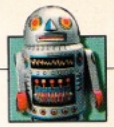

IT'S ALIVE!

One day soon your Mac could become home to dozens of digital creatures that will live and breed inside it, or perhaps run errands for you. Garrick Webster investigates artificial life and exactly what it could do for you and your Mac...

Orange-flavoured drink. Apple-flavoured drink too, if you will. Chocolate-flavoured cake covering. They're never quite as good as the real thing, are they?

Artificial intelligence (AI) has a similar reputation. In the movies and on TV, it's always going horribly wrong (as in *2001* and *Westworld)* or it's just plain flaky (like Marvin the Paranoid Android in *Hitchhiker's Guide to the Galaxy,* everyone in *Red Dwarf* from Rimmer and Kryten to the Scutters, or the big ED-209 – the over-zealous two-legged "gunslinger" robot – in *Robocop)*. Nevertheless, serious researchers into artificial life are determined to get it right.

Artificial life – or "A-life" – is a relatively new branch of science that melds the highest levels of computing and biology. For some programmers, the goal is to come up with a superior form of AI by building software models of the smartest computer currently known, the brain, and seeing if those models can learn and adapt their behaviour as we can. Others approach it in a different way, believing that computers can be useful for simulating biology.

Today, programs can very accurately simulate both organisms and their environments – creating competing, reproducing creatures, with the possibility for genetic mutation and adaptation, as well as the pressure of natural selection. Nature has taken some four billion years to get us where we are today, but computer simulation can both speed up evolution and answer new questions. What if we changed this gene here? What if that predator were eliminated from the food chain? What if this creature had a slightly larger brain? In the words of Chris Langton, who heads up the Santa Fe Institute, a key centre of A-life study, "By extending the horizons of empirical research in biology beyond the territory currently circumscribed by life-as-we-know-it, the study of artificial life gives us access to the domain of life-as-it-could-be."

What does all this mean to you? More lifelike and intelligent games, for a start – and therefore probably even more entertaining sims. Beyond that, A-life is likely to affect how you interact with your Mac, use the Internet, experience virtual reality and create computer artwork. With computer-controlled systems being used more widely, A-life could even have an impact on your real life at home, in the car, wherever you go. Perhaps surprisingly, the famous robot taxi-driver from *Total Recall* is one of the least unlikely applications of this new science. Read on and meet the main characters in the whole exciting story...

No. 39 TOTAL RECALL: GASP IN HORROR AS SINISTER ROBOT TAXI DRIVERS DRIVE VERY CAREFULLY INDEED.
CREDIT: CAROLCOITRI/STAR (COURTESY KOBAL)

Your Mac seems to have a life of its own, doesn't it? It crashes when it knows you haven't saved your work, it plays up at the worst possible time, and it keeps moving your files and icons around... Try to rectify a bad situation and it lets out a "quack" or a "ping" just to say "Hey, I'm watching you – and I'm laughing at you to boot."

Apple has always produced the most human, friendly and intuitive operating system, and sometimes it does indeed seem as if your Mac is alive. But it is, after all, nothing more than a bunch of electronic components. As artificial life creeps onto the scene, however, things could change...

The Game of Life

Actually, you've probably dabbled with A-life on your Mac already, playing *SimCity* and *SimCity 2000*. Both these classic games are actually just larger, more complex versions of cellular automata like John Horton Conway's *Game of Life* (see the box below), which was the first ever computer manifestation of what we now call artificial life.

In essence, it was lifelike because it was self-replicating. Whether any "cell" survived was determined by simple rules – if it was in contact with no other living cells, it would die of loneliness; if it was in contact with one other cell, it would remain alive; if it was touching two other cells, it would produce an offspring cell; and so on – but the exciting thing was that, turn after turn, the "colony" of cells grew and developed.

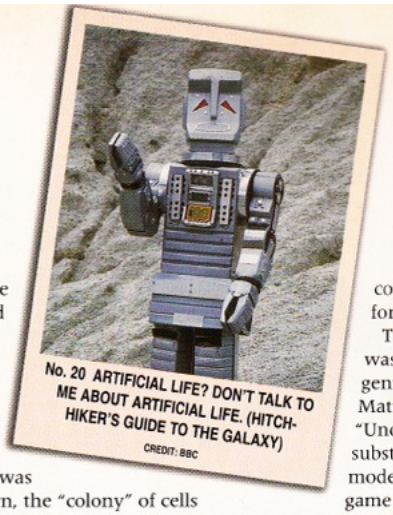

No. 20 ARTIFICIAL LIFE? DON'T TALK TO ME ABOUT ARTIFICIAL LIFE. (HITCH-HIKER'S GUIDE TO THE GALAXY)
CREDIT: BBC

Today, the main practical contact we have with self-replicating A-life is in games. Of course, the state of each cell on a *SimCity 2000* grid has many more rules applying to it than those in *Game of Life*, and usually involves user input too, but essentially it's the same. Once you've created commercial, industrial and residential zones, transport and a power supply, the little squares can spring to life. Buildings pop up, passing the power supply to adjacent cells, water pipes are built and SimCitizens appear, asking for schools, protection from crime and a fire service. As the cells change, or as you change the conditions around them, they interreact with one another. You just set up the system and watch a seemingly living city "grow" right there on your Desktop. As you dabble, it develops in new and unexpected ways, and it's very easy to get hooked.

Visit the Isle

The latest Sim game from Maxis, *SimIsle* (reviewed on page 115 this issue), is also basically a cellular automaton. However, rather than a set of rules, the state of each cell is determined by an extremely complex model of rainforest ecology.

This latest Sim-offering was designed by Intelligent Games, whose head, Matthew Stibbe, explains: "Underlying *SimIsle* is this substrate of an ecological model that divides the game into a series of cells. They're not actually rectangular – they're a bit more like polygons. What that does is it allows you to have much more organic shapes with your grid. But essentially the formulation is the same. You have a cell that's described as rain forest and the cell next to it is described as arable land, and an adjoining cell. The gaps between the cells can be roads or rivers. What you do is you work out what's happening in a cell and how that affects the cells next to it."

As Matthew Stibbe will readily admit, however, *SimIsle* is only a living system in a broader sense – we're not talking about

***SimIsle*, the latest in the popular *Sim* series, is different in that you affect the course of events by acting though agents.**

Automaton evolution

The cellular automaton is one of the best known types of artificial life program, but this sort of software also forms the basis of some of the most popular games ever.

Basically, a cellular automaton is a grid of cells. The internal state of each cell is determined by a set of rules. In *The Game of Life*, devised by Cambridge mathematician John Horton Conway in the '60s, these rules are brief and simple. If a live cell has no neighbours, it dies of exposure. If it has one, it survives, but in stasis. If it has three or more, it dies of overcrowding. But with two, it reproduces, using an empty cell adjacent to it.

Each turn, dead cells would be removed and new births would appear. Infinitely self-replicating patterns were discovered, along with structures of cells that could infinitely produce them.

Conway's game enjoyed a huge cult following in the 1970s, and its implications were astounding. Theoretically, this simple rule-based system, starting from a certain pattern of live cells, could grow into a massive living computer all on its own. In such an instance, however, the grid accommodating it would have to stretch across the known universe – or further! Conway speculated that a creature roughly as complex as a single-celled organism would require a grid measuring one million by one million cells.

