THE RACE TO BUILD INTELLIGENT MACHINES

Computers can't read the Britannica or talk about their feelings—but they're working on it

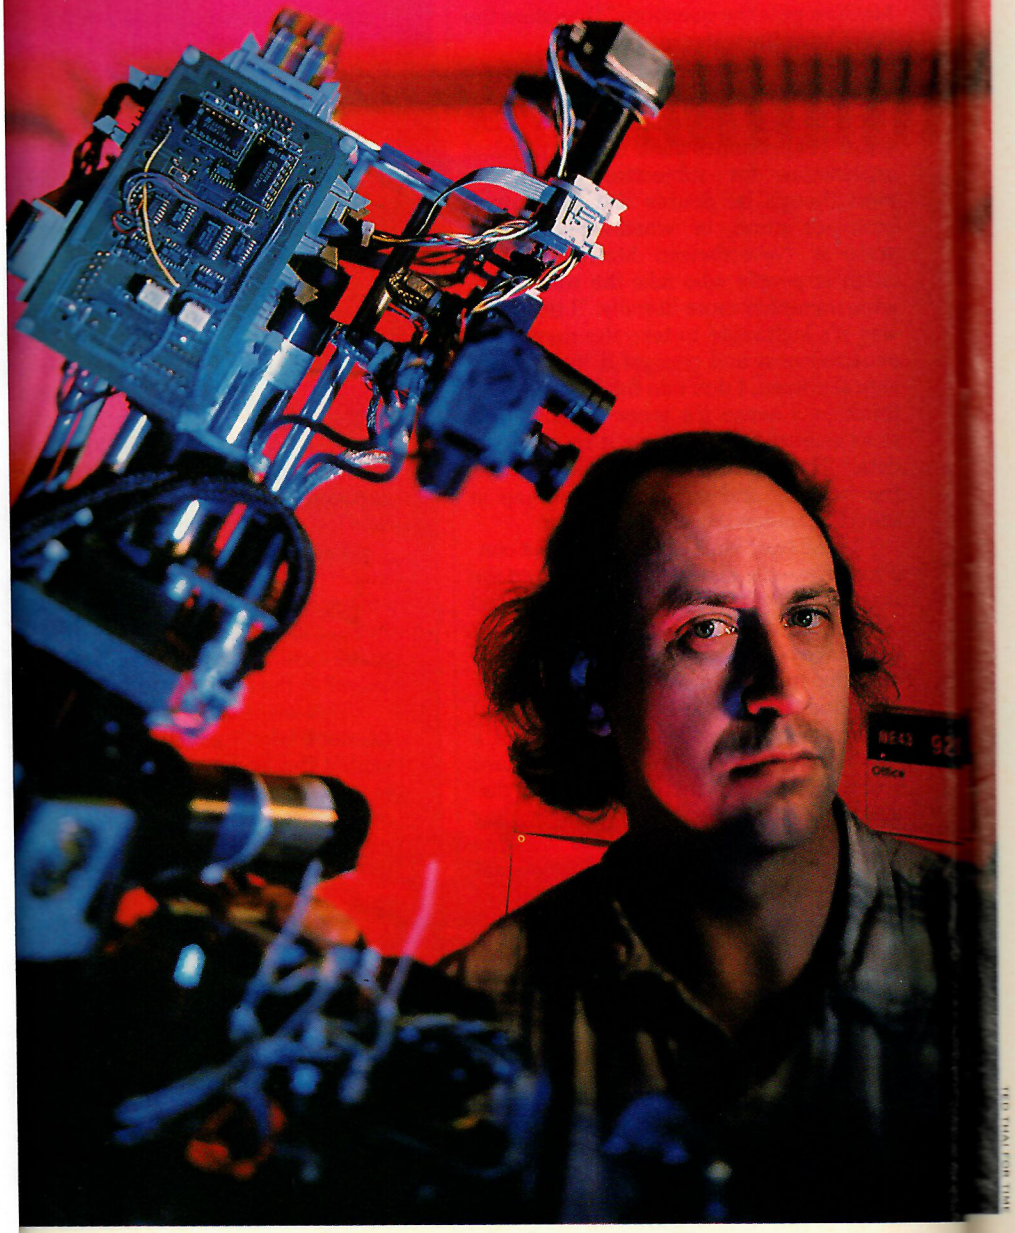

By JULIAN DIBBELL

RODNEY BROOKS GUSHES LIKE A first-time parent about the things his baby can do. "It sits there waving its arm around, watching its arm, reaching for things," he says. These are pretty standard tricks for newborn humans, of course, but then Brooks' "baby" (nicknamed Cog) isn't exactly human. It's a vaguely person-shaped concoction of metal, plastic and silicon, with cameras where its eyes should be and eight 32-bit microprocessors for a brain. Cog is an artificially intelligent computer that is trying to learn about the world the way babies do, programming and reprogramming itself through interactions with the people and objects around it. And Brooks, a professor at M.I.T., is as ambitious for his progeny as any father is for his child. He'd like to have a conversation with Cog someday about nothing in particular.

Douglas Lenat has similar hopes for his artificially intelligent brainchild—a sprawling, data-rich computer program called CYC (as in encyclopedic). But where Brooks expects Cog to teach itself about the world, Lenat is leaving nothing to chance. For more than a decade, his CYC development team in Austin, Texas, has been typing in the rules of "human consensus reality" (otherwise known as common sense) one thuddingly obvious assertion at a time. "Bread is a food," for example, or "You're wet when you sweat." CYC knows nearly a million of these rules now, and when it has another million or so under its belt, Lenat asserts, the program should be savvy enough to need no more spoon feeding. It will just swallow the encyclopedia whole and then ask questions about whatever it doesn't understand.

TWO APPROACHES, TOP DOWN AND BOTTOM UP, REPRESENT AN ALMOST THEOLOGICAL SCHISM IN COMPUTER SCIENCE

