

Co-evolutionary Design: Implications for Evolutionary Robotics*

Seth G. Bullock

School of Cognitive and Computing Sciences

University of Sussex

Brighton

BN1 9QH

sethb@cogs.sussex.ac.uk

C.S.R.P No. 384

June 29, 1995

Abstract

Genetic Algorithms (GAs) typically work on static fitness landscapes. In contrast, natural evolution works on fitness landscapes that change over evolutionary time as a result of (amongst other things) co-evolution. The attractions of co-evolutionary design techniques are discussed, and attempts to utilise co-evolution in the use of GAs as design tools are reviewed, before the implications of natural predator-prey co-evolution are considered. Utilising strict definitions of *true* and *diffuse* co-evolution provided by Janzen (1980), a distinction is drawn between two styles of evolutionary niche, Predator and Parasite. The former niche is robust with respect to environmental change and features systems that have had to solve evolutionary problems in ways that reveal general purpose design principles, whilst the nature of the latter is such that, despite being fragile and unsatisfactory in these respects, it is nevertheless evolutionarily successful. It is contested that if co-evolutionary design is to provide systems that solve problems in ways that reveal general purpose design principles, i.e. to provide robust styles of solution, true co-evolution must be abandoned in favour of diffuse co-evolutionary design regimes.

Keywords

Co-evolution, Evolutionary Robotics, Predator-Prey Evolution, Parasitism, Niches

*Presented at the 3rd European Conference on Artificial Life, Granada, Spain, 4–6 June 1995. This research was supported by the Baldwin Bewdley Trust and a bursary from the School of Cognitive and Computing Sciences.

1 Introduction

The attraction of artificial evolutionary design techniques lies in the undisputed majesty of naturally evolved systems. That the organisms which populate the world – complex, economical, and robust solutions to evolutionary problems – were generated through a process of replication with variation is a supremely encouraging discovery. The development of man-made design processes fashioned in evolution's likeness offers us the chance to harness and direct some of the creative power of evolution.

However, the products of evolution are not necessarily the organisms that catch the evolutionary roboticist's eye. For example, evolution does not of necessity strive for complexity. This is attested to by the multitude of simple organisms prevalent on this planet. Similarly, although economy is the norm amongst evolved creatures, there exist selective pressures which are capable of producing energetic, opulent and costly extravagance (Zahavi, 1975). Further, although the vast majority of the solutions to life's trials arrived at via the evolutionary process *are* robust systems towards which, as a designer, one might direct ones admiration, there is evidence to suggest that evolution is equally capable of generating systems which should fail to satisfy the evolutionary roboticist's project.

Consider the problem encountered by Clever Hans, the counting horse. Hans was required to tap his foot a number of times to indicate the answers to simple sums. Hans arrived at a solution to this problem that I will term 'fragile'. Rather than actually performing the calculation, Hans merely tapped his foot until his trainer gave a slight, unintentional signal. As Hans' trainer always gave this signal when Hans had tapped the correct number of times, Hans always got the answer right. This solution is fragile because trivial changes in the problem which should not affect Hans' performance, such as removing the trainer from Hans' line of sight, will destroy Hans' 'mathematical' ability. Note that Hans' solution is still robust under a fairly wide range of circumstances including changes in the weather, the clothing of his trainer, etc. Fragility is thus subjective, and a matter of degree.

Evolution appears to have weeded out fragile solutions, having demanded of its creations that they satisfy reproductive demands in a wide variety of hostile, unpredictable situations, and indeed this is, on the whole, the case. However, such robustness is not a product of the *evolutionary process*, so much as the *selective pressure* under which most of natural evolution's subjects have toiled. In exceptional cases, evolution can produce systems as fragile as Hans' claims to numeracy.

If such systems exist in nature (and I will attempt to show that they do), how are we to avoid generating them in our attempts to evolve systems which are solutions to design problems? How can an evolutionary roboticist guarantee robust solutions? In answering these questions it will be useful to consider evolutionary biology's notion of co-evolution and how artificial evolution measures up to it.