Increase the number of rules and internal states, and you have a recipe to make an incredibly complex automaton life. Jon von Neumann theorised about a cellular automaton in which each cell could exist in one of 29 states. No-one has ever written a program like this, but Chris Langton wrote one in which each cell could exist in eight states. It grew into something that looked like a coral colony.

ON THE CD

You'll find demos of all these on the Cover CD

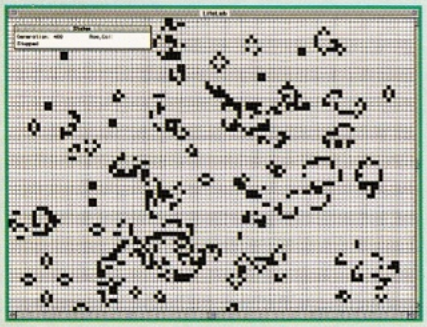

***LifeLab* is one of the fastest running examples of Conway's *Game of Life*. In this cellular automaton the rules are simple and two-dimensional.**

***Trevista* is essentially a 3D version of Conway's *Game of Life*. The rules determining life or death in each cell are therefore slightly more complex.**

In their highest form, such as in *SimCity 2000*, the rules for each cell are extremely complex. But what a fascinating game they make.

Wanted: godlike geneticist

Fancy populating a simple simulated environment with creatures of your own making, then fiddling with their genes, and generally trying to keep the whole confusing mess in order? Well, that's the kind of gameplay on offer in *SimLife*, a predecessor of *SimIsle*. You define physical environments by terrain and climate, then release plant and animal species into them to see if they would survive. Libraries of organisms are available for you to start with, or you can specify a genetic make-up of your own and see how well suited your organisms are to their environment, whether they get eaten by predators and whether they can survive long enough to reproduce...

ON THE CD

Vivarium is in the A-life Shareware folder on the CD

Or you could have a bash at *Vivarium*, a shareware A-life game on this issue's Cover CD and Disk. Here, you have only four species to dabble with, and there are no confusing genetic factors to master. You design both the rules your creature lives by and its capabilities. The senses on offer are sight and smell (olfaction). The stronger you make their senses, the more they need to eat. Likewise, the greater their movement capability, the more food they require. And finally, carnivores are hungrier than herbivores.

Our game suggestion is that you create lots of herbivores with weaker senses and movement capabilities, and then a smaller number of carnivores to feed on them. In the middle, you could add some cannibalistic omnivores. It's a challenge to get the balance right.

What's neat is that, although it looks rather simplistic, the creatures in *Vivarium* can actually learn. As they adapt their behaviour to the environment, they pass on their traits to their offspring. The ones with the best adapted behaviour are more likely to breed, and subsequent generations should therefore be genetically improved. It's a full Desktop artificial life environment!

With *SimLife* you can head back to the days of the dinosaurs and find out whether your humble dino-management skills would do a better job of keeping them alive than history did.

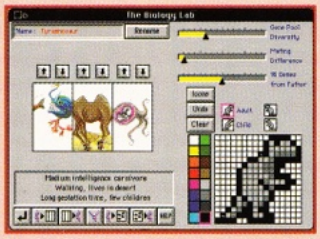

The genetics lab in *SimLife* enables you to design your own plants and animals. The design you come up with determines your life form's survival needs, its behaviour and, indeed, its genetic make-up.

The mountains are obstacles to the movement of the creatures. The ones which adapt to that fact tend to survive better. After all, you won't find a mate by staring at a rockface all day.

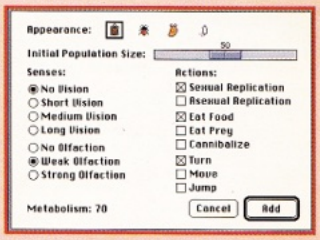

When you start playing *Vivarium*, create a new document, then select Add Lifeforms under the Options menu. Set up the capabilities of each creature you want to introduce. We usually make the mice carnivorous.

simulated creatures that breed, eat and fight to survive. "I think what we've done in *SimIsle* is we've created a model that is like a little Gaia model," he says, referring to the theory that the entire planet is a living organism. "It has so much complexity built into it that it's quite hard to predict how it's going to behave. But it has behaviour patterns, it has a kind of inherent objective.

"I would hesitate to say this, but it seems to be alive in the same way that maybe an amoeba is."

Playing *SimIsle*, however, you'll discover that it's not quite like the previous *Sim* games. Instead of interacting directly with the map as you do when you place a power station or railway line on the *SimCity 2000* grid, in *SimIsle* you implement your plans using partially intelligent agents. Each of these agents has certain skills, and each should go out into the rainforests and villages of the tropical islands in the game to do your bidding.

The use of autonomous agents that can help you with certain tasks, or even do your work for you, is very much a growing field within the study of A-life.

Secret agent man

Autonomous agents can take many forms. They can be characters in games, like those in *SimIsle*, or robots, like Robbie from *Forbidden Planet* or Number 5 from *Short Circuit*. They can be fully self-contained software utilities, like Internet search engines, or they can even be very simple strings of code, existing within a digital environment like computer viruses.

Any of these must, of course, respond to changes in its "environment" (if a search engine, for example, encounters whatever it is programmed to find, it must react); but what makes it A-life is an ability to consistently adapt its behaviour – that is, to learn. An intelligent search agent might, for instance, remember where it succeeded last time and *begin* its next search there.

Apple is developing software agents as part of future versions of the Mac Operating System. Your Mac won't quite be brought to life by Mac OS 8 (AKA Copland), or even the System version after that (codenamed Gershwin), but it will become more attentive to how you use your machine.

"In several small steps," says Onno Kluyt of Apple Europe, "we are developing intelligent agents. We don't think an all-encompassing intelligent agent in itself will work, so the process we are taking is more towards dedicated specialised agents." One, for example, might be a "telephone agent" that simply answers your phone and then records messages or acts on the messages, depending on which caller it is.

No. 33 THE EVIL CYLONS VOW TO WIPE OUT ALL MEN WITH DODGY 1970S HAIRCUTS (BATTLESTAR GALACTICA).
CREDIT: UNIVERSAL (COURTESY KOBAL)

No. 46 ROBOT NUMBER FIVE RUNS AMOK AS ITS WACKY SILICON BRAIN GETS A SHORT CIRCUIT.
CREDIT: TRI-STAR (COURTESY KOBAL)

"Same thing with agents that try to monitor and understand what you do with the machine and then provide help," Onno Kluyt continues. "They're basically the next generation of AppleScript and AppleGuide." So, for instance, if the system watches you attempt to log onto a network and fail, it will pop up and ask you if you'd like it to do the task for you automatically. It will even offer to teach you how to do it. Or if it sees

Chris Langton

Many regard Chris Langton as the father of artificial life. In 1987, he invited over 100 scientists and technicians to Los Alamos, New Mexico, to establish artificial life as a new scientific field. Today, Langton is the Director of the Artificial Life Program at the Santa Fe Institute, the pre-eminent research institution in this subject area. He also heads up the Swarm Project, which is developing a multi-agent simulation tool intended for A-life researchers. What's more, he's also Editor-in-Chief of the *Artificial Life Journal*.

