In: de Bourcier, Lemmen & Thompson eds., 1994 *The Seventh White House Papers: Graduate Research in the Cognitive & Computing Sciences at Sussex.* University of Sussex, School of Cognitive & Computing Sciences, Brighton, UK. Research Paper CSRP 350.

Dynamic Fitness Landscapes

Seth Bullock
sethb@cogs.sussex.ac.uk

School of Cognitive & Computing Sciences
University of Sussex
Brighton
BN1 9QH

Abstract : Genetic Algorithms (GAs) are typically thought to work on static fitness landscapes. In contrast, natural evolution works on fitness landscapes that change over evolutionary time as a result of co-evolution. Sexual selection and predator-prey evolution are examined as clear examples of phenomena that transform fitness landscapes. The concept of co-evolution is subsequently defined, before attempts to utilise co-evolution in the use of GAs as design tools are reviewed and speculations concerning future applications of automatic co-evolutionary techniques for design are considered.

1 Genetic Algorithms

Genetic Algorithms (GAs) are a design/optimisation technique inspired by natural evolution (Goldberg, 1989; Holland, 1975). The bare essentials of evolutionary theory (selection, reproduction, variation, fitness, etc.) are extracted and applied to artificial genetic material in an attempt to evolve solutions to problems.

A genetic algorithm works on a group of potential solutions to a problem, termed a population. Each solution takes the form of a string (chromosome) of letters (genes) from an alphabet (typically consisting of the binary digits $\{0,1\}$). The algorithm first assesses these 'genotypes' allocating a fitness value to each, dependent upon the degree to which its associated 'phenotype' solves the problem (e.g. how well the parameters specified by a chromosome characterise a useful bridge, or pack a lorry, or optimize a function, etc.).

Once this has been carried out, a new population of genotypes can be created by 'breeding' the current population. 'Parent' chromosomes are chosen with some bias towards those that are most fit. Both parents contribute some 'genetic material' to their 'offspring'. This is typically achieved through some form of cross-over operator which takes two chromosomes and produces two more by swapping a randomly chosen portion of the genome from each parent. During reproduction there is a small chance that mutation may occur resulting in a less than perfect copy of genetic material. The role of mutation in the reproduction process is to introduce some random variation to the population. The cycle then repeats as the offspring generation is assessed, and bred, etc.

Over evolutionary time populations will (hopefully) 'converge' on solutions to problems, many of which may be counter-intuitive to the designer. Convergence is said to have occurred in a population if the genes at each locus reach some criterion of uniformity across the population.

Just how the discovery of solutions is achieved is controversial, but essentially the search space (defined by the manner in which the problem is parameterized) is tested in a parallel manner which takes advantage of beneficial strings of consecutive genes (schemata) by tending to group them together to form potential solutions.

2 Landscapes

Optimisation techniques are often thought of as traversing landscapes. A potential solution is represented as a point in such a landscape, the height of which corresponds to its ‘fitness’ – the extent to which it solves the problem. Fitness can either increase with height or increase with depth. For the remainder of this paper I will adopt the former convention.

If an optimisation technique tends to move from one potential solution to another which is slightly higher then it is a good technique and will tend to reach maxima in the fitness landscape. These solutions correspond to the tops of hills, points at which any small change to the solution (changing the length of one of the struts in a bridge, rotating a box slightly in a lorry, increasing one coefficient in an equation by a small amount, etc.) will result in a poorer solution (somewhere on the hillside). For any landscape there may be a number of local maxima and will be at least one global maximum. Local maxima are hill tops that although better than their surroundings are not the highest hill tops in the landscape. They are good solutions but not the best. Global maxima are the highest ground in the landscape. They are the best possible solution to the problem.

As was mentioned in the previous section, genetic algorithms typically work with a population of solutions scattered across the fitness landscape which gradually converge on one of the maxima. The fitness landscape is fixed by the designer of the GA when she decides how she will assess the potential solutions in the population. For example, if evolving a bridge design, the GA designer may specify fitness as being a function of some measure of safety, a measure of traffic capacity, and a measure of cost. As such a function does not change throughout the evolution process, the fitness landscape can be regarded as static, i.e. a potential solution with fitness x will always have fitness x independent of its peers or the passing of time. Indeed almost all optimisation techniques traverse static fitness landscapes.

