Genetic Mountaineering

Adrian Woolfson

A New Kind of Science by Stephen Wolfram
Wolfram Media, 1197 pp, £40.00, May 2002, ISBN 1 57955 008 8

One of the most intriguing of all magic tricks, the Disappearing Handkerchiefs, was presented to King Louis-Philippe at the Château St-Cloud in 1846 by the renowned French magician Robert-Houdin. An account can be found in his Memoirs:

> I borrowed from my noble spectators several handkerchiefs which I made into a parcel and laid on the table. Then at my request different persons wrote on the cards the names of the places whither they desired their handkerchiefs to be invisibly transported. When this had been done, I begged the King to take three of the cards at hazard and choose from them the place he might consider most suitable.

One of the three cards picked by the King suggested that the handkerchiefs should be placed beneath the candelabra on the mantelpiece, the second that they should be transported to the dome of the Invalides, and the third that they should materialise in a chest buried under a particular orange tree in the grounds of the château. The King chose the orange tree, and sure enough a small chest smelling strongly of sulphur was found at its base. When Louis-Philippe opened it with a rusty key fastened to a turtledove, he found a piece of parchment that read: 'This day, 6 June 1786, this iron box, containing six handkerchiefs, was placed among the roots of an orange tree by me, Balsamo, Count of Cagliostro, to serve in the performance of an act of magic, which will be executed on the same day sixty years hence before Louis-Philippe of Orléans and his family.' 'There is decidedly witchcraft about this,' the King is reported to have said, on tearing open the enclosed envelope and displaying the handkerchiefs to the astonished spectators. The simple trick underlying this apparently complex performance was, of course, to predict which hiding place the King was likely to choose from the small selection offered him and then to make the necessary arrangements in advance.

The complexity of living things, both extant and extinct, was achieved by means of similar tricks: great feats, which on first inspection appear difficult or even impossible to achieve, but are nevertheless accomplished by simple means. This is the key to evolution: once the underlying mechanistic 'tricks' have been determined, the mystery of biological form and function is made rational and comprehensible.

That simplicity begets complexity

should not be surprising – how else could life have originated in the first place? But our understanding of how evolution accomplishes its conjuring tricks continues to mature. While Darwin's foundations were easily built on to accommodate genes and mutations, recent work has been more challenging. Stephen Jay Gould and Richard Lewontin, for example, used the spandrels of St Mark's in Venice, which exist as a necessary by-product of the process of mounting a dome on rounded arches, as a way of illustrating the anti-adaptationist idea that certain features of organisms have arisen independently of natural selection. The spandrels reflect the constraints of building a particular type of structure in a particular type of way and are, as a result, often devoid of function.

Then there is the work of complexity theorists, who argue that there are rich seams of non-programmed, self-organising and self-assembling order in the natural world that offer natural selection a helping hand by generating complex patterns without the intervention of genes. Evolution, from this perspective, is driven by an intricate interaction between natural selection, the programmed order of genes and the non-programmed 'order for free', as Stuart Kauffman has called it, generated within certain types of complex system. Historical contingency, too, has a role in the new synthesis. Simon Conway Morris has argued that if the Earth hadn't been hit by the meteorite that probably made the dinosaurs extinct, the history of life would have been quite different.

One of the most important recent insights, however, is that the capacity for evolution – 'evolvability' – has itself evolved. It is not enough for a feature to be adaptive at a given time and place. Systems must be constantly poised for change: a potential for evolvability is essential if they are to respond to the vagaries of an uncertain future. Evolvability can be understood in relation to another concept, that of an evolutionary 'landscape'. Each of an organism's characteristics, be it macroscopic (the size of a camel's hump, the shape of a bat, the length of a rhino's horn), microscopic (a feature of a neurone in the brain) or a biochemical property (such as the potency of an enzyme), may be imagined as occupying a discrete point on a metaphorical landscape that contains the array of all possible variants of that characteristic. In the case of a camel's hump, for example, the landscape would contain humps the size of the Post Office Tower, mini-humps, wobbly humps, furry humps, hairy humps, smooth humps, spiny humps and so on. We might call this landscape the 'space of all possible camel humps'. (More accurately, it is a 'genescape': the collection of genetic sequences that encode and correspond to all the variants of the physical structure in question.) When hump size or shape evolves in a population of camels, the evolutionary process explores the 'hump landscape', searching for 'better' or more adaptive humps. Similarly, the hormone insulin may be thought of as existing in a landscape containing all possible insulin molecules. Some of these would be hopelessly inefficient – the use of them would result in diabetes – whereas others would be more efficient than those with which we are familiar.

