Computational Science Technical Note **CSTN-067**

Agent-Based Evolutionary Search

K. A. Hawick and C. J. Scogings

2010

Spatial multi-agent systems provide a powerful model framework for investigating evolutionary behaviour amongst animat agents. We have developed a microscopic animat-based model in which autonomous agents are microscopically controlled by a rule-set that can be evolved using suitable operators. Our system can support over a million animats co-existing over many generations and has already been used to explore several collective phenomena including clustering; segregation; tribal warfare and battlefront formation. We incorporate simple microscopic behaviour rules based on local view information, that determine animat feeding; breeding; movement; seeking; and avoidance. We use a simple agent state model consisting of position, current health and age and a genetic code for each animat. We find a number of spatially-rich and complex, emergent patterns from the microscopic model and discuss how the model's convergence to stable macroscopic behaviour cycles is related to the localised rule parameters. We illustrate how an animat agent population of predators and prey can evolve more effective individuals by applying genetic algorithms to the species rule-sets and how our model framework and approach can be applied to sociological and predatory phenomena.

Keywords: evolution; agent architecture; convergence; complexity analysis

BiBTeX reference:

```
@INBOOK{CSTN-067,
        chapter = {Complex Emergent Behaviour from Evolutionary Spatial Animat Agents},
        pages = {139-160},
        title = {Agent-Based Evolutionary Search},
        publisher = {Springer},
        year = {2010},
        editor = {R. Barker and T. Ray},
        author = {K. A. Hawick and C. J. Scogings},
        number = {ISBN 978-3-642-13424-1},
        month = {January},
        note = {CSTN-067},
        doi = {10.1007/978-3-642-13425-8},
        institution = {Massey University},
        keywords = {evolution; agent architecture; convergence; complexity analysis},
        owner = {kahawick},
        timestamp = {2009.02.28}
}
```


Prof Ken Hawick, Computer Science, Massey University, Albany, North Shore 102-904, Auckland, New Zealand.

Technical Report CSTN-067

Complex Emergent Behaviour from Evolutionary Spatial Animat Agents

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

February 2009

Abstract

Spatial multi-agent systems provide a powerful model framework for investigating evolutionary behaviour amongst animat agents. We have developed a microscopic animat-based model in which autonomous agents are microscopically controlled by a rule-set that can be evolved using suitable operators. Our system can support over a million animats co-existing over many generations and has already been used to explore several collective phenomena including clustering; segregation; tribal warfare and battlefront formation. We incorporate simple microscopic behaviour rules based on local view information, that determine animat feeding; breeding; movement; seeking; and avoidance. We use a simple agent state model consisting of position, current health and age and a genetic code for each animat. We find a number of spatially-rich and complex, emergent patterns from the microscopic model and discuss how the model's convergence to stable macroscopic behaviour cycles is related to the localised rule parameters. We illustrate how an animat agent population of predators and prey can evolve more effective individuals by applying genetic algorithms to the species rule-sets and how our model framework and approach can be applied to sociological and predatory phenomena.

Keywords: novel frameworks; architecture; convergence analysis; complexity analysis.

1 Introduction

Understanding the emergent behaviour in highly non-linear complex systems in terms of some sort of descriptive model remains one of the ongoing goals of reductionist science. Many complex systems such as biological and social phenomena involving a number of individual components exhibit a richness of behaviour and a set of emergent properties that are not obvious from even a detailed knowledge of the workings of the individual component parts. Simulation models provide a means of exploring these emergent phenomena, whereby a system can be set up in a number of detailed starting conditions and its evolution carefully observed during a computational experiment.

One of the main problems in modelling and exploring evolutionary behaviour through simulations is the sheer size of the parameter space or fitness landscape that is usually encountered. It is extraordinary difficult to apply brute-force search methods to the phase spaces of many biologically-inspired computing models. The notion of genetic algorithms makes use of biologically-inspired mechanisms to combine and adapt existing "solutions" to find even better ones. This approach can still take very large computational resources to run simulated evolutionary models for long enough to see statistically significant changes or dramatically new solutions. It also presupposes that the starting conditions are sensible solutions and that in effect the experimenters are looking in roughly the right place - or at least a plausible position in solution space. While real biological systems proceed using thermodynamically scaled system sizes and very long time scales, even for simple computational models, it is not sufficient to start somewhere completely random and "hope" to evolve a solution somewhere in phase space. We discuss an approach and a model that we have developed, whereby we can employ evolutionary inspired methods in a very controlled manner, starting from well defined areas of model space and with a parameter space that is heavily constrained.

