Computational Science Technical Note **CSTN-055**

Pack-Hunting Multi-Agent Animats

C. J. Scogings and K. A. Hawick

2009

Multi-agent systems sometimes yield surprising emergent behaviours that are unexpected given the individual or microscopic agent prescriptions. We have investigated the emergent spatial behaviour that results when predatory animats in a multi-agent system simulation are able to hunt in a pack. The consequences of this cooperative behaviour are significant for the macroscopic spatial patterns of both predator and prey agents. Our model is based on a very simple set of rules for individual animats, which have their own individual state, and exist cooperatively or in competition in a spatial world. We show quantitatively that cooperation allows the world area to support a higher number of agents.

Keywords: multi-agent; predator-prey; animat; emergence

BiBTeX reference:

```
@INPROCEEDINGS{CSTN-055,
        author = {C. J. Scogings and K. A. Hawick},
        title = {Pack-Hunting Multi-Agent Animats},
        booktitle = {Proc. IASTED Int. Symp. on Artificial Intelligence and Applications},
        year = {2009},
        pages = {220-224},
        address = {Innsbruck, Austria},
        month = {16-18 February},
        publisher = {IASTED},
        institution = {Massey University},
        keywords = {multi-agent; predator-prey; animat; emergence},
        owner = {kahawick},
        timestamp = {2008.11.19}
}
```

Technical Report CSTN-055

Pack-Hunting Multi-Agent Animats

Chris Scogings and Ken Hawick
Computer Science, Institute for Information and Mathematical Sciences,
Massey University, Albany, North Shore 102-904, Auckland, New Zealand
{c.scogings, k.a.hawick}@massey.ac.nz
Tel: +64 9 414 0800 Fax: +64 9 441 8181

June 14, 2008

Abstract

Multi-agent systems sometimes yield surprising emergent behaviours that are unexpected given the individual or microscopic agent prescriptions. We have investigated the emergent spatial behaviour that results when predatory animats in a multi-agent system simulation are able to hunt in a pack. The consequences of this cooperative behaviour are significant for the macroscopic spatial patterns of both predator and prey agents. Our model is based on a very simple set of rules for individual animats, which have their own individual state, and exist cooperatively or in competition in a spatial world. We show quantitatively that cooperation allows the world area to support a higher number of agents.

Keywords: multi-agent; predator-prey; animat; emergence.

1 Introduction

The defining features of a "multi-agent system" have been presented [1] as:

- Autonomy: the agents are at least partially autonomous
- Local views: no agent has a full global view of the system
- Decentralization: there is no one controlling agent

We have developed a sophisticated predator-prey simulation [2] to produce a unique world of "animats" [3] that eat, breed and interact to create fascinating emergent macro-behaviours. Two species of animat exist in an unbounded 2-dimensional plain – the predators that must eat prey to survive; and the prey that must eat "grass" to survive. The model incorporates behavioural rules similar to those of Boids [4] as well as the evolutionary processes found in many simulations such as Tierra [5], Avida [6] and Echo [7]. Over a period of several years, we have also developed and incorporated a number of general techniques for managing large-scale simulations and increasing efficiency in the model [8].

Every time step, each animat executes one of a small set of simple rules, e.g. "move away from adjacent predator". The rules are placed in a priority order and by changing the order of priority, different rule sets (genotypes) can be evolved [9]. The system can manage significant numbers of animats. A typical run contains tens of thousands and the current record is over a million animats. These high population figures enable emergent macro-behaviours such as the distinct spiral patterns discussed in [2], an emergent behaviour which can not be traced to any of the simple rules guiding the lives of individual animats. A number of formation patterns can be seen in Figure 1.

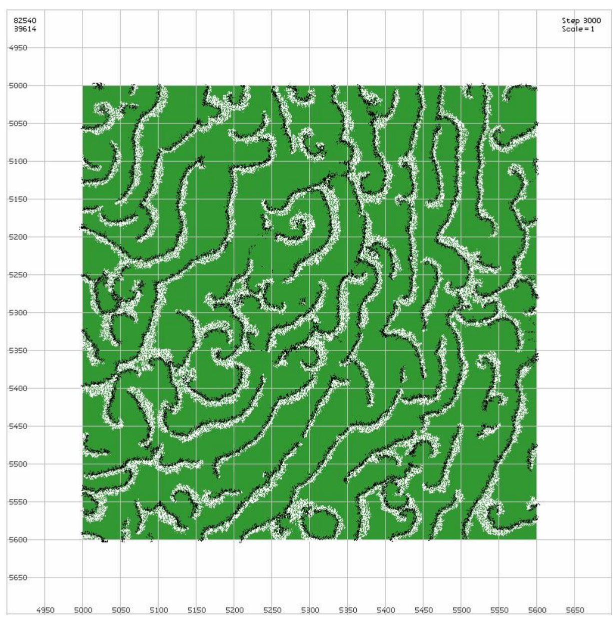

Figure 1: A typical situation at step 3000. Predators are black, prey is white and they inhabit a square "grassed area" with a grass value of 60. Various emergent formations have appeared including spirals.