2 Co-evolutionary Design

Genetic Algorithms (GAs) are a design/optimisation technique inspired by natural evolution (Goldberg, 1989; Holland, 1975). The bare essentials of evolutionary theory (selec-

tion, reproduction, variation, fitness, etc.) are idealised and applied to artificial genetic material in an attempt to evolve solutions to problems. Genetic algorithms, like other 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.

Genetic algorithms typically work with a population of solutions scattered across the fitness landscape which gradually converge on one of the fitness 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.

A fundamental problem for the designer of genetic algorithms is specifying the problem that is to be solved, in terms of a fitness function, in a manner that allows incremental steps towards a solution to be rewarded. If the fitness function, a metric which is used to assess the degree to which a solution solves the problem, is not so constructed, the genetic algorithm may spend periods of time on plateaux with no method of discriminating between competing solutions, or favour solutions which are evolutionary blind alleys leading to local maxima from which it is hard to escape, or suffer from excessively ‘noisy’ fitness landscapes which are prohibitive of incremental progress.

In contrast to these static landscapes, natural evolution works on a dynamic fitness landscape (Bullock, 1995). 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.

The possibility of artificial co-evolutionary design techniques offers the opportunity to circumvent the problems associated with static fitness landscapes 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, but can be divided into two camps.

Artificial, co-evolutionary paradigms concerned with modelling natural phenomena such as communication include seminal work by Werner and Dyer (1991), in which male and female animats co-evolved signal production and signal interpretation protocols in a simple animat-world, allowing immobile females to attract blind males. Although research paradigms involving such idealised worlds are fraught with difficulties (Cliff & Bullock,

1993), studies such as this laid the groundwork for subsequent, more theoretical, co-evolutionary studies of parental imprinting (Todd & Miller, 1993), aggressive signalling (de Bourcier & Wheeler, 1994), predator-prey co-evolution (Miller & Cliff, 1994), sexual selection (Miller, 1994), and co-evolutionary analysis tools (Cliff & Miller, 1995).

Additionally, my own research (e.g. Bullock (unpublished)), involving the use of co-evolutionary simulations as modelling tools for behavioural ecology, augments such efforts, currently exploring evolved communication, both natural and artificial. Although such work is theoretical in flavour, as the nature of artificial co-evolution is uncovered, and the intricacies of natural co-evolution are appreciated, implications of a more practical nature, particularly in the field of evolutionary robotics, seem inevitable. However, it is to the second group of research efforts that this paper is primarily addressed.

This second embryonic research programme is concerned with utilising artificial, co-evolutionary paradigms in the design of useful systems. In contrast to the theoretical drive of the co-evolutionary simulations mentioned above, these research projects are practical, pragmatic attempts to engineer solutions to real-world problems. Initial work in this area includes research by Hillis (1990) and Robbins (1994), in which parasites are used to increase the performance of artificial agents, and Husbands (1993), in which the co-evolution of shop-floor schedules was explored. Such work, however, is in its infancy.

The incremental approach of Harvey, Husbands, and Cliff (1994) can be seen as an attempt to use co-evolution in the design of autonomous agents. 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 that 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. However, a more general problem may be inherent in such co-evolutionary approaches. Before describing this problem in more detail, the natural co-evolution of predators and their prey must be considered.

3 Predator-Prey Co-evolution

Predator-prey evolutionary dynamics exhibit what behavioural ecologists have termed 'evolutionary arms races'. For example, 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. This 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 continued, *specific*, reciprocal, evolutionary change, i.e. continued evolutionary change in one specific trait of one population in response to the continued, reciprocal, evolutionary change of one specific trait possessed by another population. In contrast, diffuse co-evolution is continued, *non-specific*, reciprocal, evolutionary change, in which a 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 dumping of eggs in foreign nests is an example of true co-evolution in that the traits have evolved specifically for the purpose of brood-parasitism and defeating brood-parasitism respectively. Experimentation shows that species beyond the normal range of host-parasites show less strong egg-discrimination than those which routinely suffer egg-dumping, whilst species not subjected to such parasitism show little discriminatory ability at all (Davies, Bourke, & de L. Brooke, 1989). Further, groups within a species of parasite which differ genetically may evolve mimetic egg patterns specific to different hosts.

