

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. 84

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 contended 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

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

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 Husbards (1993), in which the co-evolution of shop-floor schedules was explored. Such work, however, is in its infancy.

The incremental approach of Harvey, Husbards, 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

ceed from Encounter,

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

time solutions

- 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.