One of the major problems facing multi-agent models is that of achieving local views for every agent because numerous global parameters are often imposed on all agents equally, irrespective of any local situation. For example, many genetic algorithms use globally imposed fitness functions to decide which organisms live or die. Such global effects may constrain emergent behaviours and we have refined our model to ensure that animat behaviour is governed by local criteria. Each animat eats, breeds or moves depending only on the proximity (and actions) of its neighbours. There is no global fitness function – there is only life or death.

Our model enables research into far more complex emergent behaviours than many previous systems because it simulates large numbers of higher-order animals. Further work is planned to study the interactions of different types of animats within the same community, e.g. "workers" and "soldiers". Our latest research project is to demonstrate that altruism can evolve naturally among animats and that it benefits both the group and the individual [10].

This paper consists of the following sections: A brief overview of our predator-prey simulation in section 2; an outline of an experiment involving cooperating predators in section 3; the results of the experiment are presented in section 4; and finally, a brief summary and conclusion in section 5.

2 The Model

Our model consists of two species of interacting animats – the predators and the prey. Animats have a very simple state: a health variable; an age; and a location. Prey consume "grass" which is assumed to be continually replenished and we can adjust the "nutritional value" and position of the grass. Predators consume only prey and other things being equal we can reproduce the well known boom-bust limit cycles predicted by predator-prey models such as the Lotka-Volterra coupled differential equations [11, 12].

Every animat carries a small set of rules that govern its behaviour and this rule set is passed on unchanged to any offspring. We have experimented with changing the order (priorities) of the rules and have investigated which rule sets generate the most successful animat groups [13]. The rule priorities used in this experiment are shown in Figures 2 and 3.

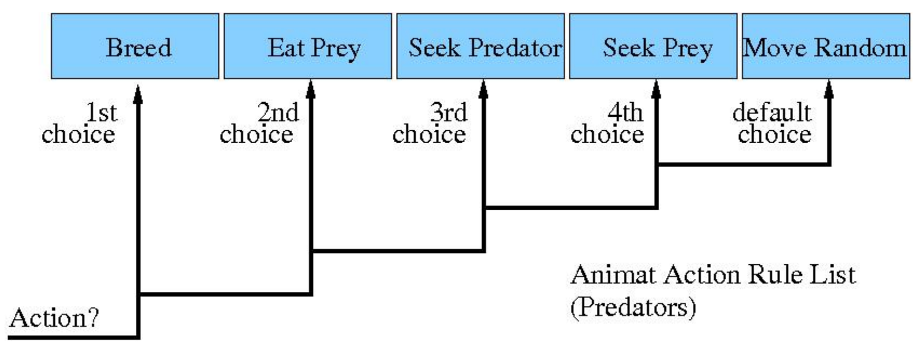

Figure 2: Predator rule set showing priority order. An animat executes one rule each timestep. If the conditions for a particular rule are not satisfied, the animat attempts to execute the next rule.

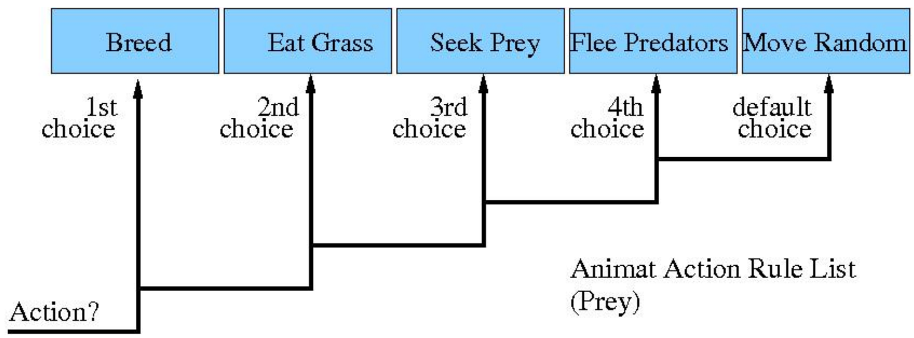

Figure 3: Prey rule set showing priority order.