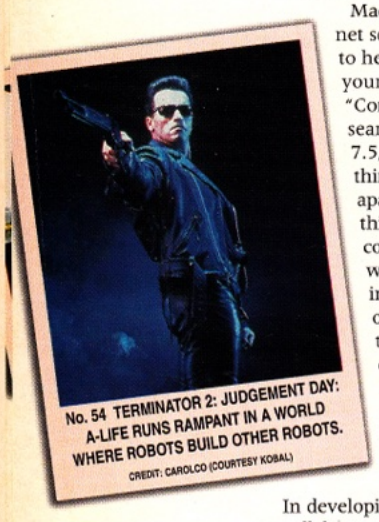

No. 54 TERMINATOR 2: JUDGEMENT DAY: A-LIFE RUNS RAMPANT IN A WORLD WHERE ROBOTS BUILD OTHER ROBOTS.
CREDIT: CAROLCO (COURTESY KOBAL)

you repeatedly going about a particular function the hard way, it will notify you of the quicker keyboard shortcuts that are available to you.

Mac OS 8 will include an Internet searching agent called *Cyberdog* to help you find things both on your Mac and out on the Internet. "Compared with the current Find search that you find in System 7.5, it will give the user two things," says Kluyt. "One is that apart from just the name of things you can also search the content of documents. Also, it will provide a relevancy ranking of the documents – based on the search criteria the user types in, it will order the documents according to how relevant it thinks these documents are."

Go fetch, doggie

In developing *Cyberdog*, Apple actually collaborated with some of the top names in Internet searching, including AltaVista (**http://www.altavista.com**). It did so not only in order to exploit their expertise in writing "Find" algorithms, but also to get help with incorporating the Internet with the Finder. Yes, with new versions of the Mac OS, folders of data on the Net will appear on your Desktop, just like folders on your hard drive (though access to them will necessarily be slower).

Over the next year or two, all this should start making your Mac even friendlier to use. However, much of the impetus for A-life research has been for less than friendly purposes. The American defence industry has been keen to develop a variety of systems, including "intelligent" robot bomb disposal teams, software designed to penetrate enemy network systems and disable them (maybe Macs could help here, if *Independence Day* is any guide), self-driven attack vehicles and even the predictable, *Terminator*-style metal warriors.

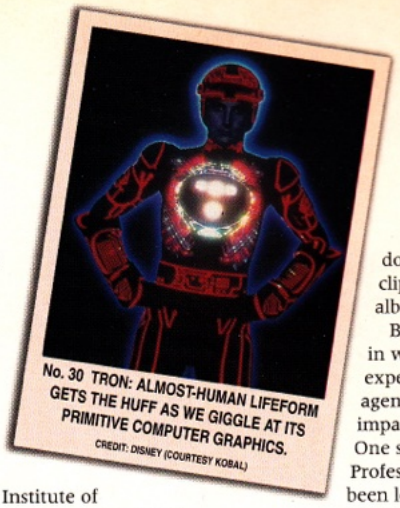

No. 30 TRON: ALMOST-HUMAN LIFEFORM GETS THE HUFF AS WE GIGGLE AT ITS PRIMITIVE COMPUTER GRAPHICS.
CREDIT: DISNEY (COURTESY KOBAL)

Nevertheless, the Autonomous Agents Group at the Massachusetts Institute of Technology, one of the major centres of A-life research, also has more peaceful projects under way. One is a match-making agent called *Yenta*. You give it your personal romantic details and a description of what it is you're looking for in a dating partner, then release it out onto the Net to seek out autonomous agents representing other lonely hearts. If two closely matched agents were to find one another, they would exchange contact information for their patrons and return home. You could then decide whether or not to e-mail or telephone the person *Yenta* deemed suitable. No need to worry about eye-contact, chat-up lines or awkward body language – *Yenta* would break the ice for you!

Another agent project being developed at MIT is *Kasbah*, which can trade and barter in a digital marketplace on behalf of businesses. You enter the highest and lowest prices you'll pay for some commodity and the quantities you require, and set the agent loose into the trading forum. The hope is that *Kasbah* will wheel and deal, then return having got the best possible price for you.

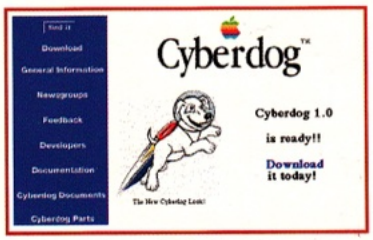

Apple collaborated with AltaVista when making *Cyberdog* in order to get expert advice on search algorithms.

Even further down the road is a system in which numerous competing agents vie to complete the tasks you set them (information retrieval or whatever) as quickly and accurately as possible. According to the Darwinian laws of natural selection, the ones that fail will simply die out. Harsh, yes, but no one said nature – or even simulated nature – was kind. The ones that learned quickest, however, would enter new rounds of competition, and some would, perhaps, specialise themselves for specific tasks.

In the *Webdoggie* project, meanwhile, an agent is being developed which offers you Net information it thinks you will be interested in, based on a survey of the tastes of thousands of other users. So, if you like Black Sabbath and you want suggestions as to similar bands to check out, the *Webdoggie* would refer to the profiles of all the other Black Sabbath fans for the most commonly suggested other bands. It might even go on to download some music clips, discographies and album reviews for you.

But inevitably, the area in which the MIT experts expect autonomous agents to have the most impact is entertainment. One system that MIT Professor Pattie Maes has been looking at is called ALIVE – Artificial Life Interactive Video Entertainment. Users of this system stand before a huge video monitor and can move about in a 16ft by 16ft area. With video cameras trained on them, their images are composited into the 3D environment displayed on the screen, and they can watch themselves interacting with lifelike artificial agents like a dog, a parrot and a hamster. The experience might be like stepping into the world of *Tron* or perhaps *Lawnmower Man*. Another project is an autonomous aerobics instructor which will actually correct users who are performing exercises incorrectly.

Quest for the grail

Anyone who's played with a multi-user dungeon (MUD) may already have come across very lifelike non-human characters. Played over the Internet, this form of role playing can involve lots of characters, both human and computer controlled. Usually, the monsters are computer controlled, but occasionally you encounter characters called chatterbots which behave so realistically that it's hard to be sure.

Such realism is pretty much the holy grail for games makers. If you've ever had the chance to play deathmatch (multi-player) *Doom* or race against human drivers in multi-player racing games, you'll know that they're infinitely more fun than the stiff, methodical single-player versions. According to Dr David Cliff, lecturer in Artificial Intelligence and Computer Science at the University of East Sussex, the crucial element is that your opponent "is trying to monitor your strategy and tactics and find flaws in the way you play and exploit those flaws. It will take quite some time, but artificial life technologies offer a mechanism for incorporating that kind of adaptability, stealth and desire to outwit the player". ➤

Pattie Maes

Professor Pattie Maes is the principal investigator of the Autonomous Agents Group at the MIT Media Lab. As well as researching ways in which artificial life technology will contribute to the development of computer entertainment (and entertainment in general), Professor Maes is working on the development of numerous sorts of autonomous agents, some of which are intelligent search agents, others that have commercial purposes and others still that are more light-hearted in nature.

With the ALIVE system, you stand in front of a large screen and interact with characters such as Silas T Dog, seen here. The dog will interpret and respond to your actions and gestures, such as pointing, in a realistic manner.

