Adaptation of Evolutionary Agents in Computational Ecologies

Paul Devine, Ray Paton and Martyn Amos

Department of Computer Science
University of Liverpool
PO Box 147
Liverpool L69 3BX, UK
email: paul@csc.liv.ac.uk
tel : 44 0151 794 3692

Abstract. We present a system developed for the application of evolutionary computation techniques to ecological simulation. A model ecosystem comprised of discrete adaptive agents and a biologically inspired environment is described and results presented. The agents are classifier based animats and the overall system has been derived from real ecologies and interactions.

1 Introduction

Ecologies are extremely complex biological systems in which adaptation is an essential characteristic. This operates over a number of physical and chronological scales by a variety of different processes and is central to the paradigm that we suggest here. Consequently, the modelling approach employed is driven by the need to implement adaptation and represent the mechanisms that contribute to it in the real world. The work presented here is highly ecological in motivation, environmental context is essential and the simulations capture salient features of real ecosystems, albeit in an abstract manner. The intention is to develop agent based modelling as an individual–based modelling technique [13] that allows adaptation at behavioural and evolutionary levels to be integrated into models of populations and communities.

The animat approach [20] has established itself as fertile ground for investigating adaptive behaviours in real, or more usually, simulated environments. Various mechanisms for adaptive functionality have been employed and their applications are well documented [14]. In common with much animat research, that described here is directed at the study of adaptation in a well defined context, namely a class of ecologically realistic herbivore environments. Other work devoted to similar systems has tended to be targeted at behavioural adaptation in either single or very few animats, often in environments rich in physical characteristics and possible interactions [18]. Our work differs in conception, it involves large numbers of coevolving individuals in an environment derived from real ecological systems. It is directed towards two levels of adaptation, one operating over somatic time (ie. during the life–cycle of an individual) and the other over evolutionary time. As evolutionary processes are essentially selectionist in nature we emulate them by maintaining a population of non-identical, reproducing animat–like agents in the simulated environment.

2 Ecological Aspects

Natural ecologies are often extremely difficult to investigate by the use of traditional formal models that rely heavily on the application of analytic mathematical techniques. This is a feature shared with many complex systems [4]. Our hypothesis is that the use of adaptive agents will allow both the construction of ecological models that may provide complementary perspectives to more traditional models and the investigation of evolutionary processes that are impossible to observe in the field due to the long timescales involved.

One of the most important features of natural evolution is the fact that populations are not usually homogeneous but rather are comprised of non-identical individuals. Even over non-evolutionary timescales individual variation may have considerable impact on the dynamics of a system, as has been observed in individual based field work [12].

It is this emphasis on ecological aspects that differentiates our system from other superficially similar systems. For example, some of the other work addressing the relationships between learning and evolution includes Latent Energy Environment based work [1], this is fundamentally different in its methodology to the work described here. The work of French and Messinger [9] is also relevant here although it is more highly abstracted and problem specific than our system.

The two extant systems which most closely resemble that described here are Swarm [10] and Echo [11], [8]. These are in many respects larger and more complex undertakings. Swarm is essentially a framework within which a large variety of complex adaptive systems may be modelled, ecologies being just a subset of these. Echo is probably bears the most similarity, but nevertheless we believe it to be quite distinct. Like our system, Echo is spatially explicit and agent based, however its treatment of agents, resources and interactions is significantly different. Echo agents interact by combat, trade and sexual reproduction. In our system interactions are limited to breeding, the implicit trophic interactions of an interferential grazing system [15] and the exchange of information between geographically connected agents.

We believe that the Echo approach allows a wider range of interactions to be modelled in an abstract manner, ours is more specific to a subset of ecological systems. Echo agents migrate by default, movement in the agents described here is 'intentional' as is direction. Resources in Echo are also more abstract than those in this system, the latter being designed to replicate many facets of a specific type of resource rather than a more general concept. We review this more fully in [6].

