

side is holoprosencephaly — a disease associated with mutations in the human version of the fruitfly's *hedgehog* gene that inactivate it⁴⁻⁶. The Hedgehog pathway is important in the development of the central nervous system and face, and holoprosencephaly stems from a failure of the embryonic brain and face to divide into symmetric halves. An extreme characteristic of this disorder is cyclopia, in which the two eyes are fused into one central structure and the remaining facial features and brain are rudimentary.

Like many other birth defects, cyclopia can be caused by either intrinsic genetic defects or the effects of environmental substances on genetically normal embryos. One such substance was discovered through epidemiological investigations of cyclopia in sheep herds of the western United States⁷. Sheep cyclopia was found to result when pregnant females ate a lily (*Veratrum californicum*; Fig. 1), and the relevant chemical from the lily was dubbed 'cyclopamine'. Although cyclopamine has no apparent effect on adult animals, it consistently leads to severe holoprosencephaly in developing embryos. As this condition can be caused by either mutations in human *hedgehog* or exposure to cyclopamine, it was suggested^{8,9} that cyclopamine acts by repressing the Hedgehog pathway. Taipale *et al.*¹ wondered whether cyclopamine might therefore be effective in treating basal-cell carcinomas.

For a compound to be effective against these cancers, it must switch off the Hedgehog pathway at a point downstream from the molecular defect — for example, downstream of mutant Patched protein. Taipale *et al.* now show that cyclopamine reverses activation of the pathway downstream of Patched and upstream of GLI. In fact, the evidence points to Smoothened as the site of the chemical's action.

The authors¹ also show that cyclopamine may be useful for treating patients with basal-cell carcinomas. These tumours grow slowly and are difficult to establish in culture. As a substitute, Taipale *et al.* used cells that were genetically engineered to lack a functional copy of *patched*. Using doses of cyclopamine that do not affect normal cells, they found that the genetically engineered cells stopped growing and that several of their malignant characteristics were reversed.

Together with the epidemiological finding that adult sheep do not suffer ill effects of cyclopamine, these results¹ are encouraging. But they do not directly address the question of whether switching off the Hedgehog pathway will cure basal-cell carcinomas. These tumours, like most other human cancers, have mutations in many genes, and it is probably the accumulation of mutations that results in malignancy. Perhaps the Hedgehog pathway must be activated to get tumour development started, but it is not certain that it needs to be kept activated to maintain

a fully malignant state. Indeed, in naturally occurring basal-cell carcinomas, shutting off this pathway might be like shutting the barn door after the horse has bolted. Treating basal-cell carcinomas with cyclopamine might be predicted to suppress the development of tumour cells and perhaps retard their growth, but not necessarily to kill them. There is some evidence that tumour cells that have been suppressed in this way might undergo programmed cell death or be eradicated by the body's natural defences. If not, then tumours treated with cyclopamine may eventually reappear when the drug is removed.

Fortunately, there is a good mouse model for Gorlin's syndrome that will help us to answer questions about the effects of

cyclopamine on tumours *in vivo*. If this drug then progresses to clinical trials, investigators would be wise to remember its source, and to avoid women of child-bearing age in their studies. ■

Allen E. Bale is in the Department of Genetics, Yale University School of Medicine, 333 Cedar Street, PO Box 208005, New Haven, Connecticut 06520-8005, USA.

e-mail: allen.bale@yale.edu

1. Taipale, J. *et al.* *Nature* **406**, 1005–1009 (2000).
2. Hahn, H. *et al.* *Cell* **85**, 841–851 (1996).
3. Sidransky, D. *Nature Genet.* **14**, 7–8 (1996).
4. Chiang, C. *et al.* *Nature* **383**, 407–413 (1996).
5. Belloni, E. *et al.* *Nature Genet.* **14**, 353–356 (1996).
6. Roessler, E. *et al.* *Nature Genet.* **14**, 357–360 (1996).
7. Keeler, R. F. & Binns, W. *Teratology* **1**, 5–10 (1968).
8. Cooper, M. K. *et al.* *Science* **280**, 1603–1607 (1998).
9. Incardona, J. P., Gaffield, W., Kapur, R. P. & Roelink, H. *Development* **125**, 3553–3562 (1998).