Multi-agent systems have proved a very interesting tool for exploring evolutionary behaviour since it is possible to encode relatively simple rules into individual agents and to perform computational experiments involving many interacting individual agents. However even this approach has its limitations, not least due to the need for large numbers of individual agents in the system to have any hope of duplicating the phenomena thought to occur in real biological or sociological systems.

Popular science accounts of genetics often fail to draw attention to the large number of components and time-scales involved in the evolution of real genetic systems[1]. In this article we describe how our software agent-based model can be used to explore complex emergent behaviours based on a relatively small and manageable space size of microscopic parameters that control individual spatial agents or animats.

Complex systems using spatial agents or animats have existed for some time – see for example [2, 3, 4] – and have yielded rich insights into emergent group behaviour in physical, biological and sociological simulation settings.

A useful starting point for us was to identify a fundamental behaviour such as predation and build up a microscopic model around it. We have refined our predator prey animat model over a number of years and it has been introduced and discussed in several previous publications including [5, 6, 7]. Unlike other models which focus on the evolution of animats and the emergence of new species, we concentrate on making explicit, well-defined changes to the microscopic control variables of the model and then analysing any new (emergent) animat collective behaviours.

In particular we have documented fascinating emergent macro-behaviours such as the defensive spirals and other features discussed in [5]. Figure 1 shows a typical example of the rich emergent spatial patterns formed in animat models. In this example the predators and prey have formed a spiral battlefront which will slowly rotate around as animats are born, eat, die and move.

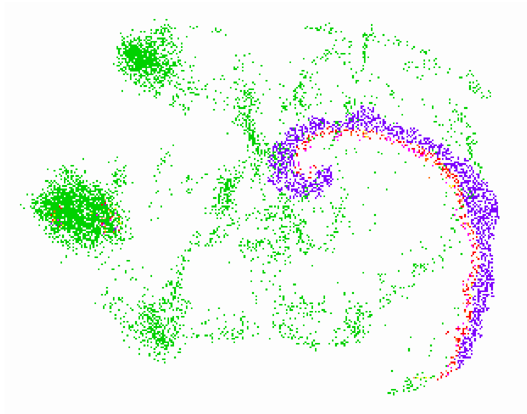

Figure 1: A complex emergent pattern showing a spiral battlefront of red predator "foxes" battling a group of blue prey "rabbits." The green residuals mark where an animat corpse remains after it has "died."

Our model exhibits some typical (and highly robust) wave front and proto-spiral pattern generation behaviours. We have been able to study these in a quantifiable manner by applying automatic feature detection techniques to the spatial animat patterns [5].

In this article we describe the detailed working of our animat model in section 2, including the features we incorporated to support evolutionary behaviour in section 2.2 and the parameter tuning that was necessary to attain stable systems in section section 2.3. We present a number of observations on the static and dynamic properties of the animat model in section 3. We review some typical properties and issues associated with genetic algorithms and other mechanisms for exploring fitness landscapes in section 4. We present some results of animat evolution in section 5, including control systems (section 5.1), genetic crossover results (section 5.2) and results when mutation was incorporated (section 5.3). We offer a discussion of ideas on the role that variable spatial conditions and environments cause in section 6 and some concluding remarks in section 7.

2 The Animat Model

Our model is essentially a predator-prey simulation that combines the replication rules of Conway's Game of Life with the behavioural rules of Boids [8] and the evolutionary processes of Tierra [9] and Avida [3] to produce a unique and exciting world of "animats" [10] 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.

2.1 Model basics

The state of each animat in the model is defined at all times by the following variables: species (predator or prey); gender (male or female); location (xy-coordinates); current age; current health; and a set of rules. The model also maintains the following global constants for each species: maximum age; maximum health; vision range (used to locate other animats); and a crowding number (if more than this number of animats are adjacent then they are crowded).

When an animat is created the current age is initialised to zero and is then incremented every time step. If the maximum age is reached the animat "dies of old age" and is removed. Similarly, an animat starts with its current health set to the average of its parents'. In early versions of the model, the current health was initialised to the maximum but it was found that this allowed populations on the brink of starvation to continue for prolonged periods. Each time step the current health is reduced and if it reaches zero the animat "starves to death". Current health can be increased by eating but can never be greater than the maximum health for that species. It is useful to define "hungry" to mean that current health is less than half the maximum and "well-fed" to mean that current health is equal to or greater than half the maximum.