3 Resources, their nature and structure

Much conventional ecological theory concerns the competition for resources and their utilisation. The model proposed here shares this emphasis. Resources are extremely diverse and an appreciation of their nature is essential. We have elected to implement food as the primary resource with the option of introducing space too. Natural ecologies posses a richer and more diverse range of resources, but these are beyond the scope of a computational model and would merely add unnecessary levels of complexity to simulations. We follow Price's [16] broad temporal categorisation and selected steadily renewed resources as most appropriate for our simulations. The other categories are also catered for in our implementation.

In addition to this temporal aspect of resources their spatial properties are also very important. Many resources are not homogeneous in their distribution but are 'patchy'.

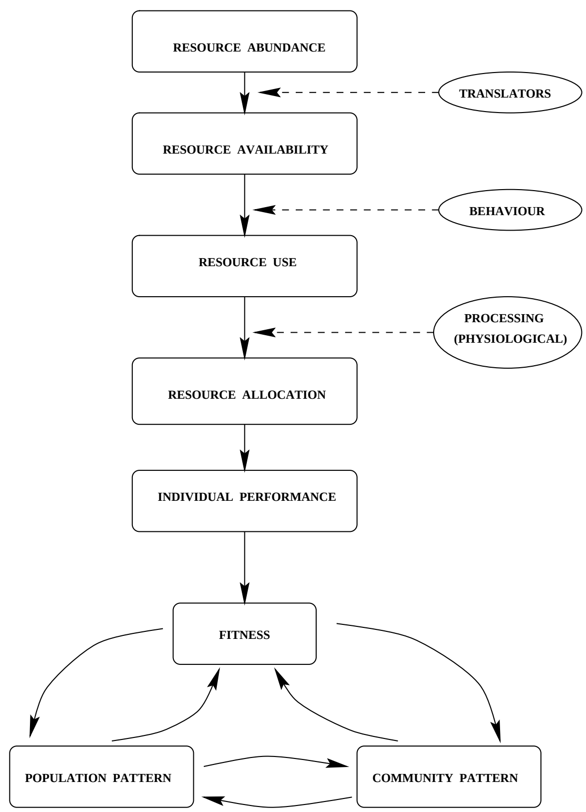

Fig. 1. Components of a Resource System. Adapted from Wiens

An ecosystem may be described as a resource system, Wiens [19] categorised the components of such a resource system, this has been adapted in Fig. 1 and forms the basic organisation of our artificial ecology. The fundamental interactions with the resource occur on an individual basis, with a consequential effect on the performance and fitness on the exploiting organism. These individual effects contribute to community and population structures, which in turn affect individual fitnesses. Clearly, over sufficiently long periods the action of selection will have evolutionary repercussions.

4 Simulation System Development

The long term goal is to produce a simulation system that incorporates many different facets of resource and consumer interactions. However, it was necessary to start with the development of a relatively simple version that nevertheless incorporated many aspects of a natural ecosystem. A purely Plant-Herbivore system was chosen, exploiting organisms and flora the major components. Superficially the absence of predation may appear to be an oversimplification, however, it is not unusual for such systems to be little impacted by predation or totally without predation. For the initial work a steadily renewed resource pattern was chosen in order to avoid seasonality and provide the agents with an environment that possessed at least some predictability in respect of resource properties.

In the context of the components of Fig. 2 the resource is flora, its abundance being dictated by various environmental factors. The translators affecting resource availability are the herbivorous agents themselves. The behaviour of these agents involves resource use and is linked both to translation and physiological processing in that the agent decides to allocate the resource to movement, breeding or its own somatic state. The effects on individual fitnesses are obvious and the consequent effects on community and population structure formed the basis of the initial investigation.

This herbivore system was developed loosely following the classification scheme of Coughley and Lawton [5] which derived its primary division from Monro's dichotomy [15]. Here 'herbivore' refers not only to cows, sheep etc. but to all creatures feeding on vegetable matter and 'grazing' encompasses all eating of plants by animals. An interferential grazing system with scramble competition was chosen for simulation.