Breeding only has a certain chance of success. This is a simple alternative to factoring in a host of complicated parameters including birth defects, nutrition, adequate shelter and so on. Note that the “Eat Prey” rule assumes that predators are always able to catch and consume adjacent prey. It is this rule that will be modified in the experiment described in section 3.

Rules are considered in a strict priority order. Each time-step, every animat attempts to execute the first rule in its rule set. However, most rules have conditions so can often not be executed. For example, predators can only eat prey if the prey is actually adjacent. If the conditions for the first rule can not be satisfied, the animat attempts to execute the next rule in the set and so on.

All animats in the model have a “current health” value. This value (in some ways analogous to “internal energy”) is reduced each time-step and if it reaches zero the animat “starves to death”. If an animat eats something (predators eat prey and prey eat “grass”) then the current health value will be increased by a certain amount, although it may never be increased past the maximum health value which is predetermined for each animat species. A “well-fed” animat has a current health value of two-thirds or more of the maximum. A “hungry” animat has a current health value of less than one-third of the maximum. Only well-fed animats can breed and only hungry animats can eat or seek prey. The concepts of health values and animats eating are discussed in [14].

3 The Experiment

Grass can be placed in specific locations on the map and each grassy area carries a specific “grass value”. When a prey animat eats the grass, its current health is increased by the grass value. This means that animats will do well on grass with a higher value and will struggle to survive on grass with a lower grass value. The snapshot provided in Figure 1 shows animats on a square grassy area with a grass value of 60. This is considered to be an above average value and the animat populations are thriving.

We then conducted an experiment to examine whether it was possible to lower the grass value but modify animat behaviour in some way in order to make best use of the reduced resources. One suggested solution was to get predators to cooperate and hunt in pairs. This was accomplished in the following two stages. The first stage was to modify the model so that predators only had a relatively small chance of actually catching prey. In addition predators were penalised a certain number of health points for engaging in a hunt (whether successful or not). If health points are equated with energy, this can be thought of as "expending energy" during the hunt. The grass value was lowered from 60 to 40 and the new penalties were introduced. This caused significant changes to the model as the predator population was greatly reduced and this, in turn, caused the prey population to increase. The complex features shown in Figure 1 vanish to be replaced by a uniform mass of mixed predators and prey – see Figure 4.

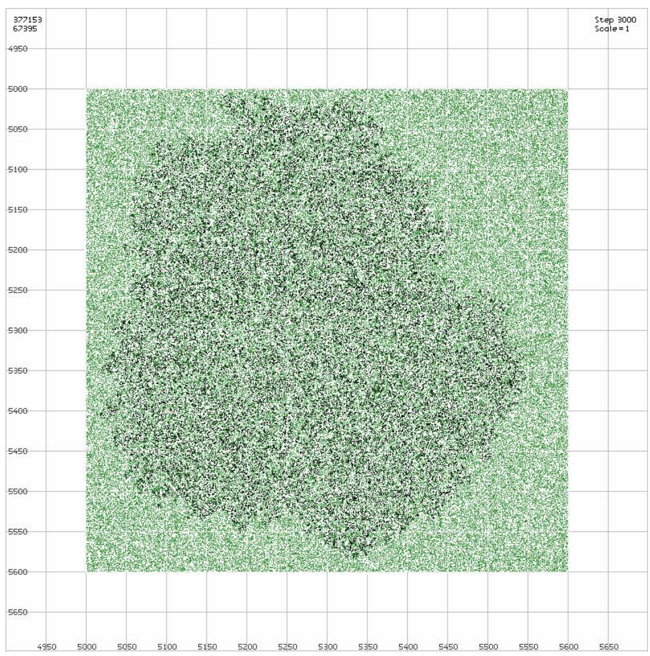

Figure 4: Step 3000 with a grass value of 40. Predators have only a 30% chance of catching prey and are also penalised for every hunt (whether it is successful or not). Predators are black and prey are white. The predator population has decreased allowing the prey population to increase. All formations and clusters have disappeared.