Wouldn't it be great if the computer-controlled characters in *Marathon* behaved more like your opponents in multi-player mode? That would be a challenge.

No. 28 BLADERUNNER: SO REAL EVEN THEY DON'T REALISE THEY'RE NOT HUMAN – THE REPLICANTS MARCH.
CREDIT: LADD COMPANY/WARNER BROS (COURTESY KOBAL)

The answer, say those on the forefront of A-life technology, is to stop giving game characters motivations or intelligence based on sets of rules. Rules result in predictable, unchanging behaviour – the little people in *Theme Park*, for example, are programmed to get hungry or thirsty, and to need the toilet after these artificial appetites are satisfied, and even to throw up if they ride the roller-coaster too soon after eating, but not one of them will ever learn the food-before-ride-equals-vomit equation. It's up to *you* to position the burger bar further from the ride or make the queue longer.

No. 19 BLAKES SEVEN: GAZE IN AWE UPON ARTIFICIAL LIFE SO SMART IT GETS BETTER LINES THAN THE HUMANS.
CREDIT: BBC

The alternative to these rule-based characters is to use brains based on neural networks (see the box below). That's what you can expect in *Creatures*, a new "life simulation" from Millennium Interactive (previewed in MACFORMAT last issue). The behaviour of each of the creatures of the title (called "norns") is controlled by a neural network capable of remembering inputs and relating them to certain outputs – in short, learning. So if, for example, you give a norn a carrot and it eats it, you can scratch the creature's head to reinforce the idea that it has done a good thing. If it throws the food down, you can smack it. If it happens to eat some poisonous food and get ill, it might remember to avoid that food next time, and so on.

Motivation for the norns is provided by an internal feedback system. Anil Malhotra, director of the Millennium CyberLife project, which came up with *Creatures*, explains. "We have a virtual endocrinal system. We've got pain increasing hormones, we've got adrenaline in there, we've got all the biochemical enzymes needed to impute that you're hungry etc, etc."

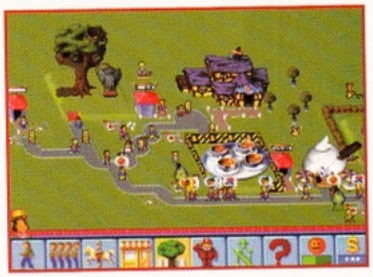

The little people who visit your funfair in *Theme Park* are rule-driven: if they get thirsty and have a drink, they need a toilet. So predictable – yet so credibly lifelike.

Ten times a second, information from a norn's virtual biochemistry washes over its neural network brain. If there's a strong lack of glucose, the creature will sense hunger, and if it has learned about food, it will try to find some. If it touches some fire, a sharp increase in pain hormones will result and it should first move away from the heat and then register that fire is painful to touch. When it plays with a ball, endorphins will go to the brain and the norn should remember that the ball is a pleasurable toy. "There are actually 250 genetically specified internal chemicals," Anil Malhotra says. "That's a substantial amount. We've even got testosterone so our creatures actually have a motivation to breed in exactly the same way you and I do."

In fact, that's not quite accurate – all it takes for a norn is a big sloppy kiss with a suitably receptive mate. But in one crucial respect, norns' reproduction is just like the real thing. "We wanted their offspring to have the capability of having a unique set of genetic characteristics," explains Malhotra. "So we created this quite complex software which created a simple but authentic model of DNA, which we call Digital DNA."

Neural networks: lifelike artificial intelligence

Neural networks promise a smarter and more adaptable form of artificial intelligence than the rule-based routines found in most software. But what are they?

Neural networks are, very basically, software constructs modelled on the way that brains work. So, instead of a long and complex list of rules that try to govern every situation the system is likely to encounter, all a neural network offers is a complex web of interconnected neurons, all of which can receive, store and pass on binary data.

Together, a very large number of neurons work together to solve problems – if there's no solution in the few basic rules, the system doesn't trip up but keeps trying until, by trial and error, it comes up with the correct response. Once it has identified a way of handling a certain type of input, the system will automatically know what to do the next time it senses such information. Hence, neural networks are capable of real learning.

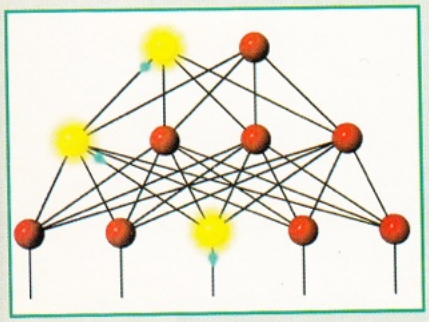

As computer-driven neural networks grow in complexity, could they become sophisticated enough to be considered new lifeforms?

Take a simple example. A neural network for optical character recognition software might be fed a graphic of a letter. It will try and "guess" the letter using a simple matching routine. If it gets it wrong, it will be given the letter again – but remembering the wrong answer each time, it will approximate towards the correct answer. Eventually, passages will form through the neurons, enabling instant recognition of letters of numerous typefaces – even ones from torn or defaced scripts.

Each neuron in a neural network can also have a non-linear response. That means the output a neuron gives needn't be equal to or in direct proportion to the inputs it receives. If, for instance, it were receiving stimulus at a rate of 10 times a second, it may give no response. Increase it to 20 times a second and it might react with a charge valued at five. Hit 30 cycles a second an it could very well send a charge to other neurons in the net worth 500, so that some kind of action is taken to reduce the input level.

Add the possibility of feedback paths within a neural network and, the theory goes, every output of the network overall can also produce further inputs internally, within the system. With closed feedback loops bouncing new impulses back and forth between the neurons, the network can begin to stimulate itself. Potentially it could even develop a memory and internal state that's affected by things that have happened to it in the past. This is an area of active research at the moment. Could it mean emotion or consciousness could arise in a neural network?

At the moment, neural networks are being employed in OCR software. Also, at least one bank in the United States is using a neural network to detect credit card fraud by identifying uncharacteristic spending patterns.

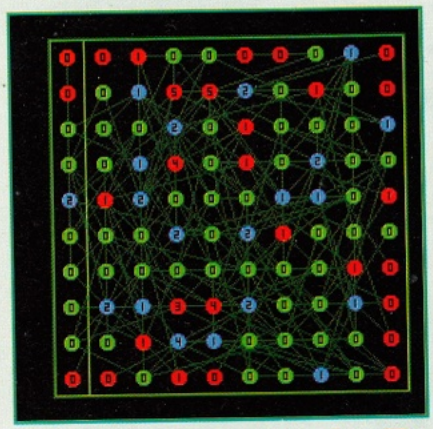

Try out the *NeuroSim* program on our Cover CD. Clicking on one of the sensory neurons down the left sends an impulse to the neurons it is connected to. This may or may not change the charge held in that neuron. When you've clicked one neuron two or three times, you'll notice some of the neurons it's connected to discharging energy, which may change the state of other neurons in the network. This shows how a neural network responds to inputs by passing on information.

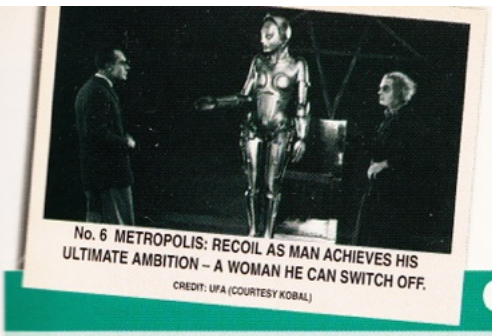

No. 6 METROPOLIS: RECOIL AS MAN ACHIEVES HIS ULTIMATE AMBITION – A WOMAN HE CAN SWITCH OFF.
CREDIT: UFA (COURTESY KOBAL)

Creatures roaming free...