5 Computational Aspects

The overall system is best described in terms of its two major components, the environment and the agent. These are described separately.

5.1 The environment

The simulation environment is relatively straightforward. It consists of a two dimensional rectangular lattice that is 'wrapped around' at its edges to form a torus. Each cell of the lattice contains one variant of the resource (ie. a species of flora) or none at all. Each variant has different properties in terms of its trophic utility and temporal type, any number of these may be implemented. This arrangement supports the construction of a large continuum of artificial environments of varying homogeneity, heterogeneity and predictability. The cells are geographically connected along their edges, not at their vertices, each cell having four nearest neighbours. Agent perception and movement in each step of the simulation is limited to its immediate cell and the nearest neighbours. Figure 2 illustrates an example heterogeneous environment.

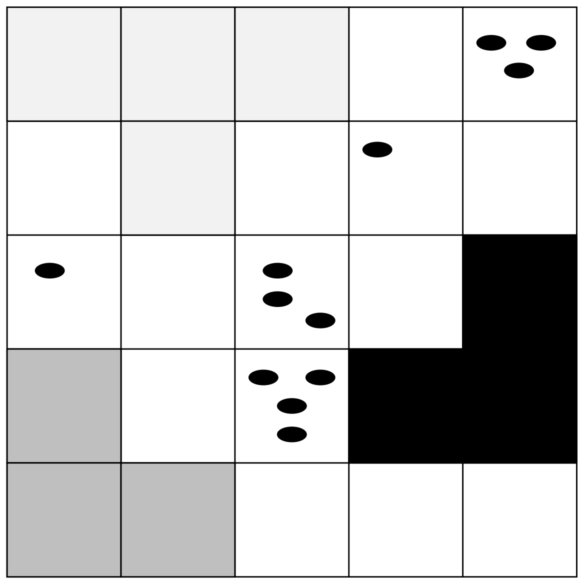

Fig. 2. An example heterogeneous environment, the black dots represent agents

5.2 Agents

The organisms that inhabit and interact with the environment are implemented as autonomous agents which adapt over both somatic and evolutionary time. Agents possess a number of simple attributes that govern their state and are capable of acting in a number of predefined, primitive ways. More complex behaviours emerge from the combination of these behavioural primitives.

The agent's state variables are derived from modals in theoretical ecology [2] but are redefined to suit the current application. These are:

Accumulated Somatic Investment (ASI). Here this is related purely to the amount that the agent eats and the energy that it expends.

$$ASI_a(t+1) = ASI_a(t) + f(d, N, e) - \gamma \tag{1}$$

Where a is an individual agent, γ the maintenance cost per round and $f(d, n, e)$, the trophic uptake includes the number of agents eating in the immediate area,N (as the system is interferential), the density of the flora, d, and its nutritional value, e. Should an agent's ASI fall to zero it is killed off. ASI is also reduced by reproducing. In some simulations it was found to be useful to place an upper bound on ASI as shall be discussed later.
Residual Reproductive Value (RRV) is a measure of the reproductive success of an agent, it has been implemented in two forms, the latter being the most important and closest to ecological theory.

$$RRV_a(t+1) = RRV_a(t) + \beta \tag{2}$$

β is a constant payoff. (2) only applies on rounds when an agent breeds. This has only been implemented in a version where reproduction is exclusively asexual. The second form has been implemented in systems where reproduction is both sexual and asexual.

$$RRV_a(t+1) = RRV_a(t) + \alpha.D \tag{3}$$

Here D is the number of living descendants an agent has and α is a fixed constant. Age is simply the number of rounds that the agent has lived through, a maximum age is usually set for a simulation, after this an agent is automatically killed off. Overall fitness is determined by the simple expedient of summing RRV and ASI. Consequently coefficients and functions are chosen to prevent one dominating the other.

