

Christopher MacLeod and **Grant M Maxwell** from the Robert Gordon University in Aberdeen, Scotland, explain how evolutionary methods encountered in nature help build electronic circuits

Minds for Robots

IN THE 1950s and 60s, popular culture was entranced by robots. There was Robby in “Forbidden Planet”, Gort in “The Day the Earth Stood Still” and many others. This fascination has continued to the present day, only the names have changed; now we have Commander Data of “StarTrek” and the NS-5 of “I Robot”.

Yet, despite the interest and the obvious advantages of having intelligent machines to do dirty, boring or dangerous jobs, we are little nearer to creating a similar technology in the real world. This article outlines a new route towards such intelligence in machines – Incremental Evolution.

Approaches to AI

Historically there have been three approaches to creating intelligent machines. The first approach was to try and copy the biological brain. However, this contains billions of nerve cells (called neurons), interlinked by trillions of connections, and unravelling it is well beyond our current scientific capabilities.

The second approach, known as symbolic AI (artificial intelligence), was to build a large database of rules and instructions to allow the machine to make complex decisions. But we know that biological brains don’t work like this. Simple rules also don’t have the flexibility to work in real-world situations and, in any case, the database soon grows to unmanageable proportions.

Finally, there is the approach of modelling simple neural circuitry using “Artificial Neural Networks” – networks of simple processors, loosely based on real nerve cells. However, the resulting circuits are not complex or flexible enough to allow intelligence to emerge.

Artificial Evolution

The problems outlined above led researchers to consider the mechanism that allowed nature to develop the biological brain – evolution. Evolution works because the fittest individuals in a society tend to survive better and therefore have a higher chance of breeding and passing on their traits to the next generation. As such, over many generations the fitness of the whole population increases.

Engineers have designed artificial versions of evolution called Evolutionary Algorithms. We can illustrate how these work using an example. Suppose we want to design a filter as shown in **Figure 1**.




Figure 1: A simple filter circuit

The circuit can be represented by a string of numbers which are its component values. **Figure 2** shows a circuit with its components and the string which represents it.

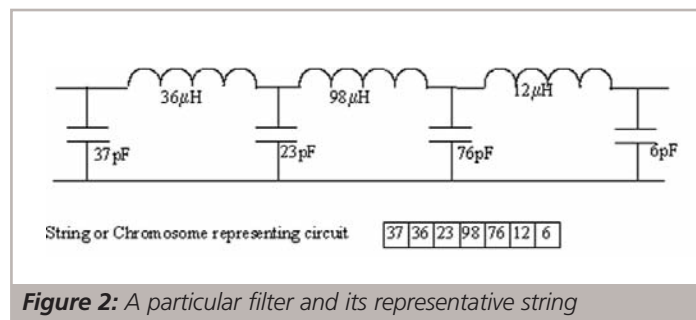


Figure 2: A particular filter and its representative string

We can generate lots of such strings (also sometimes called ‘chromosomes’) representing circuits like this – each chromosome is the component values of a single circuit. The starting point of an Evolutionary Algorithm is a number of such strings (called the ‘population’) filled with random numbers (random component values).

We then test each of these random strings to see how well it performs; in this case by comparing how good the circuit is with the perfect response. Having done this, we get rid of the worst strings; if we have a population of 99 strings say, we’d delete perhaps two thirds of them – the two thirds that are furthest from the ideal response. This would leave us with 33 strings. Next is to ‘breed’ them – just as in a population of animals, the fittest survive and breed together.

The breeding process is called ‘crossover’ or ‘recombination’. It works like this: take two random strings and select a random point on them, then cross over all the numbers to the right of that point into the opposite string and vice versa, **Figure 3** shows the idea.

Before breeding:
12 54 34 65 03 87 67
37 36 23 98 76 12 06

After breeding:
12 54 34 98 76 12 06
37 36 23 65 03 87 67

Figure 3: Crossover, also called 'recombination'

Hopefully, in some of these events, the new string will inherit some good traits from the mother string and some different ones from the father, and will therefore be better than both.

Once this is done, the original population of 99 is recreated (the 66 new strings being the result of breeding). A small percentage of the numbers are then changed at random to add some variation into the population, an operation called 'mutation' and we start the whole process all over again: evaluation, culling, breeding and mutation. Over many generations of doing this, the population will improve to the point where some of the strings can solve the design problem satisfactorily. The technique is particularly useful in circuits which don't have good design rules or where compromises have to be made in the design.

Problems

The procedure outlined here has been shown to work well in many problems. It has been used to design analogue and digital circuits as well as DSP systems and artificial neural networks.