Compared with beating the world's greatest chess player at his own game, of course, infant-like flailing or knowing about wetness might not sound like much. But programs like CYC and Cog, not chess machines like Deep Blue, currently define the cutting edge of applied artificial intelligence—the 40-year effort to build machines that think. Ten years ago, when AI was as hot as the Internet is today, researchers raced to build programs that showed deep expertise in a narrow field of endeavor—like chess, for example, or medical diagnosis. These days, however, it's the promise of breadth, not depth, that inspires the artificial intelligentsia—and drives the programs that come closest to what the rest of us might regard as thinking.

Now, with the proliferation of cheap, powerful computers and the rapid growth of the Internet, there's new interest in all kinds of "intelligent" machinery—not just chess-playing supercomputers or grandiose AI research projects like CYC and Cog. The past few years have seen a burst of entrepreneurial activity in what are called intelligent agents—programs of rather more modest IQ that are nonetheless smart enough to be released on the Internet to do small, useful chores like tracking stock prices or digging for nuggets of research data.

But even with all this activity, no one has taken up the pursuit of digital intelligence with as much audacity or ambition as Lenat and Brooks. Their parallel quests to build what may be the world's first convincingly humanlike computer programs have been compared to the dramatic 1911 Amundsen-Scott race to the South Pole; but even that analogy falls short. For the rivalry between the two researchers is not merely personal (Brooks considered naming his robot Psych! just to get Lenat's goat) but deeply philosophical as well, straddling the almost theological schism that runs down the middle of contemporary AI.

On the side of orthodoxy stands Lenat. Though CYC's unabashedly more-is-more approach has raised eyebrows in the field, its design remains true to one of the central tenets of classical AI: symbolic knowledge makes the mind go round. In other words, if you can write down the logical structures through which we apprehend the world, you're halfway to re-creating intelligence. And if you can program what you've written into a machine, even better. Hence the 170 person-years CYC's handlers have devoted to codifying what any five-year-old already knows.