Don't be deceived by the cutesy feel of *Creatures:* underneath the sweet graphics lies a highly sophisticated computer-based learning system. Watch as the creatures (known as norns) explore their world and evolve according to what they discover.

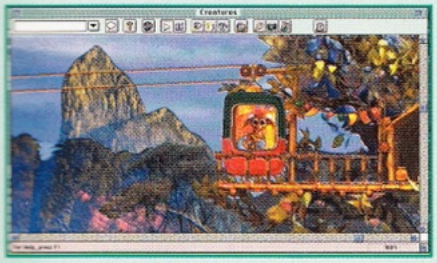

With their ability to learn, the norns in *Creatures* will eventually become able to make use of the cable car which can be found in one part of their vivarium-style environment.

The *Creatures* environment is by no means homogeneous. In one corner your norns will find a desert island habitat. Here we see a norn trying to figure out a radio device.

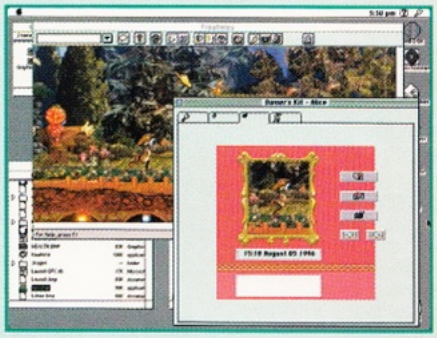

One of the function screens in *Creatures* is the Owner's Kit. There's one for each norn, and in it you can store snapshots of your pets as well as notes on their progress.

The incubator machine to the left is where you hatch your norns from eggs. In the immediate vicinity your pets will find cheese, honey, a jug of water and a lift going up to a learning machine which you can use to teach them language.

As they have independent brains and a propensity to chatter, your norns will have the ability to teach one another to talk. The younger ones will constantly say "goo" until they manage to learn some more robust diction.

If you teach your creatures well, they should learn that the garden is a good place to go for food. Here they'll find carrots, lemons and – what's that in the greenhouse? A pot of honey – good for the blood-sugar levels.

During breeding, a daughter genome is created on which are placed random characteristics from each parent. Built into the process is the possibility of a cutting error taking place, so that one characteristic is simply left out. There's also the chance that something is duplicated, or that an actual mutation occurs. "Through all three of those mechanisms, the probability arises that the daughter genome is not identical to either of its parents," continues Malhotra. "The daughter genome in the young creature is responsible for its brain development, its biochemical system and therefore its reaction rate to good and bad events, its ability to take flight, its ability to breed, etc. So you're getting real bio-diversity introduced into these creatures through the generations."

This means norns aren't just virtual pets; they're real evolution in miniature. As Malhotra explains, "Real evolution has only ever happened once, so it's very hard to go back and to know what would have happened to the evolutionary pathway if different steps had taken place. The theory is that something other than human beings would have been the outcome. Scientists are particularly interested in *Creatures,* because with potentially hundreds of thousands of users in the world, breeding over time tens or hundreds of thousands of generations, it might be possible to detect how, and for what reasons, particular diversifications in the genetic code take place."

Viruses and worms

If, as Anil Malhotra hopes, *Creatures* does succeed, the norns may come to rival the currently best known form of artificial life: computer viruses. Many claim

No. 63 STAR TREK: THE NEXT GENERATION: ANDROID SEEKS TRUE EMOTIONS IN BID FOR SPIN-OFF SERIES.
CREDIT: PARAMOUNT TELEVISION (COURTESY KOBAL)

The von Neumann Automaton

Back in the 1950s, when computers as we know them didn't exist, the highly regarded mathematician Jon von Neumann thought up a machine which would, just like organic life, be able to reproduce. He called it an automaton.

All the intelligence that controlled the machine's behaviour would be stored on a ticker tape. When read through a device controlling the machine's activities – like a nervous system – this tape would cause the machine to move around picking up mechanical components (which would be in abundance in von Neumann's imagined environment) exactly like the ones of which it would be made.

As the tape was being read, it would be duplicated and the copy of it stored. Once all the components for its own clone had been collected and assembled, the machine would insert the behavioural information into its offspring and the new automaton would go off and do the same thing.

Aside from the fact that an environment littered with precisely the right sorts of components for the automaton to find and assemble is extremely improbable, it all sounds a bit scary, doesn't it? With the theoretical possibility of minor flaws in the duplication of the ticker tape, the machines could even mutate into something more complex and powerful, *Terminator*-style. Fortunately, no-one's ever built a von Neumann automaton, nor are they likely to.

What's particularly interesting, though, is the way von Neumann envisioned the transfer of genetic material. It sounds very much like the process followed in living things, but when he wrote his theories, DNA hadn't even been discovered.

Jon von Neumann

Von Neumann came up with the idea that a self-replicating machine could be built: an automaton. Though impractical at the time, his accompanying theory that life is based on logic rather than being a God-given force underpins the whole concept of artificial life. Von Neumann's main field, however, was physics, and he was a participant in the Manhattan project, the team responsible for producing the first atomic bomb.

that viruses are the form of A-life which most closely conforms to the textbook definition of biological life.

To begin with, they exist to survive and "procreate" – when they "infect" a piece of software code, they typically modify it so that each time the software is run it duplicates or reproduces the viral code. Filling up your disk space and crippling or corrupting the host software are commonly just incidental to this imperative to reproduce. What's more, built into many viruses is also the ability to seek new programs to infect, via files exchanged between computers.

No. 13 2001: A SPACE ODYSSEY: CRY HAL! AS DESPERATE HUMANS STRUGGLE TO OUTWIT A COMPUTER'S COLD LOGIC.
CREDIT: MGM (COURTESY KOBAL)

Then there is the added pressure of natural selection in the form of the numerous virus-checkers and killers available. Some of the more advanced viruses have been programmed to attempt to fool these by mutating their own code with each new generation, "evolving" to try and keep one step ahead of the predator.

But don't be alarmed. The viral code that controls this "mutation" can itself be detected fairly easily, and modern virus killers can root out whole virus "families" without much trouble. Regular use of a reputable virus-killer like *Disinfectant* (the latest version is on the MACFORMAT Cover CD every issue) will keep your Mac safe.

The point is that the way computer viruses work is closely analogous to the life cycle of real viruses. What if A-life researchers could synthesise populations of simple, virus-like but *benign* organisms that live, reproduce, mutate and adapt within computers? In fact, that's exactly what they're doing in the *Tierra* project.

Get A-life degree

Fancy an MSc in Evolutionary and Adaptive Systems? Contact the University of Sussex on 01273 678754, or **http://www.cogs.susx.ac.uk/lab/adapt/easy_msc.html**

MacTierra

For a masterclass in dabbling with artificial evolution, try *MacTierra*. Based on Tom Ray's original *Tierra*, this piece of software turns your Mac's processor into a veritable primordial soup inhabited by primitive code creatures that need only RAM and computing cycles to survive.