However, when we try and use it to design circuits to control intelligent machines, it has only limited success. The main reason for this is that the algorithm is inefficient when applied to large circuits. This is because of the many different combinations of components and wiring possible in such circuits; we say that the circuits have a large 'search space'. So what was successful for small problems is unworkable for large ones.

A related problem is that it is difficult to design the algorithm so that it can evolve a circuit which deals with many different kinds of input and output. For example, in a robot, there may be inputs from vision sensors, bump sensors and audio sensors. If you try and integrate these together, they sometimes conflict and interfere. This is why the human brain is not one big network, but lots of little ones (sometimes called modules), all doing their own thing but co-operating together.

How Nature Does It

Although current Evolutionary Algorithms model the selection of good traits within a population, this is only one feature of natural evolution. If we look at the evolution of organisms over time, the other important aspect is 'growth in complexity', where animals have changed from simple to complex forms over time, driven by competition.

When life started on Earth it was very simple – just single cells, in fact. From then until now, it has slowly but surely become more complex. Single-celled forms were replaced by simple multicellular animals, which eventually developed into jellyfish, then worms, complex invertebrates, fish, amphibians, reptiles and finally mammals – a march of progress from simple to complex. The extraordinary animals of today are only possible because they built on the successes of the past, adding more functionality to those which had gone before.

So, it is not just about blind growth, but about building upon a previously laid, firm foundation. The past triumphs of evolution are not

discarded by nature, but new developments are built on top. This is particularly obvious in the brain, were it would have been impossible to rewire billions of neurons and trillions of connections each time it developed; the new changes had to build on top of the older ones. In fact, Richard Dawkins, in his famous book on evolution "The Blind Watchmaker" points out that incremental change is the only realistic way to evolve complex animals.

In other words, a traditional Evolutionary Algorithm simulates the selection of good traits in a population of (say) amoebas, but the amoebas can never evolve into jellyfish because the strings which represent the population members cannot become larger and more complex – and even if they did, the algorithm could not handle the extra complexity of the search-space generated.

Incremental Evolution

Incremental Evolution adds the idea of growth in complexity into Evolutionary Algorithms. We can illustrate how it works by considering the evolution of a circuit.

The algorithm works by initially making the circuit as simple as possible. An ordinary Evolutionary Algorithm is then used to evolve this simple system until it performs its task perfectly (or cannot get any better). This simple evolved circuit is then fixed and not allowed to change again.

Next, a 'module' of new components is added to the circuit; these are added to the pre-existing structure and are allowed to evolve. Again, when they can't get any better, the whole is permanently fixed and another new module is added. **Figure 4** shows this process operating on a filter.

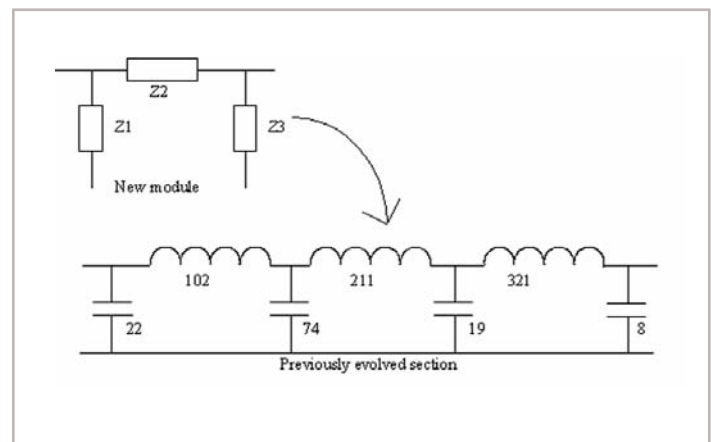


Figure 4: A new module being added to a pre-evolved section

A better understanding of this can be had by considering the evolution of a radio circuit. If we were to take dozens of transistors and other components and use a traditional Evolutionary Algorithm to design a radio receiver, it wouldn't get very far; the circuit is too complex and the search space too large.

However, if we only allowed one transistor and a few components, chances are that it would evolve a simple detector circuit, rather like a crystal set. Now, this doesn't work very well, but it does demodulate signals. If we now fixed this and allowed the algorithm to add a module, we have a good chance of evolving an RF amplifier in front of the detector. Fix all this and we might get an AF amplifier behind it. Try again and you might get a selection filter to tune it at the very front.

So, by allowing the system to grow like this, we can evolve much more complex systems than a simple Evolutionary Algorithm can – and systems which are automatically modular; for example, the filter in Figure 4 can evolve to give a much more complex multi-stage response than is likely to be produced by a single stage of evolution.

In fact, it turns out that in large systems with many connections, the search space of a traditional Evolutionary Algorithm grows frighteningly quickly, whereas in Incremental Evolution, the search space is simply that of the added module.

