Artificial life forms evolve basic intelligence

04 August 2010 by **Catherine Brahic**
Magazine issue 2772. **Subscribe and save**
For similar stories, visit the **The Human Brain** Topic Guide

Editorial: *Digital evolution and the meaning of life*

FOR generations, the Avidians have been cloning themselves quietly in a box. They're not perfect, but most of their mutations go unnoticed. Then something remarkable happens. One steps forward, and that changes everything. Tens of thousands of generations down the line, some of its descendents will evolve memory.

Avidians are not microbes, or sci-fi alien life forms. They are the digital offspring of Charles Ofria and colleagues at Michigan State University (MSU) in East Lansing. They "live" in a computer world called

Artificial life from a digital sea (Image: Gusto Images/SPL)

Avida, and replicate using strings of coded computer instructions instead of DNA. But in many ways they are similar to real life: they compete with each other for resources, replicate, mutate, and evolve. They - or things like them - might eventually evolve to become artificially intelligent life forms.

Similar to microbes, Avidians take up very little space, have short generation times, and can evolve new traits to out-compete their rivals. Unlike microbes, their evolution can be stopped at any time, reversed, repeated, and the precise sequence of mutations that led to the new trait can be dissected. "They're wonderful evolutionary pets," says Ben Kerr, a biologist at the University of Washington in Seattle.

> Avidians' evolution can be stopped, reversed and repeated. They make great evolutionary pets

They could become so much more. At the 12th annual international conference on artificial life in Odense, Denmark, this month, philosopher and computer scientist Robert Pennock of MSU will present the findings of experiments in which Avidians were made to evolve memory.
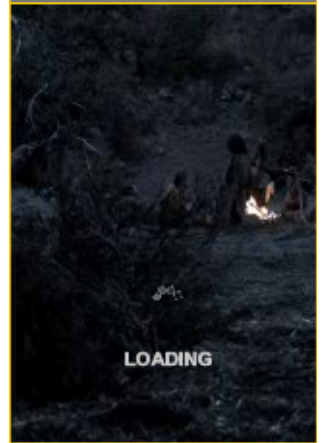
"The big question is: how did we get here? Our intelligence didn't evolve all at once," says

Pennock. "You need certain ingredients. Memory is one."

Experiments in Avida nearly always start with the simplest possible organisms, ones that can only clone themselves. To make them evolve, the experimenters release them into a competitive environment where the prize is an amount of "food" - aka processing time - which allows organisms to produce more clones.

In early memory experiments, Laura Grabowski, now at the University of Texas-Pan American, Edinburg, set up a food gradient in a computer environment made of a grid of cells. First-generation Avidians were placed at the low end of the gradient, in a cell that had minimal food. Straight ahead of them, however, lay a cell that had more.

The Avidians replicated themselves for nearly 100 generations, "living" and "dying" in the cell. Then one evolved a computer instruction to move forward. When it landed in an energy-richer cell, it reproduced more rapidly. Many thousands of generations later, some of its descendents were seen following the food gradient to its source, where concentrations were highest (*Artificial Life 2009*, p 92).

Even then the Avidians did not home in on the source. They stumbled their way along the gradient in zigzags, sensing the food and eventually reaching the source. They had evolved to ability to compare food in its current and past locations. "Doing this requires some rudimentary intelligence," says Pennock. "You have to be able to assess your situation, realise you're not going in the right direction, reorient, and then reassess."

Next, Grabowski sent a fresh batch of non-evolved Avidians on a treasure hunt. This time, cells contained a numerical code, which indicated in what direction the organisms should turn to find more food. But there was an additional twist to the task. Some cells contained the instruction "repeat what you did last time". The Avidians once more evolved into forms that could interpret and execute the instruction. "The environment sets up selective pressures so organisms are forced to come up with some kind of memory use - which is in fact what they do," says Grabowski.

This is not unlike evolution in living creatures, and the findings of the MSU computer scientists have attracted interest from biologists. "Laura's work suggests that the evolution of an ability to solve simple navigational problems depends on first evolving a simple short-term memory - and this in digital organisms that still don't exhibit something you would call learning," says Fred Dyer, an MSU zoologist who advised Grabowski. Dyer says this sort of insight would be all but impossible to obtain by studying biological systems.

