In fact, in "Building Brains for Bodies," their 1993 AI lab memo describing the project, Brooks and Stein confidently wrote that Cog "will learn to 'think' by building on its bodily experiences to accomplish progressively more abstract tasks."

That's not, of course, a universally accepted notion. Many in AI still feel that human intelligence is largely divorced from embodiment, and this hypothesis is sometimes labeled the "brain-in-the-box" approach. But how might Cog weigh in on this issue? One potential test, suggests Stein, may revolve around the concept of transitivity. The statement, "If A is greater than B, and B is greater than C, then A is greater than C," is an example of transitive reasoning. The embodiment hypothesis would argue that as Cog physically manipulates different-sized blocks or other objects, placing one on top of another, the robot could develop an aptitude for comparisons that could be bootstrapped into more abstract

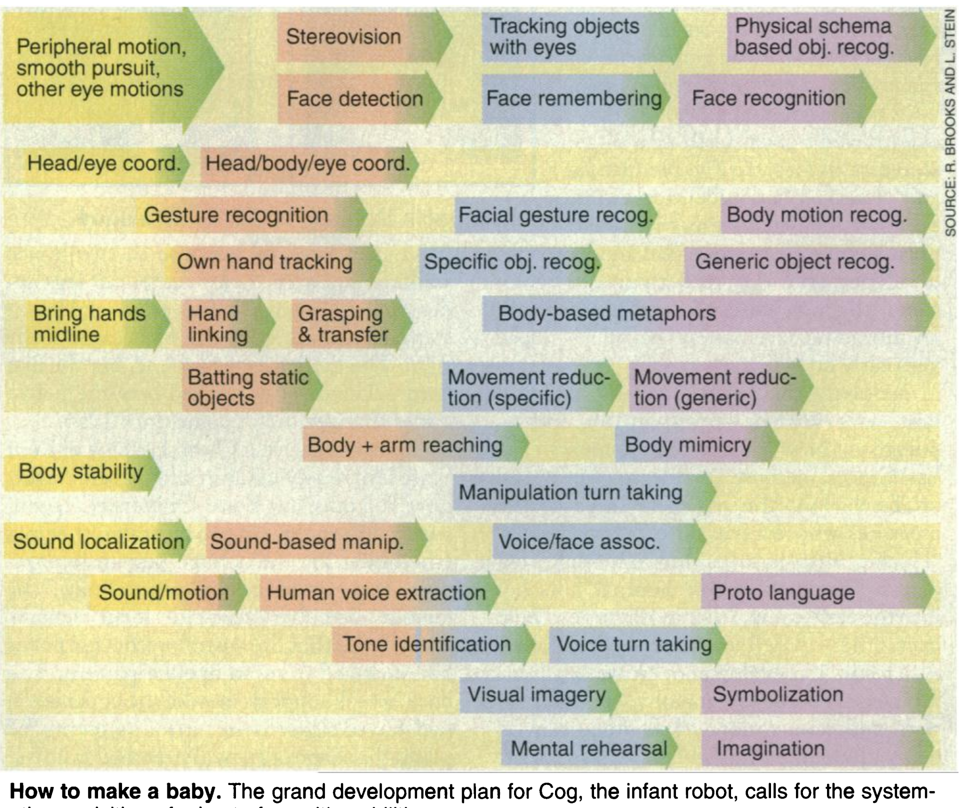

How to make a baby. The grand development plan for Cog, the infant robot, calls for the systematic acquisition of a host of cognitive abilities.

transitive reasoning; the researchers might then test for that with psychological exams designed for preverbal children.

Hype and bust?

Some admire Brooks and his colleagues for trying to shake up the community, as he did with the insect robots. "I think he wants to knock people on the side of the head again," says Grupen. Jill Lehman, an AI researcher at Carnegie Mellon University in Pittsburgh, suggests Cog is the perfect response to those who argue that researchers must integrate their successes in vision, hearing, and other relevant fields. "Someone has to try this. It's time now to stop working on the tiny little pieces and put them together," she says.

Yet others are already chalking Cog up as another blip in the disturbing "hype and bust" cycle of AI research, where overeager scientists promise to solve the mysteries of intelligence in order to attract funding for elaborate projects—projects that, inevitably, fail to deliver. "It's sexier if you build a robot, but it's not clear it's science. AI should be more disciplined, more issue-oriented, more patient," says Steven Pinker, a linguistics expert in MIT's cognitive science department. "There's so little known about the early stages of cognition that it's kind of silly to spend hundreds of thousands a year to simulate what we don't know. It's a waste of time," adds University of Rochester's Thomas Bever, an editor of the international journal *Cognition*.