ON THE CD

Simulate evolution on your Mac with MacTierra

***MacTierra* is rather complex to understand but fantastically well put together. You can inspect the genetic code which makes up each 'worm' and even monitor population levels using the graphs. For complete details see the HTML file in the *MacTierra* folder. *MacTierra* was written by Simon Fraser (smfr@ santafe.edu or http://www.santafe.edu/~smfr/).**

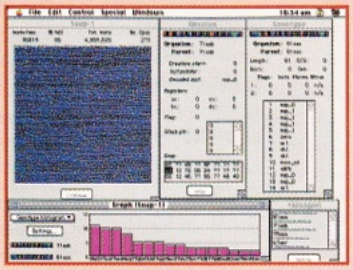

It's life, Jim, but not as we know it. [We apologise for this caption. It's tacky and obvious – but it was written by a Mac without any human intervention. Wow!]

Conceived and designed by a biologist, Tom Ray, *Tierra* is a software model of a section of the Costa Rican rainforest containing thousands, or even millions, of tiny software "organisms". Nothing more than lines of code containing instructions on how to survive and reproduce, these creatures (or "code worms") will occupy a computer's RAM, but to feed and procreate they require free computational cycles. In this respect they are unlike viruses, which consume software code and not processing power. The faster the chip in the computer they inhabit, and the fewer applications besides *Tierra* it is running, the more sustenance they get and the faster they can reproduce. With the possibility of "gene" duplication and mutation built into the way the worms reproduce, adaptations start to occur.

Tierra can run on a single machine, but the experiment is being extended via the Internet, with some 100 computers in the US and Japan running the *Tierra* software. When it's day in Japan and people are working on the host machines, the worms retreat over the Net to California to feed off the free processing cycles available. When the sun rises in America, the worms return to Japan again.

It really is a huge experiment in evolution. Who knows how advanced these worms will become? Currently, on-line *Tierra* is limited to Windows and Unix systems, but if you're interested there is a non-network Mac version called *MacTierra* on our Cover CD – see the "MacTierra" box above for more information.

Evolving into tomorrow

If such sophisticated forms of A-life were to be combined with virtual reality technologies, we could see some highly

Artificial life on the Internet

The Internet has, as you'd expect, broadened A-life horizons considerably. It's being used not just to exchange information, but to conduct huge on-line experiments.

The best by far is at TechnoSphere. This site welcomes you to design a creature and then release it into the TechnoSphere, a virtual 3D environment, and see how it will survive.

First, you decide whether it should be a herbivore or a carnivore. Then you choose its means of locomotion, eyes, body shape, mouth shape and so on. The choices on offer look more than a bit loony. Your beastie may end up moving using jagged wheels and have eyes on long stalks. (You can even download a "postcard" from the Web site showing what your creation looks like.)

When your critter's done, you name it and release it to make its way in the TechnoSphere. You can't watch its progress, but TechnoSphere automatically e-mails you when your animal dies or manages to procreate. We created two electronic beasties, a herbivore called "The Cudster" and a carnivore named "Yabbler". As far as we know, Yabbler is still alive and well, but The Cudster very quickly became a meal for a family of ravenous carnivores. That's life.

While the beasts at TechnoSphere look madly surreal, the ones at ALife Garden look more like impressionist amoebas. Coming to you from Hiroshima, the ALife Garden can contain up to around 500 creatures, which can feed, move about and breed. Evolution can occur over time. You can, whenever you like, call up a picture of one of the animals, see its size, age and parentage, and read a log of what it's been up to. Soon, visitors will be able to start up their own personal gardens containing creatures all their own.

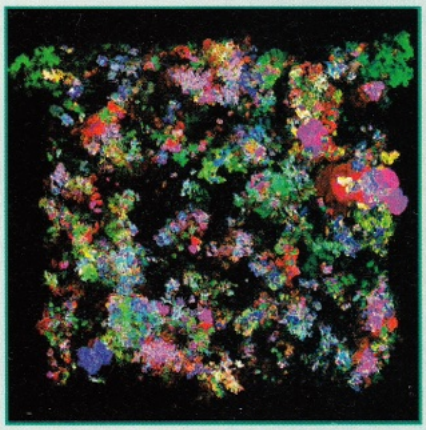

This is the ALife Garden. Basically, it's a mass of amoeba-like creatures of varying colour.

This graphic is taken from an animation about *Tierra* and it represents a digital environment presided over by a Reaper, responsible for killing-off defective organisms.

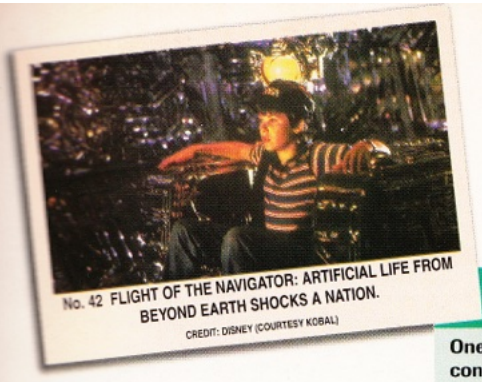

No. 42 FLIGHT OF THE NAVIGATOR: ARTIFICIAL LIFE FROM BEYOND EARTH SHOCKS A NATION.

CREDIT: DISNEY (COURTESY KOBAL)

complex real-world ecosystems being recreated for the purpose of hyper-realistic entertainment. In some quarters the aim is to create models of tree and plant growth so that the foliage in VR environments is dynamically changing and thus absolutely authentic.

However, evolution-capable A-life will have an impact on games even sooner than that. Already near completion is a game entitled *Galapagos* in which you'll have to train an insect-like creature to overcome the obstacles in its world. Developed by an American firm called Anark, it uses the company's proprietary NERM (Non-stationary Entropic Reduction Mapping) system, which enables the creature in the game to adapt to any new environment. Traditional rule-based systems of AI are thrown out the window.

Then there's *Project von Neumann*, an A-life game based on the ideas of Jon von Neumann. It's set as a space-based shoot-'em-up, where a machine that resembles von Neumann's automaton (see the box on page 37) is sent into an asteroid belt and forgotten. Years later, human outposts are attacked by superior, evolved versions of the machine. In fact, the A-life engine in *Project von Neumann* will even enable the machines to adapt during the game, making them harder to destroy. They'll also be able to consume the corpses of dead comrades to gain energy and ammunition. In turn, you too will be able to take advantage of the organic qualities of the craft in your fleet. You will, for instance, be able to "breed" ships for specific tasks, perhaps like the organic craft in the film *Flight of the Navigator*. The game is currently being written by a set of dedicated A-life hobbyists and no release date is available just yet.

***Project von Neumann*, named after one of the fathers of A-life, will feature "living" robots battling it out in shoot-'em-up action. It's to be a cross platform multi-player game.**

Further in the future, Dr David Cliff speculates, "perhaps you could be a trainer to a Gladiator or something who would enter Internet Gladiator competitions. You could monitor the performance of your agent, and retrain it to perform better in subsequent competitions. You'd also expect it to be learning while it's taking part."