At every time step, each animat executes one of a small set of simple rules such as "move away from predator". Rules are provided as a sequence of letter codes – which makes them easy to evolve and mutate when required. A new rule could be invented at any time and the rules currently in use are summarised in Table 1. Rules have conditions; for example a predator can only eat prey if (a) the predator is "hungry" and (b) the prey is adjacent. Breeding has several conditions including being well-fed and having an adjacent mate. In addition, there is a "birth rate" for each species which is the chance of breeding successfully. This provides a simple variable to control an otherwise complex mix of factors such as availability of food and shelter, possible birth complications, infant mortality rates, etc. Birth rates can be adjusted to change population levels in order to create a stable environment. Note that breeding can only be performed by females. Males can carry the breed rule in their rule set as a useless rule that will never be executed, but which can be passed on to female offspring who can execute it.

Table 1 lists the principle rules available to our animats. Each animat agent incorporates a list of rules which it tries to execute in order, trying each in turn until one succeeds. The last rule in the list may be the "do nothing" rule, which is independent of environment and therefore always succeeds. Adaptations of the rule list forms the basis of our evolutionary experiments, described below.

Normally the model is run with one set of rules for each species. This set of rules never varies and any new animats are exact clones of their parents. This is in line with our policy of carefully introducing small controlled changes to the model in order to more accurately measure the effects that such changes have on the system. Rules are placed in a priority order and each animat always executes the first rule in its list for which the conditions are satisfied. We have experimented with changing the order of priority of the rules [11] (but not the rules themselves) and thus produced different "tribes" of animats where each tribe has a different rule set (genotype) that consists of the same set of rules (for each species) but in a different priority order. The "move randomly" rule has an artificial condition in that it is only successful 50% of the time. If a rule has no conditions then any following rules in the rule set would never be executed.

2.2 Changes to the Model

In early versions of the model it was common for prey animats to breed rapidly and form dense clusters, particularly in areas where predators were temporarily absent. This often led to extremely

Code	Action	Predator	Prey	Conditions
A	move Away	No	Yes	adjacent prey (new rule)
B	Breed	Yes	Yes	well-fed; adjacent mate; birth conditions
E	Eat prey	Yes	No	hungry; adjacent prey
F	Flee predator	No	Yes	adjacent predator
G	eat Grass	No	Yes	hungry; not crowded; requires grass (new rule)
L	aLtruistic	Yes	No	adjacent needy predator (new rule)
M	seek Mate	Yes	Yes	well-fed; no adjacent mate
P	seek Prey	Yes	No	hungry; no adjacent prey
R	move Randomly	Yes	Yes	50% chance of success

Table 1: Rules currently used in the model. A rule is only executed if the conditions are met; otherwise that rule is ignored. The breeding conditions are discussed in the text.

large prey populations which prevented the formation of interesting emergent clusters and also caused a dramatic slowdown in the execution time for the model. (Note that this was prior to the introduction of "grass" as the first versions of the model assumed that prey could always find food and hence were implicitly feeding at all times.)

The concept of a crowd was introduced in order to solve this problem. If prey were too closely crowded together their chances of feeding and breeding were reduced. In addition a new rule "(A) move away from other prey animats" was introduced and this caused prey to spread out. This in turn reduced the very dense clusters that had been forming and also made it more likely that a formation of prey would be discovered (and destroyed) by predators. Crowding is still retained in the model but has less impact since the introduction of grass.

Inspired by Ronkko's model [12], grass was introduced into the system along with the new rule "(G) eat grass". Grass can be placed at specific locations on the map and can thus be used to restrict the animat populations. This is very useful as it ensures that populations do not endlessly expand and enables longer simulation runs that can then exhibit further emergent behaviour. Previous work [13] has shown that animat behaviour is not affected by the edges of the grassed area.

Grass is assigned a "nutritional" value and this is the amount that the current health of an animat will increase when it eats the grass. Thus grass with a value of 60 will increase current health by 60 although it may not exceed the maximum health which is usually 100 for prey. However, grass with a value of 30 will only increase current health by 30. This grass value is extremely useful as it allows the simulation of plentiful or scarce resources and any consequent changes in animat behaviour. Grass is never diminished and is assumed to retain the same nutritional value for the duration of the simulation. Future work may introduce grass that changes its value over time.

Figure 2 demonstrates the effect of different grass values. Animats have not survived on the inner grass values which are low (20 to 30) but are thriving on the outer sections (grass values 70 to 80) where the usual emergent spiral formations are clearly visible.