Robots

In our lab at the Robert Gordon University in Aberdeen, UK, Incremental Evolution was used to evolve a legged robot with vision sensors. The robot's sequence of evolution is shown in Figure 5.




Figure 5: An evolving legged robot

An important principle can be seen in operation here. The robot's "body" and "mind" must evolve together; the principle of Incremental Evolution applies to the whole system, not just the circuit which controls it. This, again, keeps the search space in each iteration small; you might call it a holistic form of Artificial Evolution. After all, in nature, the brain of jellyfish goes with its particular simple body plan. In the terms of Evolutionary Algorithms, we say that the "fitness function" changes along with the current evolutionary stage.

The robot started off as simple as possible, in this case as a robotic "mudskipper". All it was able to do was pull itself along the floor with two simple inflexible legs. A neural network controller (a control circuit) was evolved to allow it to do this. Next, the robot was given more complex legs, with two active degrees of freedom (a "hip" and a "knee"). Keeping the previously evolved network, another was evolved to control these new functions; the old network controlled the "hip" and the new, added one, the "knee".

Having done this, another two legs were then added, making the robot into a quadruped. Again, what was previously evolved was kept and further modules added on top of the old to control the new legs (although the

algorithm was allowed to "cut and paste" previously evolved modules into new positions).

Having got our walking four legged robot, a vision system was evolved for it. This started off as just as one pixel and was allowed to evolve slowly, using the principles above, into a five by five pixel "robotic eye", which could recognise simple patterns as shown in Figure 6.

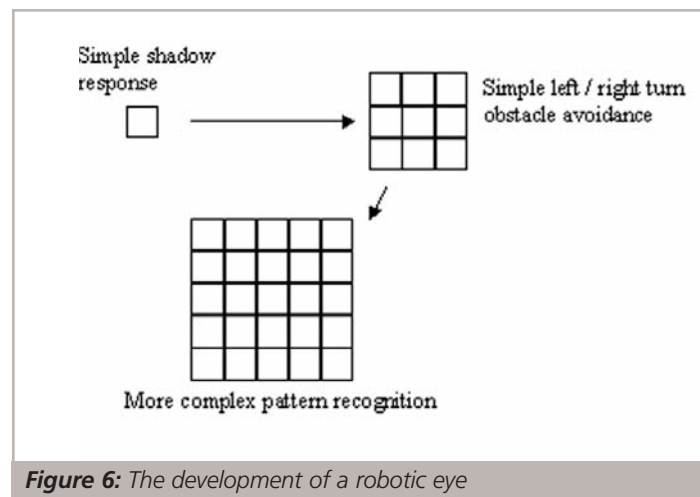


Figure 6: The development of a robotic eye

New networks were then added to interface those of the legs with the eye and produce a robot which could recognise patterns, turn and alter its gait accordingly. The final robot had over two hundred neurons in twenty two modules and could operate successfully in a variety of environments.

Figure 7 shows how one particular function evolved (in this case, the controller for a leg) as new modules were added.

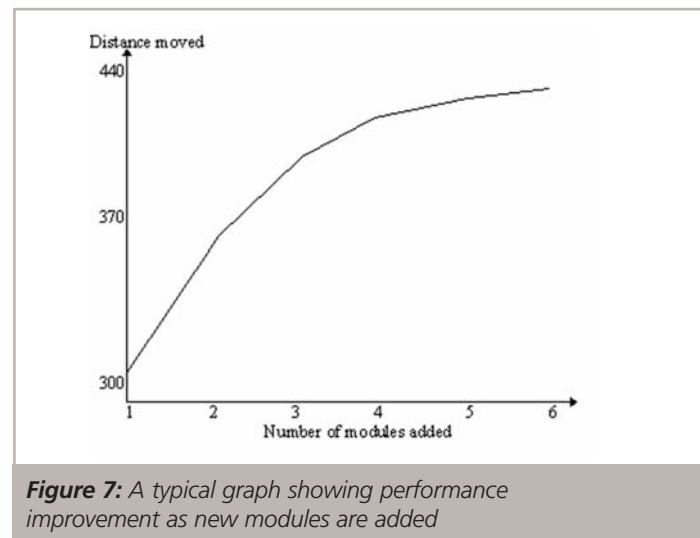


Figure 7: A typical graph showing performance improvement as new modules are added

The result of all this, is that, not only can a much more complex system be evolved, but that it is modular with parts of the network handling different functions. Hence, it has the ability to connect and integrate different sensor and actuator functions together. This modular network, however, is part of a fully functioning and co-operating whole.

Rules of Evolution

During the process of evolving the robot, it was discovered that there are some important "ground rules" for making the method work. If these are ignored then the evolution will stall and the system won't reach its full potential.