What about more pragmatic applica- ➤

Dawkins' biomorphs

One of the most eminent scientists converted by A-life is the Oxford evolutionary theorist Richard Dawkins. Using a Macintosh, Dawkins wrote a program to simulate genetic evolution. Each matchstick shape in the program had a number of genetic characteristics, which could mutate over time. Users of the program simply choose the biomorph they like and that is the one which lends its genes to the next generations. Dawkins was surprised to find that the simple matchstick-like trees could evolve into shapes that began to resemble insects.

Log on to http://www.fusebox.com/cb/morphs/docs.html and you can play with a Java applet inspired by Dawkins' original biomorph program. What sort of shapes will you help to evolve?

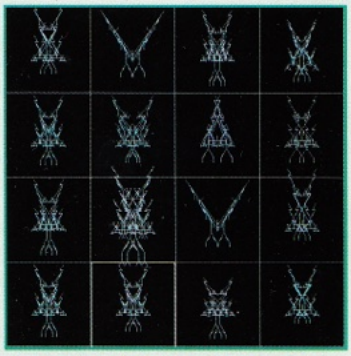

The Morphs applet at the Live A-Life Page can evolve anything from radical Christmas trees to strange insectoid creatures.

Richard Dawkins

A professor at Oxford's New College, Richard Dawkins is best known as the author of books such as *The Blind Watchmaker*, *The Selfish Gene* and *River out of Eden*. Using his Macintosh, Dawkins wrote a program that simulated genetic evolution and, from simple stick-figure beginnings, was capable of evolving elaborate insect-like graphic organisms. His work has influenced many people, including the artist William Latham, as well as A-life programmers the world over.

The biggest on-line A-life experiment is probably Tierra. Its home page explains what it's all about. Mac users can't fully join in just yet because there's no Net-savvy version for the Mac. You can, however, see how it works using a non-networkable Mac version called *MacTierra*. It's on this issue's Cover CD.

Mac software is available in great abundance at the Macintosh Artificial Life Software site – that's what it's there for. In addition to *MacTierra* and *Vivarium*, you'll find programs like *Neoterics*, *Trevista*, *CopyCats*, and *Bugs*. All excellent stuff.

Numerous Web sites demonstrate on-line "applet" versions of A-life software. The Live Artificial Life page has a whole collection of them framed. Log on and you can watch cellular automata grow, see Bots walk around or watch *Swarm's* miniature creatures swim around together in a frame.

Finally, of course, there's plenty of information on artificial life in general to be found round the Web. An essential stop if you're looking for software is Zooland. It has links to numerous Mac A-life resources. More academic in approach are sites like Artificial Life Online, Autonomous Agents Group and a paper called *Artificial-Life Simulators and Their Applications*. Log on if you feel the urge to research A-life further.

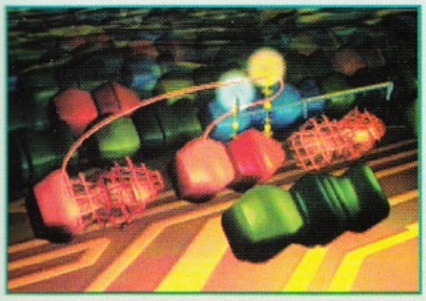

Another image from the *Tierra* video. Here we see a red parasite that's in the process of replicating itself by stealing processor power from other electronic creatures.

Artificial Life Online comes to you from Santa Fe New Mexico, a world centre of A-life excellence.

Internet A-life Index

ALife Garden – http://www.cisnet.or.jp/~alifeg/
Artificial Life Online – http://alife.santefe.edu/
Artificial Life Simulations and Their Applications – http://alife.santefe.edu/alife/topics/simulators/dret/dret.html
Autonomous Agents Group (MIT) – http://agents.www.media.mit.edu/groups/agents/
The Live Artificial Life Page – httpL//www.fusebox.com/cb/alife.html
MacTierra – http://www.santefe.edu/~smfr/mactierra.html
Technosphere – http://194.80.30.14/technosphere/index.html
Tierra – http://www.hip.atr.co.jp/~ray/tierra/tierra.html
Zooland – http://alife.santefe.edu:80/~joke/zooland/

tions? Millennium Interactive already has plans to use neural network brains, similar to the ones used in *Creatures*, in an intelligent Internet search agent. You would train such an agent by chastising or rewarding it according to the quality of the information it collected for you. Eventually, it would become completely in tune with your needs.

According to Anil Malhotra, this agent might do its searching during periods when servers are quiet. So, like the code worms in *Tierra*, it will be most active at night, when fewer people are using their machines. Perhaps you'll be able to get up in the morning to find a personalised newspaper waiting for you on your Mac Desktop, containing all the latest stories pertaining to your personal and professional interests. Wonder what it would put on page three...?

No. 58 RED DWARF: FUTURE ROBOT DISCOVERS WHAT BEING REALLY ALIVE MEANS – LAGER, FARTING AND NO SEX.
CREDIT: BBC

David Cliff

David Cliff is Lecturer in Computer Science and Artificial Intelligence at the University of Sussex in Brighton. He has worked in artificial life for some eight years creating computational models of insect nervous systems, developing neural networks to control robot systems and working on ways of adapting A-life software for entertainment purposes. We thank David for his help with this feature.

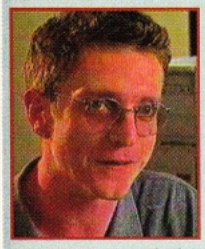

Stepping out

What about life away from your Mac? According to Dr David Cliff, there's a good chance we'll one day have intelligent agents residing in our TV sets. Once it learned what you enjoyed, your agent could help you choose what to watch from not just 40 or 50 satellite channels, but perhaps a thousand digital channels. It might alert you to shows it thinks you'd be interested in, or tape them for you, possibly without you even having to ask it to do so – a trick that today's VCRs will never be able to master.

Perhaps artificial agents could make you a safer and better driver. Dr Cliff speculates that an A-life co-driver could learn your individual driving style, and then very gently correct your worst habits – or warn you only if you were driving uncharacteristically fast or consistently making a series of unusual or dangerous errors.

Meanwhile, Toby Simpson at Millennium Interactive, who is working on *Creatures* and CyberLife, envisages A-life traffic lights. "A traffic policeman is able to control traffic better than a pair of traffic lights," he says. "Traffic lights don't respond to traffic changes, when there's a flower show or a fair or something. Biological organisms can learn new things. In the case of intelligent traffic lights, they would care very deeply about making traffic flow. If their instinct were to feel 'pain' as the queues got bigger, so it would be in their best interest to learn how to achieve a lessening queue.

"And even better, all the individual traffic lights could then talk to each other. The traffic control system in a town or a city could become one big biological system that learns and adapts itself to any problems."

Further into the smoggy future, how about a vacuum cleaner that lives in a cupboard under your stairs and comes out at night to silently clean the house? Perhaps your whole house could be controlled by a neural network in your

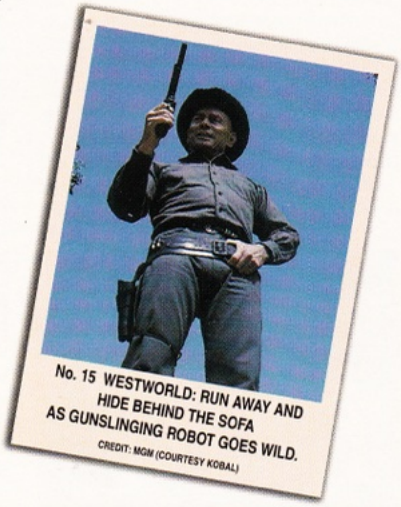

No. 15 WESTWORLD: RUN AWAY AND HIDE BEHIND THE SOFA AS GUNSLINGING ROBOT GOES WILD.
CREDIT: MGM (COURTESY KOBAL)

Art life

A few years ago, artists were fascinated by the patterns created by fractal mathematics. Now they're drawing on artificial life.