Each agent possesses its own classifier system (CS) [3], as the mechanism of adaptation and the seat of individual variation. this has been adapted from the Michigan approach and is its organisation is shown in figure 3.

The organisation of the CS is shown schematically in figure 1, it is derived from the Michigan approach [11] and has a set of discrete rules, a message board and a credit assignment system (bucket brigade algorithm). The rule sets are at the core of the system. Each agent possesses a unique rule set. This rule set S is partitioned into two subsets, S' and S". This partition separates heritable and non-heritable rules. S' is the set of heritable rules that the agent inherits from its parent or parents and that it in turn may pass on to the next generation. S" is empty when an agent is born, over the course of an agent's life history this set stores rules produced by the rule discovery mechanism and those acquired from other agents in a cultural context. S" dies with the agent, rules outside S' may only outlive an agent if they have been learned by another agent. During asexual reproduction S' is subject to mutation and during sexual reproduction crossover too. Rule discovery during somatic time is by a substring matching process which operates across the entire rule set S.

6 Simulations and Behaviour

The initial population of agents is generated by specifying life history traits and state parameters, the initial rule set is randomly generated with only the overall

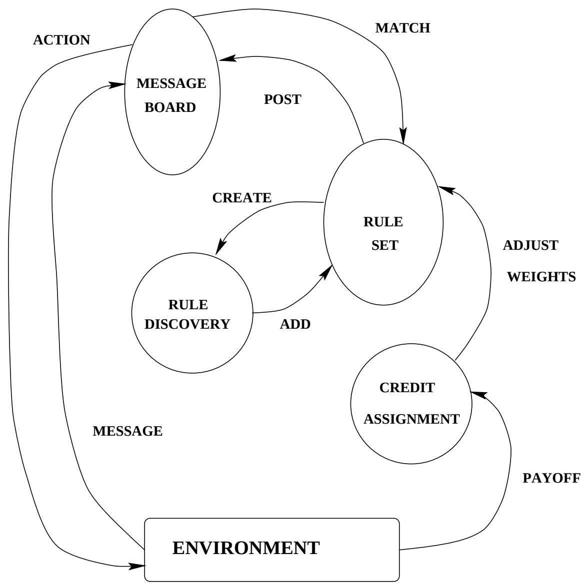

Fig. 3. Classifier System Organisation

specificity of the set being subject to control. All initial rule weights are set equal. It is also possible to 'seed' a simulation with agents evolved during earlier runs. It is important to note that the use of a random rule set effectively means that the initial population is not actually adapted to its environment, this a process which can take some time, hence the option of seeding with pre-adapted individuals.

Preliminary results relating to boom and bust population dynamics and Lotka-Volterra type cycling have been reported previously [7] and are not included here. Figures 4–9 show some typical dynamics exhibited by the agent populations. We have plotted the percentages of the populations performing specific actions each round. As a baseline for interpreting agent behaviour an agent identical in all respects to the adaptive agents other than its action selection was developed. Its action selection was made on the basis of a random number generator. The absence of any behavioural trends in the non-adaptive agent is shown in figure 4. In adaptive agents there are clear trends as the simulations proceed, figure 5 shows the incidence of trophic activity increasing. Figure's 6 and 7 show the relationship between moving and eating developing, it is interesting to note that some simulations show steady trends whilst others exhibit long periods of stability interspersed by quite rapid change. Eventually the type of behaviour shown in Figure 9 often emerges. Movement and eating appear closely linked and this behaviour occurs where Lotka–Volterra type population dynamics are most marked it coincides with the sort of coupling between movement and feeding shown in Fig. 13. As agents exhaust on area they must move to another, so a trend towards eating for a while and then moving is a good adaptive strategy.