Cog, on the other hand, embodies the principles of AI's breakaway faction, the so-called bottom-up school. Inspired more by biological structures than by logical ones, the bottom-uppers don't bother trying to write down the rules of thought. Instead they try to conjure thought up by building lots of small, simple programs and encouraging them to interact. Earlier in his career, Brooks helped put this approach on the AI map by building tiny, insectlike robots—"bugbots"—that wandered around his laboratory without the benefit of any single guiding program. Cog's "mind," similarly, is just a collection of loosely coordinated digital reflexes scattered among its eight processors, with no one place to point to as the seat of intelligence. "There's no there there," notes Brooks with a touch of pride.

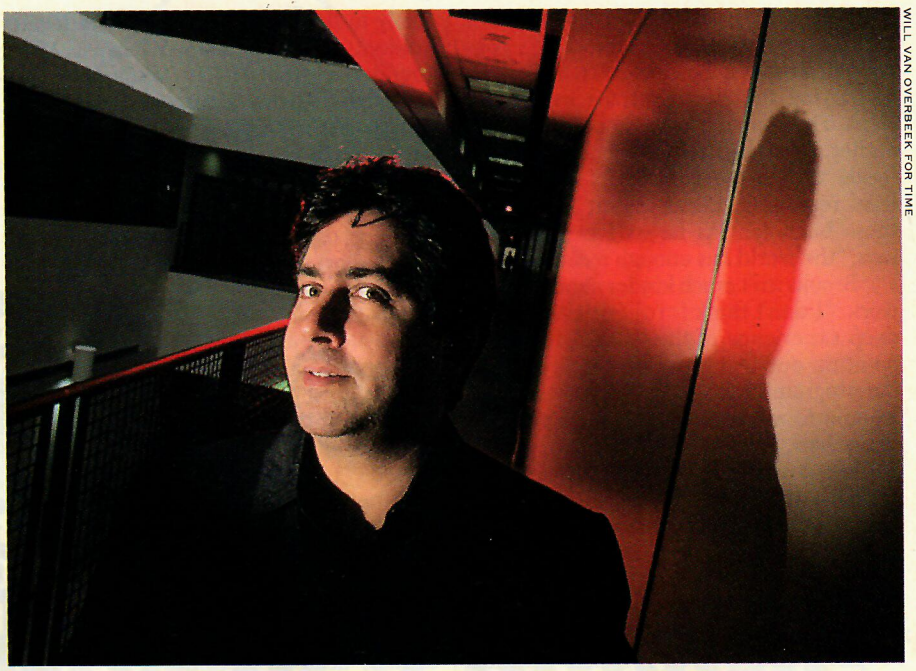

WILL VAN OVERBEEK FOR TIME

SPOON-FED KNOWLEDGE: Lenat's epic CYC project tackles one of the oldest unanswered questions in computing: Can the stuff of understanding be written down, one rule at a time?

TED THAI FOR TIME

APPLE OF HIS EYE: With eight Macintosh-type processors for a brain, Brooks' Cog tries to learn as infants do: by trial and error

Lenat might say the same, though for less flattering reasons. As far as he's concerned, Brooks is headed down a blind alley, trying futilely to retrace evolution's steps. The simple tricks that made Brooks' earlier, insectlike robots work will never "scale up" to the complexity of human-level reasoning, Lenat insists. And as for those bugbots? "As far as I can tell," he says, "the world has enough insects."

Brooks isn't much kinder about Lenat's work. "I don't think [CYC] can ever have a deep experience of the world," he sniffs, pointing out that without sensory input, the program's knowledge can never really amount to more than an abstract network of symbols.