The second stage of the experiment was to introduce cooperative behaviour amongst the predators. The model was changed such that when a predator decides to execute the "Eat Prey" rule, it seeks support from its immediate neighbours. If there is another predator immediately adjacent it will assist in the hunt and thereby double the chances of catching prey. Note that predators can only be assisted by adjacent predators. There is no "long-distance communication" and if a predator becomes isolated then it is forced to hunt on its own. (At this stage predators only hunt in pairs but the potential exists for predators to hunt in larger packs in which it is quite possible that some sort of pack macro-behaviours may emerge.) Cooperative hunting led to a dramatic increase in the predator population with a corresponding decrease in prey. New and interesting formations emerged that appeared to be a mixture of the formations previously appearing in Figures 1 and 4. Figure 5 shows these new formations in which spirals are still emerging despite changes to the size and shape of clusters.

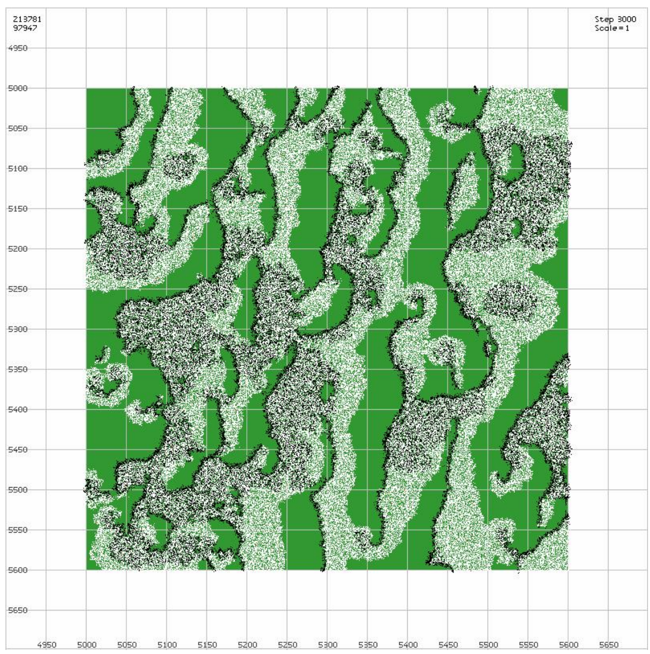

Figure 5: Step 3000 with a grass value of 40. Predators are penalised for an unsuccessful hunt but they are working as cooperating pairs and thus doubling their chances of success. Predators are black and prey are white. New formations (still including spirals) have emerged which appear to be a mixture of those shown in Figures 1 and 4.

4 Results

The results of the experiment are presented in Figure 6. The third and fourth lines (when looking at the right-hand side of the graph) show the predator

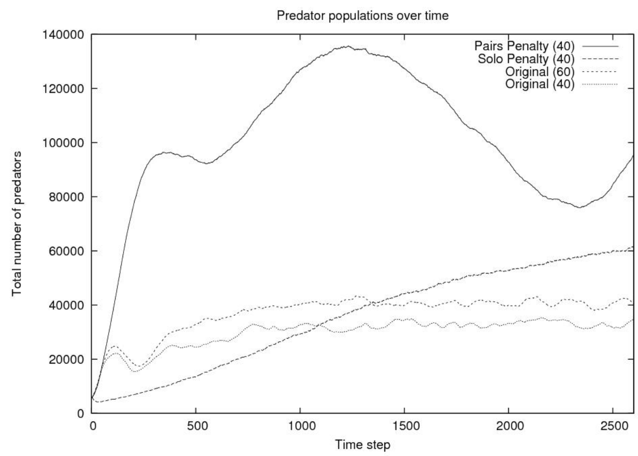

Figure 6: Results of the experiment. From top to bottom the lines represent the following populations of predators: Top line – predators when hunting as cooperating pairs and a grass value of 40; Second line (at the right-hand side) – predators not cooperating and penalised for hunting (also a grass value of 40). Third line – predators in the original model (no penalties and no cooperation) with a grass value of 60. Bottom line – predators in the original model with a grass value of 40. Fluctuations or experimental uncertainties are around ±1000 predators (between 1% and 3%).

populations in the original model. These are predators with a 100% chance of catching and consuming adjacent prey. These populations are both quite stable and show the drop in numbers when the grass value is reduced from 60 to 40.

