The robot that thinks like you...

Click to Print

05 November 2005
From New Scientist Print Edition. Subscribe and get 4 free issues.

Douglas Fox

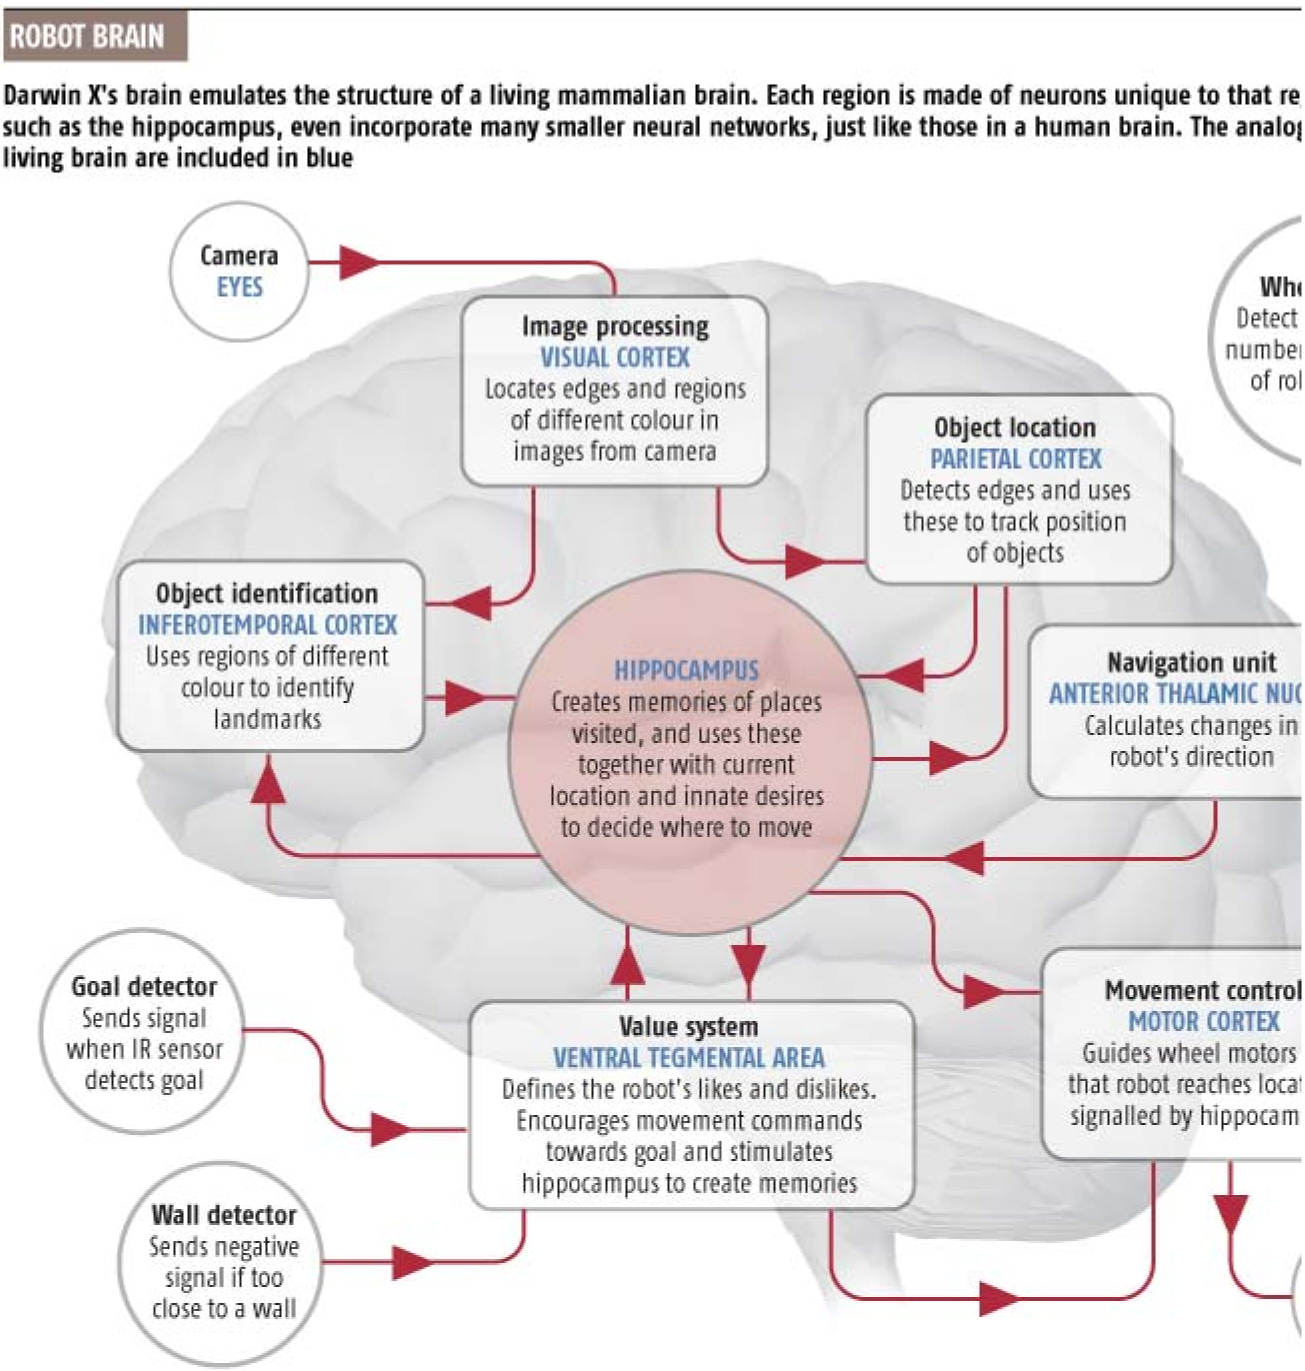

Robot brain

THE infant crawls across a floor strewn with blocks, grabbing and tasting as it goes, its malleable mind impressionable and hungry to learn. Before my eyes it is already adapting, discovering that the striped blocks are yummy and the spotted ones taste bad.

Its exploration is driven by instincts: an interest in bright objects, a predilection for tasting things, and an innate notion of what tastes good. This, after all, is how babies explore the world and discover that pink, perky objects exist, and that they produce milk. Hands-on exploration moulds their billions of untrained brain cells into a fully functioning brain.

The infant I am watching wander around its rather spartan playpen in the Neurosciences Institute (NSI) in La Jolla, California, is a more limited creature. It is a trashcan-shaped robot called Darwin VII, and it has just 20,000 brain cells. Despite this, it has managed to master the abilities of a 18-month-old baby - a pretty impressive feat for a machine.

The key to Darwin's abilities is its brain. This is an amalgam of rat and ape brains, encoded in a computer program that controls its actions. Darwin tastes blocks by grabbing them with its metal jaws to see if they produce electricity. It likes the ones that do and dislikes the ones that don't. Within half an hour of being switched on it learned to find the tasty blocks.

Darwin VII is the fourth in a series of robots that Jeff Krichmar and his colleagues at NSI have created in a quest to better understand how our own brains work - the first three versions of Darwin did not have a real robotic body to control. Darwin VII allows Krichmar to record changes in hundreds of thousands of its brain's neural connections as it explores and learns, to test neuroscientists' theories of how real brains work. "This is something that you can't do in a real brain," Krichmar points out.

If Krichmar and others like him succeed, robots like Darwin might one day be seen as the ancestors of something much bigger. Some researchers, and even the US Defense Advanced Research Projects Agency, are gambling that robots like Darwin will be the forebears of an entirely new approach to artificial intelligence (AI): building intelligent machines by copying the structures of living brains. Some groups are even designing microchips that could eventually be used to build anatomically realistic artificial silicon brains to replace the computers that power existing robots like Darwin.