In contrast, natural evolution works on a dynamic fitness landscape. Over evolutionary time the fitness of a phenotype (solution) may change radically. What was a winning strategy (e.g. eating flora of type A) becomes ‘out of date’ as conspecifics, predators, resources, etc. change through their own evolution. The resulting *co-evolution*, the evolution of systems in response to each other, can be thought of as ensuring that an organism’s evolutionary goal-posts (maxima in the fitness landscape) move. Rather than work towards the solution of some fixed problem, organisms are constantly adapting (over evolutionary time) to each other, their surroundings, etc. which are themselves adapting in response.

As an explicatory exercise, two examples of such co-evolution will be described before a more general characterisation of the phenomena at issue is attempted.

3 Sexual Selection and Predator-Prey Evolution

One of the oldest and most enduring problems in the study of animal behaviour is that of the genesis of the plethora of mating displays, colourful feathers, complex calls and songs, etc. that animals use to attract the opposite sex. With the advent of Darwin’s *On the Origin of Species* the beginnings of a solution were formulated. Darwin proposed that life-forms *evolve* through a process involving reproduction, inheritance, variation and selection. He considered the struggle for survival to be the primary selective pressure. Organisms slightly better equipped to deal with the trials of life will leave a greater number of offspring than their less well equipped peers. Over many generations populations will come to be composed of individual organisms that are well adapted to the dangers and resources of their respective niches.

However, this ‘natural’ selection cannot account for the myriad of impressive and complex sexual displays that are possessed throughout the animal kingdom. The peacock’s tail, the bower bird’s love-nest, the frog’s croak, and the stickleback’s dance do not contribute to their respective chances of survival. If anything, such behaviours and ornaments actively detract from an organism’s chances of survival through their costs in terms of energy and time (both of which could have been spent foraging or in some other useful activity), and in terms of becoming more likely to expire through predation (beautiful feathers and sonorous croaks not only attract mates but hungry predators too).

In order to explain the existence of such characteristics, Darwin proposed a separate selective pressure. Sexual selection can be considered as a corollary of the fact that survival is not the only requisite of an evolutionarily successful creature. In order to survive over evolutionary time one must reproduce. The evolution of traits which, although not beneficial in terms of natural selection (they do not increases the organism's chances of survival), improve an organism's reproductive chances can be accounted for by appealing to the notion that organisms slightly better equipped to deal with the trials of mate choice and reproduction will, again, leave a greater number of offspring than their less well equipped peers.

Darwin divided the mechanism of sexual selection into two categories, *male contests* and *female choice*. The former category contains behaviour and morphological traits that enable males to beat other males at the mating game. For example, the antlers of red deer, and the tusks of the bull walrus have evolved to serve the function of enabling a male deer or walrus to beat competitors in fights which determine a sexual 'pecking order'. Many of these traits (e.g. size, strength, fighting prowess, etc.) contribute to a general survival ability as well as specifically to reproductive success. However, male-male competition can develop in areas that do not aid an individual creature's survival.

Sperm competition is one such example. In species utilising this practice (e.g. dragonflies) females typically store the sperm of their mates for a time before their eggs are fertilised. Copulating males will attempt to expel the sperm of their mate's previous suitors through a variety of penis specialisations. To combat such expulsion techniques, methods of ensuring the retention of sperm, including the use of genital cement, co-evolve.

Female choice may, at first glance, seem less bizarre, relying as it does on the understandable notion that it pays to be picky. In most species, females put more effort into producing young than their mates[1]. They provide the vast majority of the time and energy necessary to produce a viable offspring. Given that this is the case, females that mate with poor stock are wasting valuable resources. Females with an ability to distinguish between poor and high quality mates will enjoy an advantage over their less perceptive conspecifics in terms of the quality and number of offspring produced.

Although this principal seems sound, how can it result in the preferences that we see in the animal kingdom? Why should pea-hens prefer peacocks with stunning plumage displays? Why should female frogs fall for males with loud, deep croaks? Two attempts to answer this style of question will be detailed below.

Fisher (1915) considered the case in which an infrequently occurring trait is favoured by natural selection. Females mating with males who display this trait produce fitter offspring. Thus the females with preferences for the male trait will proliferate, as will the males which possess the trait. However, offspring inheriting such a trait are not merely better equipped for survival, but are also better equipped for the task of securing a mate. Fisher showed that under certain conditions a "runaway process" could result in the latter benefit overpowering the former, resulting in males with traits that actively detract from their survival chances, and females with preferences for such traits.