Taking the landscape metaphor still further, the most adaptive or 'optimal' variants will occupy the peaks or mountainous regions of the landscape, with sub-optimal variants in the plains and valleys. In the new synthesis, natural selection is envisaged as tuning and tweaking the dynamics of evolving systems so as to ensure that the adaptive landscapes they explore are 'correlated' rather than random. In a correlated insulin landscape, the greater majority of the efficient molecules would be clustered together in a few connected mountainous regions, so that changes to the genetic sequence encoding one molecule would most likely lead to mutants that were at least as efficient as that molecule, if not more so. Landscapes of this sort are 'evolution-friendly', as they enable populations to skip from one variant of a characteristic to another while retaining or increasing their competitive edge. If, on the other hand, the insulin landscape were uncorrelated, the efficient insulin molecules would be scattered randomly across it. Routes leading to more adaptive forms would be blocked by unforgiving valleys populated by variants that were inefficient and functionally inept. Such landscapes are unnavigable by the natural evolutionary exploratory mechanisms – gene mutation, duplication, conversion and recombination – and trap features in states of local optimality rather than allowing them to converge on the peaks that represent the best of all possible global options.

Studies on flies suggest that certain proteins function as molecular capacitors, allowing mutations to accumulate without immediately showing their effects. This type of trick may facilitate evolution on sub-optimal landscapes, by allowing unfavourable phenotypes to be passed over without restructuring of the landscape. At the molecular level, there are enzymes that copy and repair DNA in an error-prone manner, which helps generate the store of mutations on which evolution depends. Genetic studies have divulged other tricks. The genes that control the shape and location of body parts, for example, are highly conserved between one species and another, indicating that the same small set of genes can generate forms as diverse as penguins, polar bears, porcupines and pythons.

Despite its many successes, the new synthesis has to some extent been hampered by the mathematical intractability of the complex natural structures at the heart of its agenda: even the most modern biological approaches are entrenched in mathematical methods developed more than three centuries ago. In A New Kind of Science, Stephen Wolfram argues that computer programs provide much better models for natural systems. Given that the extraordinary complexity of even the simplest biological system appears to preclude traditional equation-based descriptions, the possible consequences of perceiving the behaviour of such systems in computational terms are immense. Wolfram champions a new experimental methodology based on very simple computer programs that generate pixel patterns. Once begun, these are left to evolve for multiple generations and the resulting patterns are visualised on computer screens. This empirical approach contrasts with traditional methods that set out with a target behaviour in mind and devise equations to reproduce it. The implication is that by studying systems with the simplest possible structure, profound insights can be gained into much more complex ones.

The general intuition is that generating complexity is difficult, but Wolfram's experiments suggest that this is incorrect: programs based on the simplest rules can produce great complexity. In some instances, the behaviour is so complex as to be 'computationally irreducible', which means that the fastest possible way to determine what will happen is to sit back and watch the patterns unfold. It is Wolfram's contention that there is no mathematical algorithm that could short-cut this tedious running-out process, and that the behaviour of these programs defines a fundamental and insurmountable limit to what mathematics can, even in principle, achieve. Wolfram perceives this result as 'one of the more important single discoveries in the whole history of theoretical science': it suggests how 'nature seemingly so effortlessly produces so much that appears to us complex'. All natural systems, in other words, operate like simple computer programs.