But studies on complex behaviours in digital organisms don't just shed light on the evolution of organic life. They could be used to generate intelligent artificial life.

"In the past, the approach has been to start with high-level intelligence and reproduce that in a computer," says Grabowski. "This is the opposite. We're showing how complex traits like memory can be built from the bottom up, from things that are really very simple." To demonstrate this, Grabowski has evolved Avidians that move towards a light source. Her colleagues then translated the evolved "genome" into code that could control a Roomba robot. It worked: the Roomba was attracted to glowing light bulbs.

Starting simple is also what Jeff Clune, another member of the MSU dynasty, is interested in. In particular, he is focused on producing artificial brains that move robots. Clune works with a system called HyperNEAT, which uses principles of developmental biology to grow a large

number of digital neurons from a small number of instructions.

In nature, the location of a cell in an embryo often determines its function - whether it will become a heart cell or a neuron for instance. Similarly, in HyperNEAT, the location of each artificial neuron - given by coordinates - is plugged into a matrix of equations and the result defines what the cell's role will be.

This, says Clune, means that you can build complex brains from a relatively small number of computerised instructions, or "genes". In contrast, traditional neural networks have worked on a one-to-one principle: each cell in the network is encoded by a single instruction which is not re-used.

> You can build complex brains from a small number of computerised instructions or genes

It also means you can evolve brains that share structural properties with real brains. For instance, Clune has found that unlike old-school neural networks, brains evolved with HyperNEAT tend to be symmetrical and ordered - like real brains. His analysis of the networks shows this comes from having evolved symmetry and pattern-generating instructions right at the start of the series of instructions.

To test whether such brains actually perform better, Clune drops them into a virtual robot, which then has to perform a task like running across a flat surface. If the robot performs well, he selects that brain and evolves it further. As with Avidians, evolution involves copying the brain's "genes", and introducing random errors in the process to produce brains with slightly modified connections or instructions.

Clune's results, presented at the Genetic and Evolutionary Computation Conference in Portland, Oregon, last month, show that symmetrical, organised artificial brains tend to perform better at tasks like running than do non-HyperNEAT brains.

"Brains that have been evolved with HyperNEAT have millions of connections, yet still perform a task well, and that number could be pushed higher yet," he says. "This is a sea change for the field. Being able to evolve functional brains at this scale allows us to begin pushing the capabilities of artificial neural networks up, and opens up a path to evolving artificial brains that rival their natural counterparts."

"That is a lofty long-term goal, of course," he adds, "but this technology allows us to start marching towards it."

A history of life in silicon
Before Avida and before its predecessor Tierra there was Core Wars. Popular in the 1980s, the game pitted computer programmers against each other. The principle was simple: players would write computer programs that shut each other down and the last

one standing would win.

In the late 1980s, ecologist Thomas Ray, who is now at the University of Oklahoma in Norman, got wind of Core Wars and saw its potential for studying evolution. He built Tierra, a computerised world populated by self-replicating programs that could make errors as they reproduced.

When the cloned programs filled the memory space available to them, they began overwriting existing copies. Then things changed. The original program was 80 lines long, but after some time Ray saw a 79-line program appear, then a 78-line one. Gradually, to fit more copies in, the programs trimmed their own code, one line at a time. Then one emerged that was 45 lines long. It had eliminated its copy instruction, and replaced it with a shorter piece of code that allowed it to hijack the copying code of a longer program. Digital evolvers had arrived, and a virus was born.

Avida is Tierra's rightful successor. Its environment can be made far more complex, it allows for more flexibility and more analysis, and - crucially - its organisms can't use each other's code. That makes them more life-like than the inhabitants of Tierra.

New Scientist

Not just a website!

Subscribe to New Scientist and get:

New Scientist magazine delivered to your door

Unlimited online access to articles from over 500 back issues

Subscribe Now and Save

2
tweets

retweet   Like 2   PRINT   SEND   SHARE

If you would like **to reuse any content** from New Scientist, either in print or online, please **contact the syndication** department first for permission. New Scientist does not own rights to photos, but there are a variety of licensing options available for use of articles and graphics we own the copyright to.

Back to article

PRINT   SEND   SHARE