

REWINDING EVOLUTION

What recreating dinosaurs will tell us about the future of life

NewScientist

WEEKLY August 20-26, 2011

INSIDE ANIMAL MINDS

How other creatures
see our world

DIGITAL DETECTIVES

Facial recognition
and England's riots

BRAIN GAIN

Make the most of your
mental machinery

ULTIMATE DIAGNOSIS

Taking the guesswork
out of medicine

BIG BANG BLACK HOLES

They can solve the dark matter mystery

Science and technology news www.newscientist.com Focus on chemistry

No 2826 US\$5.95 CAN\$5.95



When it comes to artificial intelligence, the brain isn't everything, says Virginia Hughes

Body conscious

In January, Max Versace and Heather Ames were busy with two newborns: their son Gabriel and Animat, a virtual rat.

Like all babies, when Gabriel was born his brain allowed him to do only simple things like grasp, suck and see blurry images of his parents. The rest was up to him. From the first day his body experienced the world, his senses began to respond. He learned to follow a moving object with his eyes, tell red from yellow, and reach for his mother. Over the next couple of years, he will learn to crawl, walk, talk and, eventually, look after himself.

With any luck, Animat's development will follow a similar path. It didn't start with much programming, either. But Animat's interaction with its virtual world has already taught it how to tell colours apart and understand the space around it. As it develops, it will use its senses to learn even more.

Animat's "parents", researchers at Boston University, are trying to build intelligent machines based on the smartest machine we know of: the brain. But instead of focusing on programming the brain itself, they are taking a cue from biology. Like every human baby, and unlike the vast majority of engineered intelligence, the development of Animat's intelligence will depend on the way its body senses the world. They hope this approach will advance machine intelligence to the point that robots start to think in a more human way.

That goal has proved a formidable challenge. Take Watson, IBM's supercomputer which in February stole the crown from two human champions on the US quiz show *Jeopardy!* To do so, it used sophisticated successors of the first AI algorithms to tap massive databases stored on 90 servers. All that machinery required 80 kilowatts of

energy to run, enough to power a small village. Our grey matter is spartan in comparison. "A brain consumes less power than a light bulb, and occupies less volume than a 2-litre bottle of soda," says Dharmendra Modha, a computer scientist at IBM.

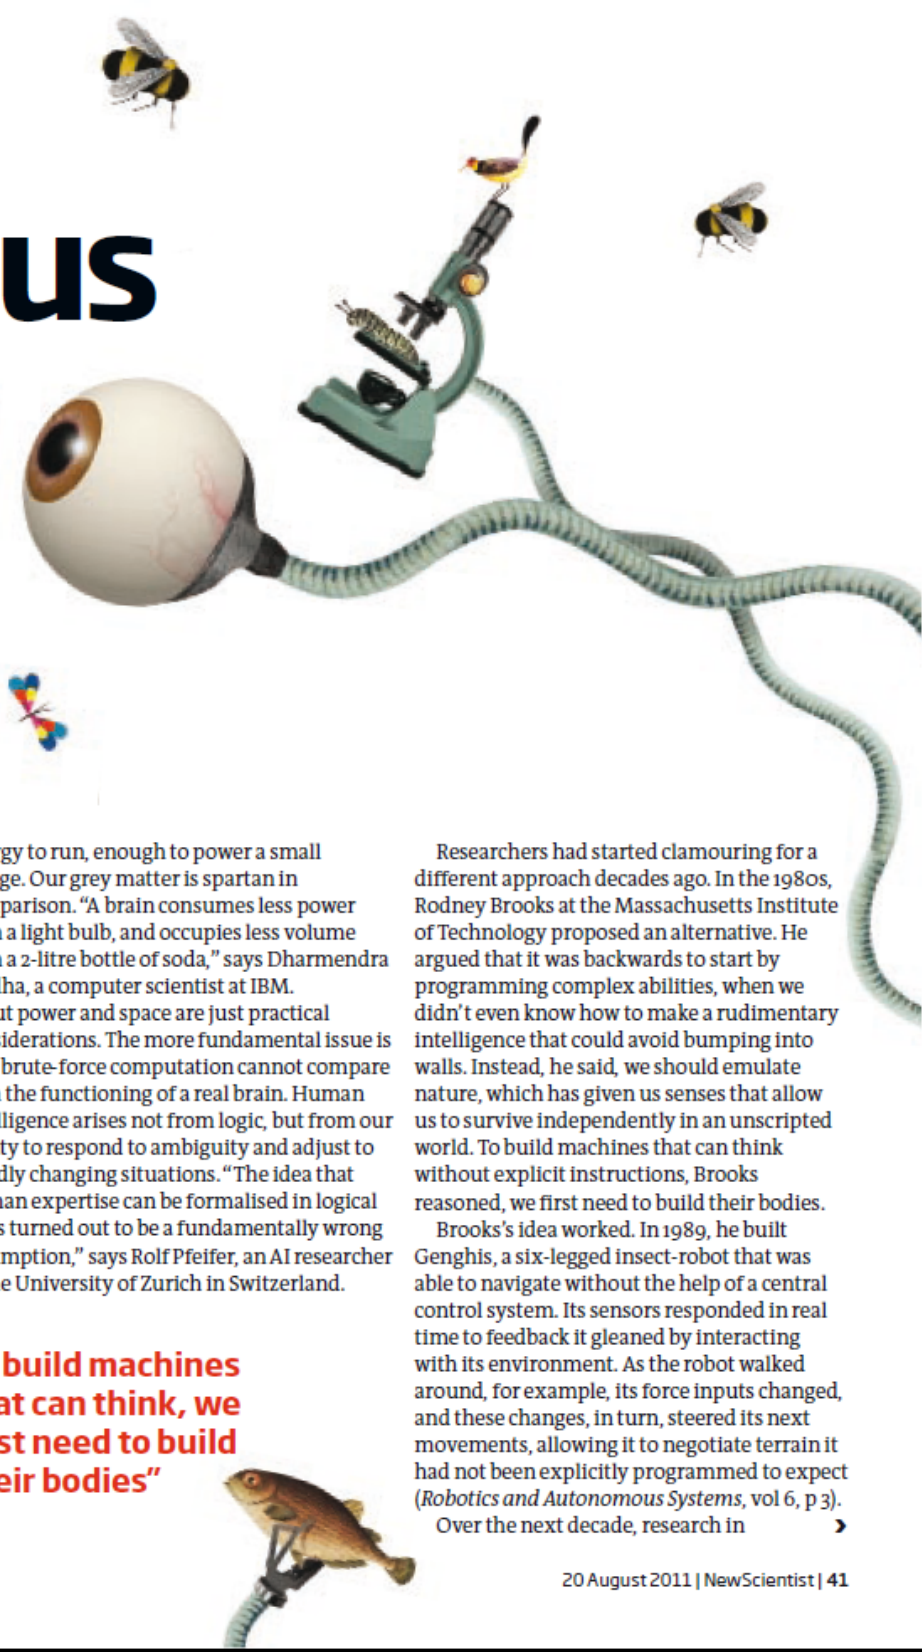
But power and space are just practical considerations. The more fundamental issue is that brute-force computation cannot compare with the functioning of a real brain. Human intelligence arises not from logic, but from our ability to respond to ambiguity and adjust to rapidly changing situations. "The idea that human expertise can be formalised in logical rules turned out to be a fundamentally wrong assumption," says Rolf Pfeifer, an AI researcher at the University of Zurich in Switzerland.