Wolfram's 'cellular automata', as these programs are known, generate pixel squares on computer screens according to a simple rule that determines their sequence in a given line by assessing the colour – black or white – of neighbouring cells in the previous line. The pattern is left to evolve through iteration of the rule. An astronomically large set of such rules can be defined, and Wolfram charted the behaviour of as many of them as he could. Despite their extreme heterogeneity, the patterns may be divided, on visual criteria, into four principal categories. The first, 'uniform and simple repetitive' patterns, lack complexity and produce simple motifs. The more complex, fractal-like 'nested' patterns repeat themselves at larger and larger scales while maintaining an overall regularity. But the remarkable and unexpected result was the discovery of 'random' patterns that are both irregular and complex, and 'mixed' patterns that contain both order and randomness, in the form of elaborate localised structures that interact in complex ways. The only definitive way of classifying a program is to run it and observe its behaviour; this must be done for many generations, as its full complexity may not initially be apparent. In some cases the form of a rule suggests its potential behaviour, and programs clustered together in 'program space' often generate similar behaviour. In terms of our earlier evolutionary metaphor, we might say that Wolfram has systematically explored the landscape of the 'space of all possible simple programs'.

Wolfram's results are not specific to cellular automata: the contention that they are universal and can be reproduced in any system in which a simple program can be implemented is the cornerstone of his 'new kind of science'. In Wolfram's world, complexity emerges seamlessly and effortlessly.

In experiments with two and three-dimensional cellular automata, programs generate shapes reminiscent of complex biological structures, such as leaves and the hexagonal patterns of snowflakes, that classical mathematics struggles to describe. Wolfram believes this attests to a fundamental and deep correspondence between simple programs and natural systems: if the patterns in cellular automata are universal, it is irrelevant whether they are generated in a computer, in the detail of a zebra's stripes or in the spiral of a seashell. Systems that are computationally equivalent should behave in a similar manner.

Once this idea has been appreciated, the process of biological evolution can be reworked as a matter of programming, allowing evolution to be understood as a romp through the mathematical space of all possible simple programs. The evolution of an organism then corresponds with a process in which neighbouring regions of the program space are searched for better solutions. It is Wolfram's view that, for all but the simplest features, searches on these adaptive landscapes are heavily constrained. In order to generate a landscape favourable for the evolution of one feature, it may be necessary to destroy the landscape of another; or, when multiple features are optimised simultaneously, conflicting constraints may produce poor, compromise landscapes. This means that once a successful program has distinguished itself, it becomes difficult to find one that is better, and the power of natural selection is undermined. The result is that, in many cases, features conform to a local optimum that may not come close to the global optimum. Wolfram concludes that many of an organism's features arise because they are easy to produce, rather than because they confer a selective advantage. Furthermore, the problem presented by conflicting constraints on natural selection will tend to drive biological systems to avoid complexity rather than nurture it. Many of the complex features in biological organisms, therefore, 'arise not because of natural selection, but in spite of it'.

Thus Wolfram goes on, controversially, to assert that 'many of the most obvious examples of complexity in biological systems actually have very little to do with adaptation or natural selection.' Their existence is, in his view, an inevitable consequence of the fact that large numbers of simple rules generate complex behaviour – the landscape of simple program space is littered with such programs. If, on the other hand, complexity were generated only rarely, the random historical search for the necessary rules would have been more difficult, and complexity would have been the exception, rather than the norm. The imprint of a process of natural selection, within this frame of reference, is simplicity, not complexity.

Wolfram illustrates these ideas by noting the close correspondence between the extensive range of pigmentation patterns found on the shells of molluscs and the range of patterns produced by simple, randomly chosen programs. Asserting that these patterns are devoid of selective value – they are obscured by the molluscs' muddy environment – Wolfram explains their existence as a result of completely random changes in the underlying genetic programs. Indeed, he goes as far as to suggest that programs yielding mixtures of colours of the type found in molluscs might coincidentally produce highly complex patterns, so that the pattern itself is nothing more than the articulation of a constraint – one of Gould and Lewontin's 'evolutionary spandrels'.