In contrast, Zahavi (1975) proposed that rather than the reduction in an organism's chances of survival being an unfortunate result of runaway sexual selection, females might actually benefit from mating with "handicapped" males. He claimed, and was subsequently supported by mathematical proofs (Grafen, 1990) that showed, that handicaps (e.g. the peacock's tail) which increase an organism's chances of perishing due to predation or starvation are honest signals of its quality.

His logic revolves around the observation that costly displays can only be made by those able to afford them. In this context, displays that are costly in terms of fitness (they reduce ones reproductive chances) can only be made by those organisms with a high enough fitness to afford them. Sexual selection will thus result in loud, energetic, opulent, costly, extravagance.

Space does not permit a fuller coverage of the issues raised by these theories and the interested reader is referred to Gould and Gould (1989) and Krebs and Davies (1993) for more thorough treatments of the topic. What remains pertinent to this essay is that, as in the male contests described above, co-evolution

---

[1] Sea-horses are a notable exception. The female deposits her eggs in her mate's pouch, the eggs are fertilised, and gestate within this sac until the male 'gives birth' to them. This role-reversal also affects their mating ritual. Male sea-horses are coy and shy, whilst females must actively win their affection by displaying.

drives the female's mate preferences and the associated male traits, and in many cases drives them in such a way as to create organisms which have deviated from their strictly survival oriented ancestors in an attempt to satisfy the constantly changing demands of sexual selection.

Sexual selection is by no means the only example of such co-evolution. Predator-prey evolutionary dynamics also exhibit what behavioural ecologists have termed 'evolutionary arms races'. The development of higher acuity in a predator may be countered by the evolution of camouflage in a prey, teeth and claws provoke carapaces and scales, toxins demand antidotes, etc., etc.

Such arms races result in highly developed behavioural skills and complicated morphology. Such complexity is the result of the increasing demands placed on organisms by their environment (including their conspecifics, predators, prey, etc.). The hunting skill and speed of the peregrine falcon, for example, could not have evolved without the concurrent evolution of the perceptual capacity and escape capabilities of its prey.

Does any kind of co-adaptation qualify as co-evolution? Daniel Janzen (1980) distinguishes between true co-evolution and what he terms "diffuse" co-evolution. He defines the former as specific, reciprocal, evolutionary change, i.e. continued evolutionary change in the trait of one population in response to the continued, reciprocal, evolutionary change of a trait possessed by another population. In contrast, diffuse co-evolution is non-specific, reciprocal, evolutionary change, in which the trait of one population changes over evolutionary time in response to a *group* of traits possessed by another population (which may contain several species).

For example, the evolution of egg-mimicry and egg-discrimination in species of bird that respectively perpetrate and suffer the depositing of eggs in foreign nests is an example of true co-evolution (which is common in brood-parasitism and host-parasitism) in that the traits have evolved specifically for the purpose of brood-parasitism and defeating brood-parasitism respectively. Conversely, the hard shells of many crustaceans have evolved in response to a general threat from predators with a variety of body crushing/piercing techniques and are thus examples of diffuse co-evolution (See Krebs & Davies, 1991, ch.6 for further examples).

In fact, under Janzen's definitions, many instances of co-adapted predator-prey traits *cannot* be classed as the product of co-evolution. For example, the ultrasonic sound detectors in lacewing moths, which are a specific counter-adaptation to the sonar hunting technique developed by their bat predators, may be the product of a one-way adaptation on the part of the moths with no reciprocal evolutionary change in the bat predation mechanism. If this is the case then the co-adapted traits cannot be termed co-evolved.

Armed with such notions of co-evolution we can proceed to examine the prospects of attempting to apply co-evolutionary techniques to the use of genetic algorithms as design tools.

4 Co-evolutionary Design

A fundamental problem for the designer of genetic algorithms is specifying the problem that is to be solved in a manner that allows incremental steps towards a solution to be rewarded. If a problem is not specified in such a manner the genetic algorithm will have no feedback with which to drive its search and will essentially perform randomly until it finds a solution that can be rewarded. Co-evolution circumvents such problems by automatically moving the GA's evolutionary 'goal-posts', gradually changing the problem as the population moves over a dynamically changing fitness landscape.