Conversely, although the hard, ridged shells of many crustaceans are the result of the continued, evolutionary escalation of shell breaching mechanisms employed by their predators, these impregnability traits have evolved in response to a group of traits in a number of predatory species (a variety of shell piercing/crushing techniques), rather than one specific trait (e.g. pincers) and are thus examples of diffuse co-evolution (see Endler, 1991, for further examples).

In fact, under Janzen's definitions, it is hard to find *any* instances of co-adapted predator-prey traits that can be classed as the product of *true* co-evolution. For example, the speed and strength of feline predators did not evolve solely in response to the escape tactics of their prey, but under a more general, "diffuse" selective pressure that favoured faster and stronger cats for a variety of reasons. Similarly, the discriminatory abilities of predators tend to be general ones rather than the egg-specialised systems of cuckoldry sufferers.

In contrast, there is evidence to suggest that much parasite-host co-evolution is *true* co-evolution. As well as the egg-dumping research mentioned above, ant colony parasitism by alien queens, and mammalian ectoparasites all feature co-evolved traits (Endler, 1991).

Endler explores a number of reasons to account for the differences between predator-prey and parasitic arms races. Firstly, whilst predators typically exploit many prey species, most parasites are host-specific. This ensures that specialisation by predators is difficult whilst parasites experience a "more consistent and simpler selection pressure" (pp.195), allowing the genesis of specific adaptations to their host species to occur more easily. Predators practicing apostatic selection, preferentially over-exploiting the more frequent prey types, and switching between preferred prey-types as prey densities change, will also experience less selective pressure for specialised traits. As parasites kill their prey more rarely, and thus spend more evolutionary time concentrating on one species, apostatic effects are a less significant pressure.

Predators encounter generalised defences more frequently than specialised ones as they typically occur earlier in the predation sequence – the series of stages that must be passed through before a predator consumes a prey item. The predator must pro-

ceed from Encounter, through Detection, Identification, Approach, and Subjugation to Consumption of the prey, and will be confronted with various defences along the way. For example, at the level of Encounter, hiding, alertness, and acuity are typical defences utilised by many prey types, but at the Subjugation/Consumption end of the sequence defences are typically specific to prey types (e.g. emetic unpalatability causing vomiting, spines, dispensible body parts, mobbing, etc.). This implies that adaptation to general defences may be more profitable for predators than evolving to defeat more specialised traits, but does not generalise to parasite-host scenarios.

The life/dinner principle (Dawkins & Krebs, 1979) suggests that the selective pressure on prey is much stronger than on predators. If a predator fails to secure a prey item it has lost its lunch, but if a prey item fails to avoid being secured by a predator, it has lost its life. As Dawkins and Krebs point out “A fox may reproduce after losing a race against a rabbit. No rabbit has ever reproduced after losing a race against a fox” (pp.59). Again this effect is weakened in the parasite-host situation as the host is killed less frequently.

Finally, prey species benefit in their struggle with predators from their superior numbers and, typically, their shorter breeding cycle. Both these factors allow faster evolution of prey than predators. As parasite numbers may equal and in many cases exceed those of their hosts, and their breeding cycles are not necessarily longer than those of their hosts, parasite evolution may proceed at rates equal to or greater than those of their hosts. This ensures that host adaptations may be matched blow for blow by parasites and in some cases comfortably ‘out-evolved’ by fast-breeding, numerous, parasitic organisms.