Firstly, the components in each module added must be as flexible as

possible, in the case of electronics the algorithm should be allowed to choose these to be any component (resistor, capacitor, inductor, etc). In the case of a neural network, the flexibility of the neurons must be as high as we can make them.

Secondly, the algorithm evolving the system must be able to choose which connections are made as well as the component values and types. It was found that if this wasn't allowed to happen, the connections interfered with each other and stopped the system reaching its full potential.

Finally, usually there are a minimum number of components needed in each new module, if there are not enough then the system won't evolve correctly. This is similar to the first and second points, which indicate that the functionality of the added modules is critical and that each module must be as flexible as it can be.

It's also worth noting that even when the rules are followed it is possible for evolutionary "dead ends" to be reached, where further improvement is not possible. This also happens in the world of nature. In this case, modules must be removed and the system allowed to evolve again but in a different direction.

Applications

There are also many possible applications of Incremental Evolution in other areas of engineering. One of the most interesting is in biomedical engineering, in the development of prosthetic limbs. In this case, there is a direct incremental route from a simple system with all the joints of the limb fixed in place, to a complex system with all the joints free and controlled. As each joint is freed the controlling circuitry can be allowed to develop by adding to its structure, in the same way as in the robot.

Other examples include mechanical engineering; for example in

aerospace engineering. Here too it is fairly obvious that sections of an aerodynamic design can be added incrementally.

There may even be lessons which biologists can learn about the early development of life and what works and doesn't in terms of evolutionary mechanisms.

Today we try and design complex systems inspired by biology, like intelligent robots, by analysing how humans and animals work and copying them. But when we see a human, we see four-and-a-half billion years of evolutionary progress. Each stage of that journey was built on the previous stage. Humankind did not come into being suddenly, without warning. It began its long trek as a single-celled organism in the primeval seas. We, our behaviour and our physical makeup, are a result of that development.

What, therefore, makes us think that we can build sophisticated intelligent systems like robots from scratch? Surely they too must build up their bodies and their behaviour in a similar way, just as we did, piece by piece, one step at a time. ■

FURTHER READING

www.rgu.ac.uk/eng/compint/page.cfm?page=12909

- S Muthuraman, C MacLeod and G Maxwell, "The development of Modular Evolutionary Artificial Neural Networks for Quadruped Locomotion", *IASTED ACS 03*, p268-273.
- S Muthuraman, G Maxwell, C MacLeod, "The Evolution of Modular Artificial Neural Networks for Legged Robot Control", *Artificial Neural Networks and Neural - Information Processing*, Springer (LNCS 2714), 2003, p488-495.
- Sethuraman Muthuraman, "The Evolution of Modular Artificial Neural Networks", PhD Thesis, The Robert Gordon University, 2005

PCB-POOL®

SERVICING YOUR COMPLETE PROTOTYPE NEEDS

1 EUROCARD
(160 x 100 mm)
+ Tooling
+ Photoplots
+ VAT

€49

Price example
Any size and contour possible!

Optional:

- Soldermask
- Fast-turnaround
- Silkscreen
- 4-Layer Multilayer
- 6-Layer Multilayer

DOWNLOAD OUR FREE LAYOUT SOFTWARE!

Over **15.000** Customers

LEAD FREE
100% OF THE TIME

Beta

Freephone **0800-3898560**

Tel: +353 (0)61 701170
Fax: +353 (0)61 701145
pcb-pool@beta-layout.com

PCAD CAMBRIDGE PRETEL EDWIN ORCAD CircuitsCode PHOTOLITH Electronics Easy-PC Sprint

Simply send your files and order ONLINE:

PCB-POOL.COM

Thermostats for UPS Battery Back-Up Systems

When a power-cut hits a company, they need to be sure that their UPS system performs immediately and for as long as possible until mains power is re-established.

The life of a battery back-up system is partly dependant on its stable temperature control which is often achieved using electronic means. As an alternative to electronic sensors, the Matsuo MQT thermostat offers UPS system installers a simpler and more cost-effective solution.

Matsuo thermostats have a unique twin-bimetal system that creates a control which is capable of switching over 1 million cycles. Accuracies of $\pm 1.5K$ and differentials down to $3K \pm 1K$ are readily available, with repeatability as small as $\pm 0.2K$. Basically they function on a par with electronic controllers.

The thermostat's IP64 plastic housing can be readily fixed to the side of a battery and so each one can be safely monitored.

Call us today for free evaluation samples and application advice.

ATC Semitec Ltd
Tel: 01606 871680 • Fax: 01606 872938
E-mail: sales@atcsemitec.co.uk
Web: www.atcsemitec.co.uk