The analogy between genomes and computer programs is widely accepted, and the computational framework has proved useful in interpreting a wide range of experimental data. Unlike the simple programs studied by Wolfram, the genomes of even the simplest creatures are very complex, but if that complexity is broken down into multiple, hierarchically nested sets of sub-programs, the analogy appears more reasonable. The actual similarities between genetic programs and Wolfram's cellular automata remain to be demonstrated. But simple, maximally compressed programs of the sort that Wolfram concerns himself with are likely to be brittle. They would be poor substrates for evolution, since the effects of their mutation would most likely be catastrophic. So, although simple programs are able to generate complex behaviour in Wolfram's mathematical universe, one might expect that in the real world of biology, more complex and more intrinsically evolvable programs would be favoured. The simple programs that Wolfram studies are in fact exactly of the type one would expect nature not to utilise. There is another, far more interesting family of possible programs located in what we might call 'moderately simple program space'. It is here that we are more likely to discover the programs actually deployed in biological systems: slightly more complex than Wolfram's simple programs, but more evolution-friendly and possessing the indispensable property of evolvability. It is reasonable to predict that natural selection has tuned the complexity of genetic programs so as to confer on them an optimal mix of stability and changeability, gently guiding them to the dynamic territory that Christopher Langton has called the 'edge of chaos'.

On a more positive note, Wolfram's hypothesis is falsifiable and provides considerable impetus for further experimental study, both in silica – on the computer screen – and in 'wet' biological systems. The visual representation of his results is novel, though others have experimented with more rudimentary genetic programs in the past. The methodical thoroughness of his exploration of simple program space is a considerable achievement and he has made a genuinely important contribution in demonstrating that even the simplest genetic programs can produce complexity. This, following our gradual recognition of the ease with which complexity can be generated in non-programmed self-organising systems, marks another significant step towards the demystification of the complexity of biological structures.

The problems remain, however. The book lacks any sense of historical or intellectual context: at no point does Wolfram so much as indicate the existence of other work in the field to which he is almost certainly indebted. There is no bibliography or sense of provenance, no indication whether an idea is new or old. Instead, there is the dogmatic, didactic drone of a single, unchecked and often pompous voice, repeatedly intoning phrases such as 'it is my strong suspicion that' or 'I have increasingly been led to believe'. As regards the scope of the book, Wolfram might have been expected to range from the macro-level of organs and appendages through the micro-level of cells to the nano-level of molecules and metabolisms. In fact, he focuses exclusively on the macroscopic level at the expense of the molecular and microscopic levels, where the most powerful evolutionary forces operate. And his trivialisation of the key processes of adaptation and natural selection is, in my view, fundamentally misguided.

One is left with the feeling that it is Wolfram, not Darwin, who has got it wrong. Yes, complexity can be generated from simplicity and the potential routes to complexity, both programmed and non-programmed, may be far more numerous than ever imagined. The generation of complexity in systems programmed by simple rules may be so routine that much of nature's complexity may in the first instance be accidental. But there are good reasons to uphold the explanatory power of selection in the face of these ideas. Programs exist in an unforgiving and capricious world: environments change continuously, and any system too brittle to respond to such change will snap; spandrels may well exist, but they are rarities, and even features that begin as spandrels will, in the end, become subject to a process of selection; and the Achilles heel of evolution by natural selection may well be the occasional degrading of adaptive landscapes through the demands of multiple competing constraints, but one would expect evolution to favour biological structures engineered in such a way as to minimise such situations. The principal role of selection is most likely to select from among self-organising systems, and the programs that enslave them, those elements that are good potential substrates for evolution. Evolutionary tuning will maximise their evolvability and strive for an optimal balance between order and complexity. The close embrace of non-programmed self-organisation, genetic programs and natural selection, buffeted all the while by historical contingency: this, at every level, is responsible for the great complexity of biological form.