One of Britain's foremost computer artists, William Latham, is famous for his work in the field of organic art. Ever since his days as a student, Latham has been interested in the idea of "breeding" shapes. He started off with pen and paper, but very soon the multiplicity of forms he was coming up with necessitated the use of a computer. Between 1988 and 1994, Latham worked as a Research Fellow at IBM, where with mathematician Stephen Todd he wrote a program called *Mutator*.

With *Mutator*, Latham can create interesting, otherworldly shapes by breeding colour and form – that is, combining genes from four, five, six, even ten parents. The process is evolutionary with the artist acting as a sort of aesthetic predator, eliminating unwanted or ugly offspring, or using a genetic algorithm that determines which shapes live and which ones die according to defined criteria, or, introducing mutation factors.

Then of course there are the forms themselves. Again, Latham borrows from nature. The shapes being melded together include bone structures, ammonite shell patterns, fractal ferns, vertebrae bones, palm fronds, branches, rib structures, leaf shapes, tree growth structures, tentacles and so on. The works really can look quite stunning. Earlier this year Latham released a screensaver creator for Windows which enables PC users to generate patterns and 3D animations much like his own. It's called *Organic Art* and we think it should be on the Mac.

Another renowned A-life artist is Karl Sims from MIT. Sims is interested in both animation and the breeding of forms. In 1991, he produced an award-winning computer animation called *Panspermia*, in which a seed crash-lands on a planet and germinates. It grows, spreading itself across the land mass. Tangled forests arise, covering the planet. Eventually, it develops cannon-like seed pods, and when these explode new seeds are launched into space to take over new planets.

To produce his animation, Sims didn't use conventional software. Instead, he wrote plant growing simulations containing 20 genetic factors as the basis for foliage development. These covered fractal limits, branching factors, scaling and so on. Eventually, he rendered the

Computer creator William Latham doesn't make art – he breeds it using artificial life equations.

Latham blends natural forms like shells and plant shapes to make something familiar yet alien.

No. 17 STAR WARS: LEARN THE MEANING OF FEAR AS ROBOTS WONDER IF THEY'LL BE IN THE NEW FILMS.
CREDIT: LUCASFILM/20TH CENTURY FOX (COURTESY KOBAL)

Virtual sheepdog

Sounds incredible, but it's true: we've heard down the wire that one A-life researcher has come up with the dream of building an artificial sheepdog. Dr David Cliff, who brought us this news, explains what's so interesting about this off-the-wall project.

"The amount of intelligence in a dog is not normally thought of as high, but when you think about all the things you'd have to specify an artificial sheepdog to be able to do, it rapidly becomes a very, very difficult task.

"They have to cope with all different lighting conditions, different weather, different types of ground – bumpy ground or flat ground, and also mushy, muddy ground or hard ground that's been baked dry in the sun. And they have to be able to filter out the whistles of the sheepdog trainer from other noises, and they have to be able to identify the rough direction in which different sheep are going in order to predict their movements and interfere with them.

"Specifying all those kinds of conditions in a rule-based system, where you have to think of every eventuality in advance, can be very, very difficult. So," Dr Cliff concludes, "it might be that if you wanted to build a robot sheepdog, these artificial life technologies – where, when you're issuing a punishment or reward, the system is adapting itself – might be a more appropriate technology to employ."

Wonder if it will be allowed to compete in *One Man and his Dog*?

No. 18 DR WHO: SWOON AS COMPUTERISED CANINE BECOMES A TIME LORD'S BEST FRIEND.
CREDIT: BBC

Mac – security systems, climate controls, coffee machines, cleaning robots, curtains, lighting and so on. Unbelievable, you say? In the United States, some office buildings are already run by central AI systems.

But what about the Frankenstein factor? Could all this A-life technology run away with itself, and perhaps even turn hostile? Could your TeasMade develop an attitude problem and spit boiling water at you, a disgruntled traffic light decide it would rather be a Christmas tree, your K9 start staining the carpets? Could A-life become so sophisticated that, like the Replicants in *Blade Runner*, it starts resenting you for being more real than it? And could it start building bigger, meaner new generations of itself, as in the *Terminator* movies?

Dr David Cliff thinks not. "The prospect of machines which have the capability to build copies of themselves as machines is probably not impossible in principle – there's no theoretical reason why it couldn't be done – but the complexity of achieving that is just outrageous. The cost of doing so would require some billionaire to just go mad and devote all of his or her funds to the task, and even then they would probably fail. I think for economic reasons it's unlikely ever to happen. Maybe someone will build one out of Lego to prove that it can be done."

And what about the hostility? "In almost all serious science or engineering," continues Dr Cliff, "there are professional bodies who have codes of ethics and standards of practice which require that the output from the science or engineering is not dangerous to humans. This is why engineers build bridges that don't fall down, or every now and then people build space rockets that don't blow up. Occasionally, things go wrong," he concludes, "but I think self-replicating machines that are talked about in some nightmarish science fiction fantasy will always remain fantasy."

Besides, there's no more reason to expect a dystopia than an A-life utopia. For every hostile Replicant, there's an R2D2 or C3PO. SF visions are full of brilliant, kind and noble artificial beings, from *Star Trek's* Data to Isaac Asimov's Adam Link to... well, *Red Dwarf's* Kryten.

Okay, so perhaps there's room for a *little* scepticism...

MF

Read about it

Check out Steven Levy's *Artificial Life: The Quest for a New Creation*. Published by Penguin and costing £8.99, this will give you the entire rundown on the history of artificial life. Theories from the top minds are discussed, alongside the most important software.

finished work in 3D. Hence, A-life techniques were used to generate the animation.

In another work entitled *Primordial Dance*, Sims used a technique called artificial evolution to generate textures, colours and shapes. Like Latham, Sims defines the system and breeds the artwork, keeping the pieces he likes to breed with others, killing off the ones he doesn't. More recently, Sims has been working on installations which you can use to breed artwork by choosing your favourite pieces from big screens then employing them to propagate new generations.

Of equal interest are the patterns created by flocking birds, herds or antelope on the hoof and the glidingly fluid motion of large schools of fish. Simulations of these natural phenomena have been achieved using software models such as *Boids*, by Craig Reynolds. In this program, each individual fish or bird observes and reacts to the other individuals around it as well as the environment. Combined, the members of the flock determine its overall behaviour as they fly around obstacles, change direction and react to other flocks.

The effects can be compelling and have been used to animate group behaviour in movies. In *Batman Returns*, the movement of a flock of bats was choreographed using software like *Boids* rather than artist-drawn animation. Similarly, the stampeding wildebeest in Disney's *Lion King* moved according to a software model.

Karl Sims' *Primordial Dance* relied on a natural selection process to achieve its results.

Sims' *Panspermia* is a stunning animation where life blossoms and evolves on an alien world.

Panspermia was developed by planting the seeds for life within a computer and letting them grow.