2.3 Fine Tuning

One of the significant problems of this type of model is execution speed. In particular, when every animat has to locate all its neighbours in every time-step, the increase in execution time becomes exponential. It is partly for this reason that many early models (for example Avida [3] and Tierra [9]) contain organisms that act in isolation with no knowledge of their neighbours.

We have spent several years fine tuning our model and optimising execution speed. Initially n animats were placed on an unbounded (effectively infinite) plain but the $O(n^2)$ search for neighbours soon rendered this impractical. Then the plain was divided into coarse-grained squares ("the grid") with several animats in each square. Animats could then search for neighbours only within the square they occupied – except for those that were positioned near the border of a square. Finally a fine-grained matrix was introduced with each animat occupying a cell and searching only for neigh-

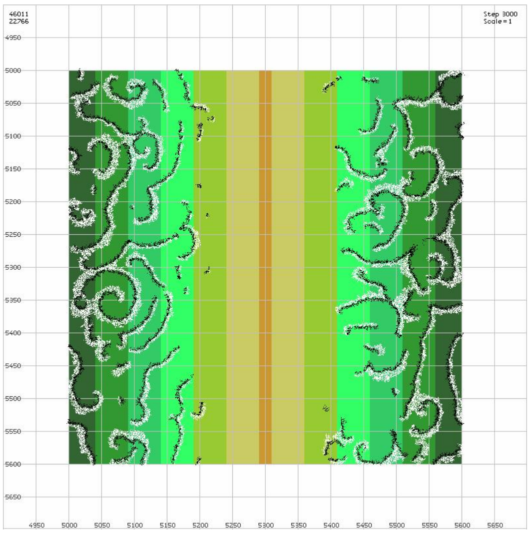

Figure 2: Clusters of animats at step 3000 on a background of varying grass values – 80 on the outer edges and 20 in the centre. Predators are black and prey are white. Animats can not exist on the lowest grass values in the centre but are thriving on the higher values on the outer edges. Note the emergent spiral clusters.

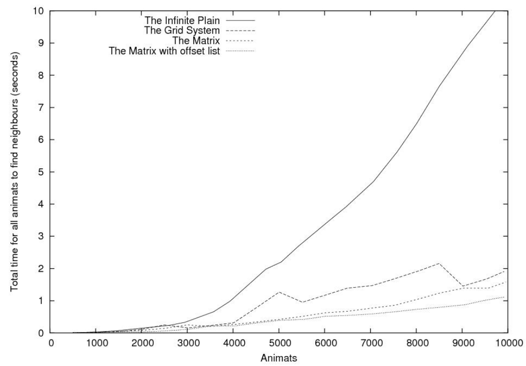

Figure 3: Run time of the model using the four different spatial approaches. The algorithms show continuous improvement although the "Grid System" shows execution time peaks at points where animats "spill over" into nearby grid squares. This figure first appeared in [6].

bours within the vision range. The speedup obtained by these measures is shown in Figure 3.

In order to reduce speed further the need for cumbersome distance calculations was removed by the introduction of a table of offsets. In order to find a neighbour the animat simply has to check all cells in the offset list. In the example in Figure 4 the vision range is assumed to be 3 (it is normally 20 to 60). In this case the offset list would contain (0,0)(-1,0)(0,-1)(0,1)(1,0)(-1,-1)(1,-1)(1,1)(-2,0)...(0,3)(3,0). Animat A would then only search through the shaded cells in order from nearest to furthest. Further detailed description of the fine tuning of the model can be found in [6].

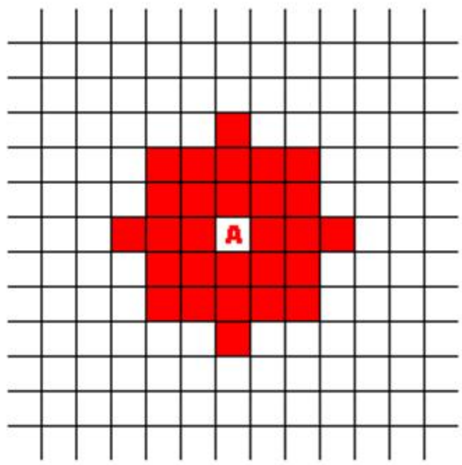

Figure 4: The radial lookup table prescribes those cells that might contain another animat. The cells are sorted by distance from the centre so that only the minimal list of possible locations necessary to find the nearest neighbours will be traversed.