The dream is that these new brains, embedded in robotic bodies of silicon and steel, will go to a level beyond today's artificial intelligence systems. By sensing their environments as they explore and learn, they will develop the ability to survive in the constantly changing real world of imperfect information that we navigate so effortlessly, but which computers have yet to master. They will learn to do anything from mundane household chores we'd rather not do, to driving the kids to school, and even autonomously explore Mars or run nuclear waste facilities, all without human intervention. "All you would have to do is teach them," says Juyang Weng, a developmental roboticist at Michigan State University in East Lansing.

These systems will arise, say the researchers, by emulating the brain's neurons and the way they are connected to each other. In animal brains neurons are linked to form huge reconfigurable networks that behave like filters, transferring, modifying or blocking signals that they receive. Though living brains have been studied for decades, we still don't know exactly how they achieve the amazing abilities of the human mind.

That hasn't stopped computer scientists from trying to imitate them. The idea of an artificial neural network that could perform computations was proposed as long ago as 1943, by Warren McCullough and Walter Pitts at the University of Illinois. In the decades since, efforts to create intelligent machines from these networks progressed in relative isolation from the study of how real brains work. As a result, these artificial neural networks bear little relation to the structure of neural networks in real brains. But in the past few years, neuroscientists and AI researchers have started collaborating more closely, and their labours are beginning to bear fruit. Their conclusions challenge two decades of research into artificial neural networks.

It all boils down to this: existing artificial neural networks, such as those used in many computer systems today, are totally inadequate for creating anything resembling animal, let alone human, intelligence. To do that, you have be as faithful as possible to the real thing. And for the first time that's what several groups around the world are trying to do: emulate both the structure and the function of living brains in detail.

In all neural networks, both artificial and real, structure and function are intimately linked. The pattern of connections between neurons determines how well the network performs a particular task. If you train an artificial neural network to recognise abnormal cells in smears test, for example, it adapts by adjusting connections between individual neurons until external feedback indicates to the network it is doing the job well. But unlike the human brain, these systems are optimised to perform a single task. "It is a small part of what might be happening in the brain, a tiny portion of an intelligent action," says

Igor Aleksander of Imperial College London.

To get the adaptive, flexible behaviour you see in animals, you need to imitate the design of a whole brain, the body it lives in and the drives that motivate it, Krichmar says. "A brain-based device provides them all; a traditional neural net simply doesn't."

Neuroscientists have identified hundreds of different neural areas within mammalian brains. In effect each is a specialised neural network unto itself. It is only when you recreate these areas and start interconnecting the different modules that complex behaviour emerges that no single part of the system could achieve on its own, Aleksander says.

Back in Krichmar's lab, I watch as Darwin VII approaches another block. On a computer screen in front of me a map of its brain lights up. Groups of its simulated neurons that recognise patterns light up to identify the stripes on the block. Darwin has 18 simulated neural areas, including a lower visual area that detects edges and a motor area that directs movement. Its brain even simulates a virtual squirt of the pleasure-inducing neurotransmitter dopamine whenever it tastes a yummy block. This positive reinforcement encourages Darwin to repeat actions that lead to it finding those tasty blocks in the first place. I watch the result of that trained response as Darwin recognises a striped block. The mere sight of a block causes Darwin's motor area to light up in anticipation, triggering a reflex response. It lunges forward to grab it.

It's a neat demonstration, but since Darwin's brain is buried in a computer simulation, I wonder how much it really does resemble a real brain. To show me, Krichmar pulls up a computer file describing Darwin's brain. "We choose these values based on what we know of animal neurons in these brain areas," he says. These numbers represent a model of a brain based on the neural activity measured in laboratory experiments on rat and ape brains. It describes the topology and types of connection between individual neurons and the neural regions of Darwin's brain.

While the structure of traditional artificial neural networks bears little resemblance to living brains, the details of Darwin's brain are grounded in reality, Krichmar says. He and his colleagues laboriously assembled these details by scouring published research. From this they extracted a wide range of data, including functional MRI scans of brains showing points of activity in response to different sights or smells, and neural connections mapped by painstakingly injecting single nerve cells with dye and then, by microscope, tracing the cell's hundreds of branches that lit up.