4 Implications

The differences between the evolutionary niches of the predators and parasites described above, and the resultant differences in both their behaviour and morphology hold a lesson for evolutionary robotics. Simplifying the findings described in the previous section, one can construct two idealised, stereotypical niche styles, Predator and Parasite. Whilst Predators cope with environmental change through recourse to flexible, general-purpose skills, Parasites have settled into niches that allow them the luxury of specialised, brittle life-styles. In the terminology of Section 1, Parasites are fragile solutions to the trials of life whilst Predatory solutions are robust¹.

What is wrong with the Parasitic lifestyle? In what sense is it brittle or fragile? It is evolutionarily successful, admirably economic, and at times breathtakingly complex. Granted, from the point of view of the behavioural ecologist the Parasite is no worse than

¹It must be reiterated here that the notions involved in this discussion are not dichotomies. They are gradated. No evolved system is truly general purpose. For example, although the vision system possessed by humans is a very flexible and general mechanism, it has limits, and may throw no light on the design principles of some extra-terrestrial vision system consisting of a tri-focal, x-ray-sensitive device. Similarly, no naturally evolved system is completely parasitic. Even Clever Hans' fragile foot-tapping displays a certain degree of flexibility. Further, judgements of fragility and robustness are subjective in that they are made with respect to some pre-given task or domain. If one is evolving discriminatory abilities as part of a limited-scope egg recognition system, one may be convinced of the robustness of a quite simple, egg-specific visual system. However, if a system is evolved as part of an endeavour to design a more general visual mechanism, more of it might be required before its robustness could be asserted.

any other style of niche. However, I would assert that evolutionary robotics is aiming for Predatory solutions to the problems it addresses rather than Parasitic ones. Parasites succeed in a manner that is inadmissible from the point of view of evolutionary robotics. As we saw in the previous section, they succeed through superior reproductive advantage and relatively simple selective pressure.

These two factors allow Parasites to eschew general-purpose solutions to the problems confronting them, in favour of specialised solutions that can be changed at a rate comparable to that of their hosts' counter-solutions.

Imagine a Parasitic bridge design, it has evolved specifically to exploit the river it spans, and the traffic that crosses it. If one challenges such a Parasite with environmental change within its lifetime it will most likely fail. Such a bridge will not adapt to span a different river. "So what?", one may retort, "It was evolved to span River X and carry Traffic Y". Further, if one examines such a Parasite in order to discover some general design principle, one will most likely be left dissatisfied. The characteristics of such a bridge which contribute to it being a good bridge will be indistinguishable from the characteristics which make it a good Bridge-Over-The-River-X-With-Traffic-Y.

Such solutions are analogous to that of Clever Hans. Whilst perhaps being adequate, they will not reveal general purpose principles because they do not succeed through the adoption of general-purpose principles.

If artificial co-evolutionary design is to avoid generating solutions analogous to the parasites discussed above, care must be taken to avoid scenarios in which artificial agents enjoy a numerical and reproductive advantage over the problems they are posed, and in which the selective pressure they experience is relatively simple. The slow-moving fitness landscapes that result from such scenarios pose no threat to the Parasitic lifestyle, being equivalent to the true co-evolutionary situations described in the previous section. The only fitness landscape which fosters Predatory solutions is one that is relatively fast moving and promotes diffuse co-evolution.

Interestingly one of the first studies of artificial co-evolution (Hillis, 1990) featured just such a fast-moving landscape. Hillis attempted to evolve list-sorting algorithms for fixed-length lists and assessed each solution with a sub-set of the total number of possible lists. The population of solutions was arranged on a toroidal grid and mating was limited to near neighbours. This resulted in fit algorithms spreading across the grid until either one algorithm filled the entire world or a number of algorithms maintained territories separated by a boundary of non-viable offspring.