"To build machines that can think, we first need to build their bodies"

Researchers had started clamouring for a different approach decades ago. In the 1980s, Rodney Brooks at the Massachusetts Institute of Technology proposed an alternative. He argued that it was backwards to start by programming complex abilities, when we didn't even know how to make a rudimentary intelligence that could avoid bumping into walls. Instead, he said, we should emulate nature, which has given us senses that allow us to survive independently in an unscripted world. To build machines that can think without explicit instructions, Brooks reasoned, we first need to build their bodies.

Brooks's idea worked. In 1989, he built *Genghis*, a six-legged insect-robot that was able to navigate without the help of a central control system. Its sensors responded in real time to feedback it gleaned by interacting with its environment. As the robot walked around, for example, its force inputs changed, and these changes, in turn, steered its next movements, allowing it to negotiate terrain it had not been explicitly programmed to expect (*Robotics and Autonomous Systems*, vol 6, p 3).

Over the next decade, research in





neurobiology, cognitive science and philosophy suggested that Brooks's ideas applied much more broadly. In the late 1990s, George Lakoff, a cognitive scientist at the University of California, Berkeley, proposed that human intelligence, too, is inextricably linked to the way our bodies and senses interact with the environment. According to Lakoff and his supporters, our "embodied mind" explains not only rudimentary intelligence, such as how we learn to visually recognise objects, but even complicated, abstract thought (*Artificial Intelligence*, vol 149, p 91). Here, at last, was the key to building a sophisticated, human-like intelligence.

There was just one problem: embodied AI is difficult to upgrade. Improving a robot's sensor-laden body requires not just programming extra functions but the painstaking disassembly and reassembly of the sensors themselves. Despite these obstacles, a few researchers found the idea too compelling to abandon. In 2009, Owen Holland at the University of Sussex in Brighton, UK, created a humanoid robot modelled on some of the principles that had led to Genghis, called Eccerobot. However, Eccerobot is still incapable of even basic movements, and has exhibited no sign of intelligence. And so, even as advances in computing power have supercharged conventional AI, embodied AI limps on in ever-smaller circles.

Then, a couple of years ago, Versace and his team began to realise that there was hope for embodiment after all – if you skipped the physical body. Thanks to powerful new graphics cards, video game designers can simulate anything, including a robot's body, the environment it lives in, and even the complex physics underlying the interactions

between the two. "Graphics cards are starting to change the game of what questions you can pose," says Ennio Mingolla, a cognitive scientist on the Animat team.