Lenat may have the last word—at least for now. With its 10-year head start over Cog, the CYC project is much closer to spinning off practical applications, and its timing couldn't be better. The World Wide Web's chaotic infobloom is starting to strain the limits of today's popular but simpleminded search engines (which work, for the most part, by matching up key words). But CYC, with its ability to make commonsensical leaps of logic, can connect a request, say, for pictures of "happy people" with the caption, "A man watching his daughter learn to walk."

There might be a market for inferential power like that, especially if it could be yoked to some of the autonomous semi-intelligent agents that are already buzzing around on the Internet. But Lenat isn't the only researcher poised to bring real intelligence to intelligent agents. If anybody is likely to beat him to market, it's a former student of Brooks' named Patricia Maes, founder of the M.I.T. MediaLab's Autonomous Agents group and of Agents, Inc., based in Cambridge, Massachusetts. Maes and her students have devised imaginative ways to use bottom-up AI to personalize information delivery—"your news" as opposed to "the news."

Perhaps the most fascinating of these applications—and certainly the most accessible to Netheads—is Firefly (*http://www.ffly.com*), Agents' flashy new "music-recommendation system." Firefly lets you rate records, tapes and CDs and then pools those ratings to create a "map" of your musical tastes. With thousands of users' tastes pinpointed on that same map, it becomes simple to extract dead-on recommendations from the ratings lists of your closest "neighbors." From the mingling of large numbers of relatively simple pieces of information, in other words, an uncanny sort of acumen emerges, able to make suggestions even your best friend might miss.

ALREADY, PRIMITIVE INTELLIGENT AGENTS ARE BUZZING AROUND ON THE INTERNET

More intriguing, though, in the long run, are Maes' experiments with genetic programming—a software version of survival of the fittest. In one program designed to help users sort their E-mail, Maes let slightly different agents compete for the user's approval; the most favored were allowed to "mate" and pass their "genetic code" on to the next generation. More than just a whimsical mimicry of living processes, the tactic may actually be, in Maes' view, the only way for agents to keep pace with people's rapidly changing tastes in a network like the Internet. "In my vision of the ultimate software agent system," she says, "there are all these 'life' forms evolving by themselves and specializing toward whatever you happen to be interested in."

The idea of harnessing Darwinian evolution to help humans do their programming is gaining currency, especially in the biologically oriented milieu in which Brooks and Maes operate. Farsighted proponents of this school imagine huge populations of digital agents meeting and mating in increasingly complex global networks—creating in their progeny artificial intelligences that exceed even the descendants of CYC and Cog.

We may meet them sooner than even these visionaries imagine. Within the next few weeks, in Kyoto, Japan, an ecobiologist and radically bottom-up computer theorist named Tom Ray will initiate an open-ended experiment he calls a "digital biodiversity reserve." A single, tiny, self-reproducing program will be loosed into a "virtual Internet" spread among hundreds of computers around the world. If all goes as earlier trial runs suggest, Ray's artificial "organisms" will quickly populate the network and begin to evolve.

After that, even Ray doesn't know what will happen. Perhaps the population will reach stasis and stagnate at the level of pond scum. Or perhaps Ray's digital beings will set off down the same sort of evolutionary path our species has traveled, only at electron speed. And if that happens, what then? We may find ourselves face to face with an artificial intelligence so thoroughly immersed in the silicon realm, so distant from our curious, carbon-based concerns, that we cannot even hope to converse with it.

That isn't what AI set out to achieve 40 years ago. But artificial intelligence tends to be a moving target; once a computer can do something—like play tic-tac-toe or even chess—that behavior no longer seems to require much thought. If anything unites AI's increasingly diverse strategies for building machines that think, however, it's that they all require us to stretch our notions of what human thought really is. "We need to be prepared," warns Ray, "for an intelligence that is very different from our own." ■

TED THAI FOR TIME

SURVIVAL OF THE SMARTEST: M.I.T.'s Maes uses Darwinian programming to hone her software agents; those that do their job best get to pass their code on to the next generation