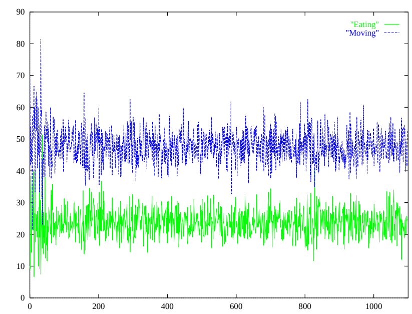

Fig. 4. Relative behaviour frequencies in random agent populations

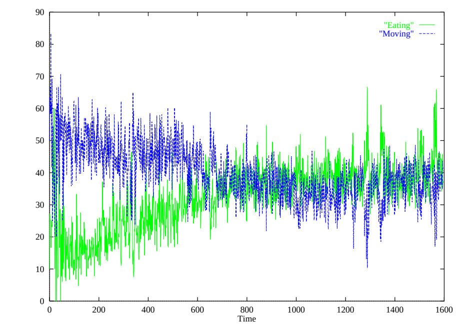

Fig. 7. Movement and Eating incidences changing gradually over time

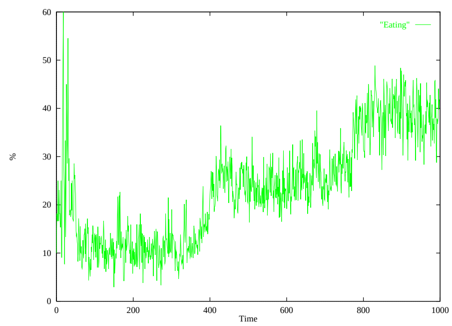

Fig. 5. Percentage of the population eating per round, non gradual change

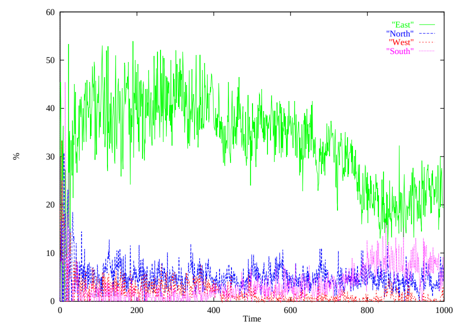

Fig. 8. Incidence of the three movement options per round

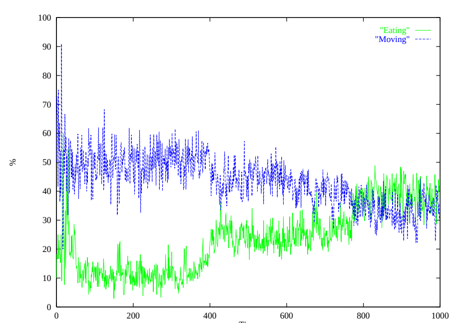

Fig. 6. Percentage eating and moving per round

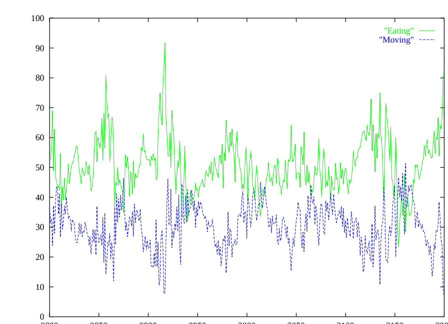

Fig. 9. Movement and eating options become synchronised

Whether this strategy is an Evolutionarily Stable Strategy (ESS) [17] is open to debate. Though frequently observed this trend appears and then disappears only to reappear again later, as it is invadable it cannot be considered evolutionarily stable.

When the components of movement are considered independently it appears that so long as movement is present the direction is unimportant, some directions dominating for a while and then being replaced by another. This may well be a consequence of the relatively small environment used for these experiments, only 25 areas. So long as agents move they will arrive at regenerated areas.

7 Interim and Prospectus

Preliminary work has established that the system does exhibit biologically and ecologically relevant behaviour. It is now to be used in the investigation of three ecological questions: sexual versus asexual reproduction; the impact of epigenetic information exchange between individuals in population and community structure and whether or not it affects the underlying rate of change in the rule sets; the development of different scales of migration and the performance of migrants. There is sufficient flexibility in both the agent and environment to permit a number of hypotheses to be investigated with the system in the future.