Artificial life

From robot dreams to reality

Rodney Brooks


Artificial life is a diverse field of research, but a common theme is teasing out the fundamental principles of life by building detailed working models. One of the most ambitious goals of artificial-life research is the construction of living systems out of non-living parts. Most of the artificial systems built remain strictly inside a computer, but on page 974 of this issue¹ Hod Lipson and Jordan Pollack take a first step towards bridging the gap between computer models and physical reality. They describe a system that evolves locomotive machines inside a computer and then automatically builds them, using rapid-prototyping technology, so that they can move around in the physical world.

Over the past 20 years, computer algorithms inspired by genetics² (genetic algorithms) have become common tools in

solving optimization problems. Usually the optimization target is a mathematical procedure of known form, but with many undetermined parameters. A population of candidate procedures is generated by expressing different sets of these parameters as binary strings. The performance of each procedure is evaluated and those that are 'fitter' or do better according to some specific criteria are allowed to reproduce. Reproduction may be sexual, by taking the strings of two fit parent procedures and generating the binary string for the descendant by using cross-over and perhaps mutation, or it may be asexual, in which case just mutation is used. Over many generations, the population tends towards more successful procedures.

Over the past decade many people have experimented with evolving populations of 'artificial creatures' in simulated environ-

Figure 1 A robot 'evolved' by Lipson and Pollack¹ to produce horizontal motion. The body parts and control circuits are evolved inside a computer — thousands of robots over thousands of generations — and then rapid-prototyping technology is used to turn them into reality. Some of the winning designs are surprisingly symmetrical, which may be explained



by symmetrical machines finding it easier to move in straight lines. This particular machine uses antiphase synchronization to move — while the upper two limbs push the machine forwards, the central body is retracted, and vice versa. Movies of this robot and others are available at <http://www.nature.com> as Supplementary Information¹.

ments or virtual worlds. Here, rather than the binary string representing the parameters of some procedure, the string is an artificial genome, which encodes a control circuit (or nervous system) for a simulated robot. So, over time, better-performing robots slowly emerge. In some cases there is an implicit fitness evaluation as creatures fight it out for virtual resources necessary for survival. In other cases there is an explicit 'fitness function' applied to each generation of creatures, forcing evolution in a desired direction.

There are also experiments in transferring robot control systems evolved inside a computer onto physical versions of the simulated robots. But there is much debate and conflicting evidence over how well these transfer experiments work. In some cases³, the 10,000 or so fitness evaluations over tens to hundreds of generations for a population size of ~100 have all been done on physical robots — but such experimental tenacity is understandably rare.

A glaring omission from most of these experiments is that the 'body' of the robot is usually considered to be constant while just the nervous system evolves. An inspiring exception is the work of Karl Sims⁴. He evolved creatures that could swim and crawl in a virtual world that follows the laws of newtonian physics. His fitness function rewarded horizontal motion, and out popped creatures with locomotive ability, although some tweaking was required to get everything to work as planned.

An early version of the fitness function did not penalize vertical motion and was only applied to a few seconds of existence, whereupon really tall creatures evolved that were good at falling over and even tumbling. For a while, creatures evolved that moved along by beating their bodies with their limbs — they were taking advantage of a bug in the part of the simulated physics that encoded conservation of momentum. Sims' system coevolved nervous systems and body plans, but his creatures were all purely computational.

A couple of years ago Paolo Funes and Jordan Pollack⁵ tried to convert computer models into physical reality by evolving not creatures, just structures — the simulated structures were selected for their strength. They then built physical versions of the structures by hand from real Lego blocks and confirmed that the structures were much stronger than human-designed ones.

Lipson and Pollack now take this idea a step further. They evolve locomotive systems in computational space and use rapid-prototyping technology to automatically produce multi-linked structures that only need motors to be snapped on by hand. The successful designs evolved surprisingly different ways of generating motion, but there is reasonable correlation between the predicted locomotive abilities of the models and those measured from the physical robots (Fig. 1).