A different kind of hardware

Darwin is a work in progress. The biological data and computing power necessary to build such a machine are only now becoming available. Huge gaps remain in our understanding of the human brain, so a team lead by Olaf Sporns, a neuroscientist at Indiana University, Bloomington, has proposed a project inspired by the Human Genome Project to map the neural connections throughout the human brain (*PLoS Computational Biology*, vol 1, p 42).

But it is going to take more than just simulating neural networks in software to make significant progress towards genuine new forms of artificial intelligence. Brain-based systems run very slowly on computers because brains and computers work in fundamentally different ways. Conventional computers funnel their calculations through one or a few processors at best, whereas mammalian brains distribute calculations across billions of neurons that operate in parallel. To get a significant improvement in speed, and therefore capability, new hardware will be needed that can imitate the way brains compute.

To this end, some researchers have begun developing silicon devices that imitate the behaviour of real neurons. Their processing units behave like neurons in that they respond to inputs of different value with a range of output values, rather than just switching on or off as in conventional computers. The chips can even change the interconnection between processors in real time, something that is impossible with existing microchips.

Even without these new computing machines, brain-based devices have been making steady progress. Each of the four robots NSI has built so far is designed to test a specific aspect of brain function. Using results from experiments with Darwin VII, Krichmar showed that different robots learning to recognise the same object may recruit entirely different sets of neurons to do so. In this respect, the robot's mind is a faithful approximation to a living brain.

Using his most sophisticated robot to date, Darwin X, Krichmar has been studying how brains use landmarks to navigate. Starting from different points in a room each time, the robot's task was to find a hidden floor marker, using stripes painted on the walls as landmarks. As the robot learned the layout of the room, it developed "place neurons" in its virtual hippocampus that helped it home in on the marker three times as fast on the 10th attempt as it did on the first (*Proceedings of the National*

Academy of Sciences, vol 6, p 2111). The place cells had not been implicitly encoded Krichmar says. "They just emerged from the model." He says the same thing can be seen in the hippocampal cells of a rat.

Weng has seen similar results with a robot named SAIL (Self-Organising Autonomous Incremental Learner). It operates in a slightly different way to the Darwin robots, focusing less on copying brain anatomy and more on mimicking how connections between individual neurons organise themselves as a brain matures.

When SAIL is switched on for the first time, Weng takes it for a stroll, leading it through the corridors of the building. After a day of this kind of guided activity, he lets SAIL navigate under looser supervision. If the robot hits a wall it gets negative feedback signals from bumper sensors, if it gets to the end of the corridor without hitting a wall, Weng presses a button on the robot that "rewards" it.

To begin a new experiment with SAIL, such as having it learn to navigate outdoors or associate spoken words with specific objects, Weng can wipe the robot's mind clean. For its first baby steps after each reboot, SAIL's visual cortex is completely undeveloped, but within a few days the simulated neurons have self-organised into a menagerie of clusters, each with a specialised task such as identifying simple visual features like edges. Other areas assemble these edges into complex features such as doorways or tea trolleys. Similar feature-detecting regions arise in living brains.

This is the crucial breakthrough, says Anil Seth, a theoretical neuroscientist at NSI. "Complex phenomena arise from relatively simple elements." Darwin and SAIL show it is as true for artificial brains as it is for real ones.

That this phenomenon occurs in animals was demonstrated conclusively in an experiment performed by Colin Blakemore and Grahame Cooper at the University of Cambridge and published in 1970. They raised kittens in pitch darkness, except for 5 hours each day during which they were put in pens painted with either vertical or horizontal stripes. After five months, the kittens were let loose in a normally lit room. Those that had been exposed only to horizontal lines would repeatedly walk into table legs, while animals exposed to vertical lines couldn't see horizontal edges. Each was effectively blind to edges in the direction they had not been exposed to during the formative period.