The results become interesting when the predator "Eat Prey" rule is adjusted to (a) reduce the chance of successfully catching adjacent prey to 30% and (b) a penalty is introduced (health points are deducted) for hunting – irrespective of whether the hunt was successful or not. Predators are now less likely to catch prey and, initially, predator numbers are the lowest in the graph and increase only very slowly. However, because the prey population is not reduced at the normal rate, it too increases and thereby makes it more likely that an individual predator will have an abundance of adjacent prey. This means that (somewhere between steps 1000 and 1500) the predator population is able to overtake those of the original model. However we postulate that this increase is not sustainable and the predator population will reach a critical point and then decline. This boom-bust cycle has been well documented in previous discussions of the model [13]. The top line of the graph clearly shows the benefits of hunting in pairs as the predator population is significantly higher, even though the usual boom-bust cycle is present.

5 Summary

As we have shown, the model exhibits some quite dramatic changes when cooperation is introduced in the form of pack-hunting. The main outcome is that the animats are able to make much better use of the existing “land mass”. The interactions between groups of animats are driven by a mix of boundary effects and microscopic intermingling. The result is that there is still a high degree of spatial complexity in Figure 5 compared to that of Figure 4. On the other hand there are fewer regions of no animat activity – as there are in the non-cooperative model shown in Figure 1.

We are presently developing some other multi-agent analysis metrics such as animat autocorrelation functions. These are expected to emphasise quantitatively what seems clear enough qualitatively in the configuration snapshots. We are also investigating variable hunting-pack sizes.

We believe models such as ours, that posit a spatial collection of greater than 10^5 animat agents, are a good framework for investigating emergence and other cooperative effects in multi-agent systems.

References

[1] Wooldridge, M.: An Introduction to MultiAgent Systems. Wiley (2002)

[2] Hawick, K.A., Scogings, C.J., James, H.A.: Defensive spiral emergence in a predator-prey model. In Stonier, R., Han, Q., Li, W., eds.: Proc. Complexity 2004, Cairns, Australia (December 2004) 662–674

[3] Wilson, S.W.: The animat path to AI. In Meyer, J.A., Wilson, S.W., eds.: From Animals to Animats 1: Proceedings of The First International Conference on Simulation of Adaptive Behavior, Cambridge, MA: The MIT Press/Bradford Books (1991) 15–21

[4] Reynolds, C.: Flocks, herds and schools: A distributed behavioral model. In: SIGRAPH ’87: Proc. 14th Annual Conf. on Computer Graphics and Interactive Techniques, ACM (1987) 25–34

[5] Ray, T.: An approach to the synthesis of life. Artificial Life II, Santa Fe Institute Studies in the Sciences of Complexity **xi** (1991) 371–408

[6] Adami, C.: On modeling life. In Brooks, R., Maes, P., eds.: Proc. Artificial Life IV, MIT Press (1994) 269–274

[7] Holland, J.H.: Echoing emergence: Objectives, rough definitions, and speculations for echo-class models. In Cowan, G.A., Pines, D., Meltzer, D., eds.: Complexity: Metaphors, Models and Reality. Addison-Wesley, Reading, MA (1994) 309–342

[8] Scogings, C.J., Hawick, K.A., James, H.A.: Tools and techniques for optimisation of microscopic artificial life simulation models. In Nyongesa, H.,

ed.: Proc. IASTED International Conference on Modelling, Simulation and Optimization, Gabarone, Botswana (September 2006) 90–95

[9] Hawick, K.A., Scogings, C.J., James, H.A.: Spatial emergence of genotypical tribes in an animat simulation model. In Henderson, S.G., Biller, B., Hsieh, M.H., Tew, J.D., Barton, R.R., eds.: Proc. 2007 Winter Simulation Conference (WSC2007), Washington DC (9-12 December 2007) 1216–1222

[10] Scogings, C.J., Hawick, K.A.: Altruism amongst spatial predator-prey animats. Technical report, Information and Mathematical Sciences, Massey University (February 2008)

[11] Lotka, A.J.: Elements of Physical Biology. Williams & Williams, Baltimore (1925)

[12] Volterra, V.: Variazioni e fluttuazioni del numero d'individui in specie animali conviventi. Mem. R. Accad. Naz. dei Lincei, Ser VI **2** (1926)

[13] Hawick, K.A., James, H.A., Scogings, C.J.: Roles of rule-priority evolution in animat models. In: Proc. Australian Conference on Artificial Life (ACAL 2005), Sydney, Australia (December 2005) 99–116

[14] Scogings, C.J., Hawick, K.A., James, H.A.: Tuning growth stability in an animat agent model. In: Proc. IASTED International Conf. on Applied Simulation and Modelling (ASM 2007), Palma de Mallorca, Spain (29-31 August 2007) 312–317