The evolution in Lipson and Pollack's experiments goes on within their computer — there is no fitness evaluation in the physical world. And the computational parts of the robots stay within that same computer even when they have been physically built. This means there can be no feedback from the physical world into the evolutionary process. At best, this system is like a virus that uses other more complex machines (which in this case are not life-forms themselves) to carry out reproduction. There is some way to go before self-reproducing robots can exist in the real world. But for artificial-life researchers there are aspects of the present study that satisfyingly resonate with the ways in which living systems develop.

First, these particular robots cannot be built by conventional manufacturing techniques. The rapid-prototyping technology solidifies polymers in place, so that ball and socket joints are constructed with the ball already inside the socket. The parts are never separate, and if they were they could not be assembled without damaging them. This is not far removed from the way that biological systems grow. Second, the evolutionary strategy used in these experiments starts with a blank or 'null' genome and randomly mutates it into one that generates a working

machine — so there is no in-built bias from seed machines in the population. One could say these machines have evolved 'naturally', without human intervention.

People have long thought about building technology from non-biological materials that can reproduce themselves — the US space agency NASA convened a panel in the 1960s to investigate the possibility of seeding the Moon with a small self-reproducing factory. Although we are still a long way from that goal, Lipson and Pollack have at last demonstrated a computational system that designs functional machines and builds them with almost no human intervention. The resulting machines cannot match the complexity of the rapid-prototyping machine designed by human engineers that is required to do the actual fabrication. Nevertheless, this is a long awaited and necessary step towards the ultimate dream of self-evolving machines. ■

Rodney Brooks is in the MIT Artificial Intelligence Lab, 545 Technology Square, Cambridge, Massachusetts 02139, USA.

e-mail: brooks@ai.mit.edu

1. Lipson, H. & Pollack, J. B. *Nature* **406**, 974–978 (2000).
2. Holland, J. *Adaptation in Natural and Artificial Systems* (Univ. Michigan Press, 1975).
3. Cliff, D. *et al. Adapt. Behav.* **2**, 73–110 (1993).
4. Sims, K. *Artif. Life* **1**, 353–372 (1994).
5. Funes, P. & Pollack, J. *Artif. Life* **4**, 337–357 (1998).

Bacterial genomics

Pump up the versatility

E. Peter Greenberg

On page 959 of this issue, Stover and colleagues¹ publish the genome sequence of the resilient bacterium *Pseudomonas aeruginosa*. This microbe is feared by every patient with the genetic disorder cystic fibrosis because it colonizes the lungs of most people with this disease. Once established, it slowly but surely causes more and more damage to the lungs, eventually leading to the patient's death. This opportunistic pathogen also infects people whose immune system is compromised in any way. It rarely infects uncompromised hosts, and occurs naturally almost everywhere — in lakes, streams, soil and even our drinking-water supply. *P. aeruginosa* (Fig. 1, overleaf) shows great nutritional versatility; like a goat, it can live on almost anything.

Traditional approaches have not provided us with the tools needed to fight this bacterium effectively. For reasons that are not clear, *P. aeruginosa* infections resist treatment with antibiotics. Hence the hope that the *P. aeruginosa* genome sequence¹ will reveal new ways of tackling this organism.

The completed sequence, described by Stover and colleagues, is the result of a unique collaboration between a group of

academic investigators, a pharmaceutical company and a charitable foundation, the US Cystic Fibrosis Foundation. The project, which began three years ago, was financed jointly by the Cystic Fibrosis Foundation and the company PathoGenesis, and has proven to be a model of scientific integrity and interaction.

Academics at the University of Washington in Seattle sequenced and analysed the genome, using the 'shotgun' sequencing and compiling strategy developed by Craig Venter and colleagues². Periodically, the available sequence data were posted on the Internet³. Scientists at PathoGenesis — with the help of 61 *P. aeruginosa* experts from around the world — then 'annotated' the genome. What this means is that they picked out sequences in the genome that were likely to be genes and used bioinformatics to compare these predicted genes with known genes from *P. aeruginosa* and other species. This allowed them to work out what the functions of the *P. aeruginosa* genes might be.

The collaborators hoped to gain insights into the basic biology of *P. aeruginosa*, to find targets for drug development, and to spur interest among scientists in studying this