SAIL's brain begins as a homogeneous mass of overly connected neurons - just as the brain of a new-born baby is excessively connected until learning prunes some connections and reinforces others. It was designed to follow two rules well known to neuroscientists. The first is that neurons which fire together wire together: connections between neurons that fire at the same time when triggered by the same stimuli will strengthen, and vice versa. The second is sometimes called winner-take-all inhibition: if several neighbouring neurons respond to the same visual input, such as a vertical edge, then the strongest, the so-called "winner" neuron, will inhibit its neighbours from responding to vertical edges in future. This ensures that these neighbouring neurons remain free to respond to other stimuli.

Not all in the mind

In themselves, SAIL's edge-detecting abilities may not have great practical significance, as they are easy to design into conventional computer vision systems. But Weng says the same simple developmental rules should help robots like SAIL tackle more complex problems that computer programmers have difficulty solving, such as using a combination of memory, vision and hearing to navigate in unfamiliar terrain. "We're relieving humans from a tedious and intractable design task," he says.

Having a physical body is another crucial element steering the development of Darwin's and SAIL's brains, just as it does for living ones. In 1963, Richard Held and Alan Hein at Brandeis University, Massachusetts, published a report on an experiment that showed how crucial this is in living brains. They raised 10 pairs of kittens so their only exposure to light came when they were attached to a circular, rotating table top, like a lazy Susan. One kitten always rode enclosed in a box on the edge of the table. It could see but not walk. The other kitten was attached to the opposite side of the table by a harness, so it could walk in circles, rotating the table top and the first kitten with it. Both kittens saw the same view of the room as they turned, but one's visual experience was caused by its own activity, whereas the other's was not. This small difference had a huge impact. The kitten that pulled the turntable round developed normal depth perception within a few days; the kitten that merely rode on it did not.

The kittens developed such different responses because of the way the nervous system is wired. As a kitten steps forward, the nerves driving that movement notify the brain of that step, so it can adjust its interpretation of visual information. Disembodied artificial neural networks trapped in a desktop computer can't experience this. But being mobile, Darwin and SAIL do. They use movement to enhance what they learn from what they see. Darwin VII learns to recognise striped blocks by moving, so that it sees them from countless perspectives. SAIL uses movement to distinguish door frames from posters and other vertical edges. "You really do need a body," says roboticist Oliver Brock at the University of Massachusetts, Amherst.

Most AI robots of the past two decades used elaborate modelling systems to describe the real world around them, but failed on simple common-sense tasks. To make them function, programmers had to tackle the monumental task of anticipating all the likely objects in a robot's environment and how they might change. Will AI robots that combine movement with a brain-like neural network fare better? "I give a resounding yes to that," says Aleksander, who sees this approach as a way out of the hole into which traditional AI has dug itself.

Even SAIL, which starts each task with completely blank neurons, was able to overcome problems that seriously taxed other AI systems. First its brain developed areas to represent straight edges, then it combined these into relevant categories such as doorways and cabinets. And it did all this without a programmer's guidance.

Does that mean all other efforts to develop AI should be dropped in favour of developing systems that combine robot-like movement and neural-style processing? Stephen Muggleton, a computer scientist at Imperial College London, isn't convinced. "You may make more rapid progress by starting from well-defined mathematical processes," he says. These can be inspired by the brain, but they shouldn't be copied from it.

But Seth insists that in an unpredictable world, mimicking the brain on a detailed level will provide advantages that other approaches cannot. For example, real brains often have a lot of redundancy when it comes to performing a particular task. Similar redundancy might allow a brain-based robot to continue to function even if part of its brain is damaged.

Slowly Darwin and his friends are leaving their playpens and heading out into the real world. The latest version in the Darwin series is learning to play soccer riding on a Segway (*New Scientist*, 12 February, p 34).

Could such a robot one day even become conscious? Seth and Krichmar are both convinced it's possible. "I don't see any reason why it couldn't happen," says Krichmar. But as I watch Darwin VII pootle around the lab, it's clear that having a brain is only half the problem. Would you entrust your child's life, or the safety of nuclear waste disposal site to a clumsy if intelligent trashcan?