He subsequently found that if the test-sets were placed under evolutionary control and distributed across a similar toroidal grid, with each test-set scored according to the the number of lists in the test-set that were incorrectly sorted by the algorithm in the corresponding grid location, the resultant co-evolution produced better sorting algorithms in less time. Grid areas populated by a certain algorithm would promote test-sets which defeated it, thus applying strong selective pressure to even the best algorithms. Similarly, easy test-sets would be driven to extinction by fit algorithms, thus guaranteeing the generation of the most difficult test-sets.

As the rates of evolution of both algorithm and test-set were comparable, the fitness landscape of each population was changing at a reasonably fast rate. Over evolutionary

time solutions to the sorting problem were not allowed to exploit the weaknesses of one test-set but were placed under diffuse selective pressure by a number of different test-sets.

5 Conclusion

In conclusion, two broad claims have been made within this paper. The first is that co-evolutionary design techniques are a useful way of dealing with the problems associated with static fitness landscapes. The second is that co-evolution will only produce fragile solutions unless it is of a diffuse nature. I intend neither of these points to be dogmatic as they are by no means applicable to every research programme. For example, if one can construct a good static fitness landscape, one that allows a GA to move towards, and eventually reach, an optimal solution, then co-evolutionary techniques may be superfluous (although even in this case they may accelerate the optimization process). Similarly, if, for some reason, fragile solutions are acceptable, the demand for diffuse selective pressure need not be met.

Indeed, in the case of the Parasitic Bridge described in the previous section, it may well be that a fragile solution is perfectly acceptable, as the general design principles underlying bridge building are well known. No extrapolation from the Parasitic Bridge to further bridge building efforts is necessary.

However, if ones evolutionary design is in any way exploratory, if one is working in a domain which is little understood, in a domain in which early successes must justify investment in the similar procurement of later successes, if one thus wishes to make claims of solutions arrived at through evolutionary design concerning the potential for such evolutionary design, to make general claims concerning classes of solution and the underlying principles that unify them, then one had better be sure that the solutions upon which one bases such claims are not Parasitic. Because Parasitic solutions admit of no such claims. Clever Hans was not the first step on some road to a theory of animal addition.

Acknowledgements

Many thanks to Inman Harvey, Phil Husbands and Richard Vaughan for discussion of the issues addressed herein, to Julia Hutchison for her numeracy, and Dave Cliff for comments on a previous draft.

References

Bullock, S. G. (unpublished) Sex, Lies and Ideolects: Simulating the evolution of deceptive courtship displays. In *The Simulation of Adaptive Behaviour as a Behavioural Ecology Modelling Tool*. Unpublished Research Proposal.

Bullock, S. G. (1995). Dynamic fitness landscapes. In de Bourcier, P., Lemmen, R., & Thompson, A. (Eds.), *The Seventh White House Papers: Graduate research in the Cognitive and Computing Sciences at Sussex*. University of Sussex.

Cliff, D., & Bullock, S. G. (1993). Adding “Foveal Vision” to Wilson’s Animat. *Adaptive Behaviour*, 2(1), 49 – 72.

Cliff, D., & Miller, G. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Morán, F., Moreno, A., Morelo, J. J., & Chacón, P. (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*. Springer.

Davies, N. B., Bourke, A. F. G., & de L. Brooke, M. (1989). Cuckoos and parasitic ants: Interspecific brood parasitism as an evolutionary arms race. *Trends in Ecology and Evolution*, 4, 274 – 278.

Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society Volume B*, 205, 489 – 511.

de Bourcier, P., & Wheeler, M. (1994). Signalling and territorial aggression: An investigation by means of Synthetic Behavioural Ecology. 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. 436 – 472. MIT Press/Bradford Books.

Endler, J. A. (1991). Interactions between predators and prey. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology: An Evolutionary Approach* (3rd edition). Blackwell Scientific Publications.

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

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*. MIT Press/Bradford Books.

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 co-evolution?. *Evolution*, 34(3), 611 – 612.

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/Bradford Books.

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*. MIT Press/Bradford Books;

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*. MIT Press/Bradford Books.

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.