8 Acknowledgements

We would like to thank Prof. G. A Parker FRS for his support and interest. This work was funded by the Engineering and Physical Science Research Council.

References

1. D. Ackley and M. Littman. Interactions between learning and evolution. In Farmer, Langton, Rasmussen, and Taylor, editors, *Artificial Life II*, pages 488–509, Reading (Mass.), 1991. Addison-Wesley.
2. M. Begon. A general theory of life-history variation. In R.M. Silby and R.H. Smith, editors, *Behavioural Ecology - Ecological Consequences of Adaptive Behaviour*. Blackwell, 1985.
3. L. Booker, D.E. Goldberg, and J.H. Holland. Classifier systems and genetic algorithms. *Artificial Intelligence*, 1990.
4. R. Constanza, L. Wainger, C. Folke, and K. G. M. Ler. Modeling complex ecological economic systems: Toward an evolutionary, dynamic understanding of humans and nature. *BioScience*, 43:545–555, 1993.
5. G. Coughley and J.H. Laughton. Plant-herbivore systems. In R.M. May, editor, *Theoretical Ecology*, pages 132–166. Blackwell, Oxford, 2 edition, 1981.
6. P. Devine. Biologically inspired computational ecologies: A case study. Technical report, University of Liverpool, 1996.
7. P. Devine, G. Kendall, and R. Paton. When herby met elvis – experiments with genetics based learning systems. In V. J. Rayward-Smith, editor, *Modern Heuristic Search Methods*, pages 275–292. J. Wiley and Sons, 1996.

8. S. Forrest and T. Jones. Modeling complex adaptive systems with echo. In R. J. Stonier and X. H. Yu, editors, *Complex Systems: Mechanisms of Adaptation*, pages 3–21. IOS Press, 1994.
9. R. French and A. Messinger. Genes, phenes and the baldwin effect. In *Artificial Life IV*. MIT Press, 1994.
10. D. Heibeler. The swarm simulation system and individual-based modeling. In *Decision Support 2001: Advanced Technology for Natural Resource Management*, 1994.
11. J. H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, 1992.
12. M. Huston, D. DeAngelis, and W. Post. New computer models unify ecological theory. *Bioscience*, 38(10):682–691, 1984.
13. O. P. Judson. The rise of the individual-based model in ecology. *Trends in Ecology and Evolution*, 9:9–14, 1994.
14. J. A. Meyer and A. Guillot. From sab90 to sab94: Four years of animat research. In Cliff, Husbands, Meyer, and Wilson, editors, *From animals to animats 3. Proceedings of the third international conference on simulation of adaptive behaviour*. MIT Press, 1994.
15. J. Monro. The exploitation and conservation of resources by populations of insects. *Journal of Animal Ecology*, 36:531–47, 1967.
16. P. W. Price. Alternative paradigms in community ecology. In P. W. Price, C. N. Slobodchikoff, and W. S. Gaud, editors, *A New Ecology, Novel Approaches to Interactive Systems*, pages 353–383. J. Wiley and Sons, 1984.
17. J. Maynard Smith. Evolution and the theory of games. *American Scientist*, 64, 1976.
18. T. Tyrell and J. E. W. Mayhew. Computer simulation of an animal environt. In J. A. Meyer and S. W. Wilson, editors, *From Animals to Animats - Proc 1st Int Conf on Simulation of Adaptive Behaviour*. MIT, 1990.
19. J. A. Wiens. Resource systems, populations and communities. In P. W. Price, C. N. Slobodchikoff, and W. S. Gaud, editors, *A New Ecology, Novel Approaches to Interactive Systems*, pages 397–346. J. Wiley and Sons, 1984.
20. S. W. Wilson. The Animat path to AI. In *From Animals to Animats - Proc 1st Int Conf on Simulation of Adaptive Behaviour*, 1990.

This article was processed using the LaTeX macro package with LLNCS style