What are the prospects for such an automatic co-evolutionary approach? Initial work in this area is thin on the ground. Simulations of undirected co-evolution have been undertaken (e.g. Werner & Dyer, 1991), but have little relevance here as they typically seek neither to explicate co-evolution nor utilise co-evolution in the solution of some design task.

The incremental approach of the Evolutionary Robotics Research Group at the University of Sussex can be seen as a first attempt to use co-evolution in the design of autonomous agents (Harvey, Husbands, & Cliff, 1994). The agents involved initially face a simple sensory-motor problem, which is incrementally made more difficult in an effort to coax complex behaviour from systems which could not be evolved from scratch. Such scaffolding techniques are reminiscent of the parent-child interactions which facilitate infant development (Rutkowska, 1994).

However, the hand-cranked nature of such scaffolding requires the presence of a human designer 'in the loop' and, potentially, the tasks of specifying the incremental goals that allow evolution to reach solutions to complex problems could itself become as problematic as designing the agents manually.

First attempts at utilising automatic co-evolutionary design includes work by David Hillis (1990) and Phil Robbins (1994), in which parasites are used to increase the performance of artificial agents, and Phil Husbands (1993) at the University of Sussex, in which the co-evolution of shop-floor schedules was explored. Such work, however, is in its infancy.

Before the full potential of co-evolutionary design techniques can be realised, the burgeoning body of work exploring artificial co-evolution must be consolidated. At Sussex, studies of predator-prey co-evolution (Miller & Cliff, 1994), sexual selection (Miller, 1994), and parental imprinting (Todd & Miller, 1993), have already been carried out and further research seems both worthwhile and inevitable. Open questions, such as the paucity of true co-evolution in natural predator-prey ecologies, in comparison to the relative abundance of such evolutionary dynamics in parasitic relationships, seem amenable to investigation through the artificial means employed within this style of research. The possibility of fruitful collaboration between the simulation of artificial co-evolution and the study of naturally occurring co-evolution seems to be a set of goal-posts worth shooting for, and one that will not be moving in the foreseeable future.

References

Fisher, R. A. (1915). The evolution of sexual preference. *Eugen. Rev.*, *7*(184 - 192).

Goldberg, D. E. (1989). *Genetic Algorithms in search, optimization and machine learning*. Addison-Wesley.

Gould, J. L., & Gould, C. G. (1989). *Sexual Selection*. Scientific American Library.

Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, *144*, 517 – 546.

Harvey, I., Husbands, P., & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*. A Bradford Book; MIT Press.

Hillis, D. (1990). Co-evolving parasites improve simulated evolution as an optimisation procedure. *Physica D*, *42*, 228 – 234.

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. MIT Press.

Husbands, P. (1993). An ecosystems model for integrated production planning. *International Journal of Computer Integrated Manufacturing*, *6*(1 & 2), 74 – 86.

Janzen, D. H. (1980). When is it coevolution?. *Evolution*, *34*(3), 611 – 612.

Krebs, J., & Davies (1993). *An Introduction to Behavioural Ecology* (3rd edition). Blackwell Scientific.

Krebs, J. R., & Davies, N. B. (Eds.). (1991). *Behavioural Ecology - An Evolutionary Apporach* (3rd edition). Blackwell Scientific.

Miller, G. F. (1994). Exploiting mate choice in evolutionary computation: Sexual selection as a process of search, optimization, and diversification. In Fogarty, T. C. (Ed.), *Evolutionary Computing: Proceedings of the 1994 Artificial Intelligence and Simulation of Behaviour (AISB) Society Workshop*, pp. 65 – 79. Springer-Verlag.

Miller, G. F., & Cliff, D. (1994). Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behaviour*, pp. 411 – 420. MIT Press.

Robbins, P. (1994). The effect of parasitism on the evolution of a communication protocol in an artificial life simulation. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*. A Bradford Book; MIT Press.

Rutkowska, J. (1994). Emergent functionality in human infants. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*. A Bradford Book; MIT Press.

Todd, P. M., & Miller, G. F. (1993). Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior*, *2*(1), 5 – 47.

Werner, G. M., & Dyer, M. G. (1991). Evolution of communication in artificial organisms. In Langotn, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II - SFI Studies in the Sciences of Complexity*, Vol. X, pp. 659 – 687 Redwood City, California. Addison-Wesley.

Zahavi, A. (1975). Mate selection – a selection for a handicap. *Journal of Theoretical Biology*, *53*, 205 – 214.