Accelerated learning

He and the rest of the team seized on these advances to "cheat" at embodiment. Instead of toiling over a real body, they built a virtual one, whose synthetic sensors would interact with a painstakingly rendered virtual environment (*IEEE Computer*, vol 44, p 21). This way, they reasoned, they could reap all the advantages of embodied AI with none of the drawbacks. If it worked, they would be able to hit fast-forward on the evolution of an embodied intelligence.

Animat was born on 11 January. That was the day Versace's team hooked up its brain, which is made up of hundreds of neural models – for colour vision, motor function and anxiety, among others – all of which are faithful

imitations of biology. That means they don't contain a list of explicit commands, just as Gabriel's brain doesn't calculate his cot dimensions to figure out where to reach for a toy. "Everything is designed as closely as possible to its biological counterpart," Versace says. "There are no rules, per se."

So, like Genghis, Animat depends on the feedback it gets from its virtual "body", which is equipped with the kinds of sensors found in skin and the retina, to learn and move. Unlike Genghis, however, every bit of Animat can be upgraded in the blink of an eye.

Its environment also obeys the laws

BETTENDORF



Swim to safety

The virtual Animat is programmed to have a real rat's colour vision, navigation abilities and hatred of water. The only way out of the water is to find the hidden platform by the white pole



Experiment 1



Animat swims around randomly and does not find the platform

Experiment 2



Animat swims in a different pattern and eventually finds the platform

Experiment 3



Animat can now find the platform faster, guided by the colours of the poles

Experiment 4



On the fourth attempt, Animat swims directly to the white pole to find the platform

of real-world physics, including gravity, giving Animat realistic sensory information. Light hits its virtual retina, for example, giving it colour vision, and correctly calibrated forces – for things like water and air pressure – are exerted against its simulated skin. Different combinations of such inputs drive Animat's reactions.

Animat's virtual world is a giant blue swimming pool surrounded by many poles, all different colours (see diagram). Like real rats, Animat hates water, courtesy of an anxious streak Mingolla included in its neural models. The only way to escape the water is to find a small platform hidden somewhere below the water's surface. The test of Animat's intelligence is how quickly it learns to find that anxiety-alleviating platform.

The first experiment looked like a failure: Animat swam frantically in random patterns for an hour before the researchers gave up and ended the test. When they dropped it into the pool a second time, however, Animat's swimming pattern changed. This time, after 45 minutes of swimming in a new pattern, it stumbled onto the platform. Its anxiety levels decreased sharply when it left the water, and that reward strengthened the connections that led it there. It now knew the colour of the poles near the platform, for example, and the approximate path it took to get there.

And sure enough, the third time it was chunked into the water, Animat spent far less time swimming before it found the platform, because it looked for the right coloured poles. On the fourth try, it didn't even hesitate: it swam directly to the platform. "You could see the intelligence of the animal right away," Versace says. "That was the Eureka moment."

Within a few years, Versace's team plans to create more Animats with expanded

capabilities. As well as visually recognising objects, these creatures will hear and even communicate with each other. The optimistic time frame wouldn't be possible without the fast track virtual bodies provide.

Brain transplant

Auspicious as these early experiments may be, however, the virtual world is just a rehearsal space. The real test will come once Animat's virtual body is swapped for a real one. After all, the ultimate goal is a robot that can move independently in the real world.

The advantage of such a machine over a traditionally built AI is that it would not require every possible scenario to be anticipated in the programming. That would save millions of lines of instructions. This means less hardware, and less hardware means less power is needed to run it all. And that, ultimately, is the point of embodied cognition: a machine that can move, make its own decisions, and survive in the real world on very little juice – just like we do.

The possibilities for such machine intelligence explain why, earlier this year, NASA came calling. A Mars rover with

"The virtual world is just a rehearsal space. The real test will come once Animat's virtual body is swapped for a real one"



biological intelligence could learn to use its neural networks for vision, to balance itself, and to navigate away from rough terrain, eliminating the need for constant human supervision. "Such a rover would also be able to plan its route, explore, sense its battery level and safely return to base," Versace says. The team is already preparing a virtual Mars for Animat, complete with craters.

Because Animat is designed to learn like its biological counterparts, some familiar questions crop up. Can Animat feel pain? After all, Animat gets negative reinforcement, in the form of intense anxiety, and positive reinforcement in the form of instant relief when it reaches the hidden platform. "It basically starts to build these representations of good and evil," Versace says. So when should he start worrying about animal cruelty?

The point might not be as silly as it sounds. Feeling could be a crucial bridge between intelligence and consciousness. Some cognitive scientists think that basic mechanisms of reinforcement, such as anxiety and relief, are exactly how human consciousness arises. The "feel" of consciousness – the internal experience of seeing red or feeling pain – derives not from higher cognition, but from simple interactions with the world. "Feel is not something that's generated by our brains, but something that we do, a way of interacting with the world," says Kevin O'Regan of the Institute of Neuroscience and Cognition at Paris Descartes University in France.

Neither Versace nor Ames believe Animat will develop consciousness. But they don't mind: they have their hands full with the embodied biological intelligence they created the old-fashioned way. ⁿ

Virginia Hughes is a science writer based in New York