While the Cog team shrugs off such criticism, Bever's mention of money does raise an issue that may ultimately stunt Cog's development. Until now, the project has largely been financed through a nest egg of unrestricted grants that Brooks had built up, but that reservoir will not last forever, he says. A first try with a large grant proposal at the National Science Foundation failed, they say, despite encouraging reviews. "The one big thing against us, in the current funding environment, is that this is not an application-driven program," explains Brooks, decrying what he calls the narrowing vision of U.S. funding agencies. "I believe Cog will have practical spin-offs, but Cog is not about practical spin-offs. Cog is about basic research with long-term strategic goals," adds an even more frustrated Stein.

Like worried parents, the two want the resources to bring their child up right. And like realistic parents, they expect Cog's life to include failure along with success. "We're overreaching; I'm perfectly willing to admit that," says Brooks. "We will fail in many dimensions," he adds, "but I think there's enough there to succeed a little bit in some dimensions." He and Stein just hope that Cog quickly learns from its mistakes—and in the process, educates its own parents about where they went wrong in raising it.

–John Travis

ANTHROPOLOGY

Putting a New Spin on the Birth of Human Birth

Humans do any number of things better than other animals, but giving birth is not one of them. Among the apes, our closest relatives, females bring infants into the world through a roomy birth canal with little fuss. In contrast, human babies often spend hours corkscrewing their way down a narrow birth canal, finally emerging head down, away from the mother—the only primates to do so. That makes human birth a risky business. Because babies don't bend backward, mothers can't pull them out without risk of serious injury, nor can they clear their newborns' airways if they are in trouble. Says University of Delaware anthropologist Karen Rosenberg wryly, "it's not the type of system you would invent if you were designing it today."

The process wasn't invented today, of course; it evolved over millions of years. But exactly when and how it did so has for decades perplexed anthropologists, who lacked the fossil evidence that could answer those questions. Over the past 8 years, however, they have been able to reconstruct a few bones associated with the birth canal from human ancestors dating as far back as 3 million years. At a symposium at last month's meeting of the American Association of Physical Anthropologists in Denver, researchers used those bones to begin tracing the evolution of human birth.

Those reconstructed bones, however, haven't given birth to a single scenario. In fact, they've produced a pair of decidedly nonidentical twins: Some researchers at the symposium presented new evidence that modern birth developed only very recently in our evolutionary history, while others countered with an intriguing speculation that it developed very early on.

Resolving this "when" question is important, says Wenda Trevathan, an anthropologist at New Mexico State University in Las Cruces, because of the insights it can produce about the social abilities of the creatures who evolved into human beings. "Human birth is so painful and risky," explains Trevathan, whose analysis puts her into the early camp, "that mothers need help from others to deliver a baby successfully." As a result, its development created a powerful selective force for empathy, communication, and cooperation—skills important to being human. And when those traits emerged is another date that anthropologists would love to pin down.

One thing researchers do know is why modern birth gives women such a difficult turn. The human pelvis, which surrounds the birth canal, crimps that canal partway down. At the top, the canal is widest from side to side. The longest dimension of a baby's head is from the nose to the back of the skull, and so the baby enters the canal facing sideways. But lower down, the canal changes its shape so that the longest dimension is from front to back. As a result, the infant must rotate 90 degrees. And there's one more twist: The baby's head is broadest at the back, but the lower canal is a bit broader at the front. So the infant enters the world facing down.

Apes, which have small heads and relatively larger pelves, don't have to go through these gyrations. In 1960, anthropologist Sherwood Washburn of the University of California, Berkeley, speculated that humans took this turn for the worse because the lineage was caught in an "obstetrical dilemma." The pelvis narrowed in response to two-legged walking, since this helps center our legs under our bodies. But as babies' heads and brains started getting bigger, the fit became really tight. The theory seemed sound, but the first hard evidence of when and how the pelvic girdle changed didn't come until 1986.

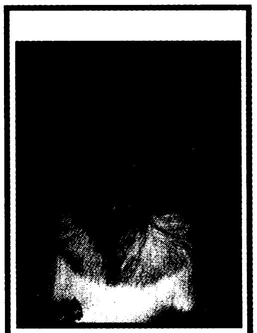

MICHAEL DICK/ANIMALS ANIMALS

Proboscis Monkey

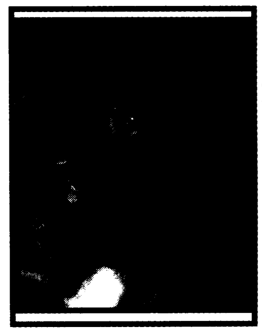

DOUG WECHSLER/ANIMALS ANIMALS

Macaque

That was the year Owen Lovejoy of Kent State University in Ohio and Robert Tague, now at Louisiana State University, reconstructed the pelves of two australopithecines (the oldest known nonape human forerunners), including one belonging to "Lucy," the famous 3-million-year-old fossil female. They found that the australopithecine pelvis had changed from an apelike pattern. The back, which supports most of the upper body, had moved closer to the hip joints, giving the pelvis the shape of an oval stretched from hip to hip. The change helped "to adjust posture in a biped," says Tague.