3 Animat Model Observations

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 that would not be possible with lower numbers. The interaction of thousands of animats (of both species) as they follow their rules has produced fascinating emergent features in the form of macro-clusters. We have analysed and documented these emergent clusters in [5]. The most significant of these are distinct spiral patterns, an emergent behaviour which can not be traced to any of the simple rules guiding the lives of individual animats. Several spirals are visible in the figures in this chapter.

One of the biggest problems facing this type of simulation is what we call the "Hand of God" effect in which a model contains numerous global parameters that are imposed on all animats 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. We have refined our model to ensure that animat behaviour is governed almost entirely by local criteria. Each animat eats, breeds or moves depending only on the proximity (and actions) of its neighbours. In our model there is no global fitness function – there is only life or death.

This system enables research into far more complex emergent behaviours than many previous models of this type because it simulates higher-order animals. For instance we have demonstrated how altruism can benefit both groups and individuals when resources are scarce but that selfishness tends to dominate when resources are plentiful. In order to introduce altruism into the model we introduced a new rule "(L) Altruistic". When an animat executes this rule it finds an adjacent neighbour of its own species which has a lower current health level than its own. Both animats then receive the average of the two health levels. For example, an animat with a current health level of 60 has an adjacent neighbour with a current health level of 40. The first animat executes the altruistic rule and both animats receive a current health level of 50. It is interesting to note that this rule is very restricted; it is only executed if an immediate neighbour has a lower current health level and no attempt is made to seek out other "needy" animats that may not be adjacent. And yet the rule made a significant difference when resources (grass value) were low. A full description of the altruism experiments can be found in [14].

Another study of higher-order behaviours involved a study of the effects of predators working in teams (or "wolf packs") when hunting. The model was altered such that the "eat" rule was not always successful and there was a chance that predators might fail to catch and eat prey. Various techniques of cooperating hunters were then introduced and analysed. These experiments are discussed in [15]. This type of experiment demonstrates the versatility of our model which can easily be adapted to investigate a multitude of behaviours. Further work is planned to study the interactions of different types of animats within the same community, such as "workers" and "soldiers".

4 Evolution Algorithms

Genetic algorithms provide a mechanism for exploring the phase space of a model in search of optimal solutions [16]. There are two biologically inspired mechanisms that are commonly used: combination and mutation. Genetic crossover is a specific means of combining genetic information from the two parents to produce a new offspring that hopefully takes the best characteristics from its parents. Mutation is a simple random mechanism to introduce completely new information and avoid some of the problems of solutions being stuck in local minima of their fitness landscape.

The key problem in applying genetic algorithms is in setting up an appropriate genetic representation of the properties of individual life forms (solutions) so that a sensible interpretation can be made of "mutation" and "crossover." It is by no means trivial to follow the biological metaphor exactly and establish a genetic strand of pseudo DNA to which these notions can be directly applied.

Our model has a physical landscape which is essentially featureless. Individual animats do encounter a physical environment but its properties are solely a consequence of animats' neighbours. Predators need prey, and need other predators to breed with. Prey are assumed to eat "grass" which is replenished at a uniform rate in our model. This is analogous to a complex entropic system through which a source of energy (sunshine) continually flows. Energy is not conserved in our model system – it is not intended to be as it is not a closed system [17].

Through careful experimentation with our microscopic parameters we have arrived at a set of feasible prey and predator parameters. We can initialise our system with viable animats and not have to wait in hope for artificial life to emerge from the "primordial soup." In this sense we have vastly reduced the area of fitness landscape we need the evolutionary model to explore.

While we might have introduced a satisfactory predator type and a satisfactory prey type, by applying evolutionary operations we hope to find the "perfect predator" and the "perfect prey." However as both species are evolving together this means we have a situation of changing goals or in the language of genetic algorithms, the niches in our fitness landscape are varying too.

The spatial aspects of our landscape provide an interesting mechanism for mixing species but this also means that different niches might arise in different physical parts of the model system. We proceed on the assumptions that we can in fact model a big enough system of animats, and that these effects are not significant when averaged over many independent starting conditions and run times. We are able to model systems of around 10^6 animats over several thousand time steps on a typical desktop computer overnight. Parallel computing techniques make it relatively simple to average over around 100 separate independent starting configurations and runs.

5 Evolution Experiments

Our model has been designed to use genetic algorithms to evolve animat rules and rule sets. Each rule is supplied as a string of characters in order to maximise the scope for mutation and crossover. However, we have adopted a cautious approach, preferring to edit one small aspect of the model at a time and then measure the effect that each change has on the model.

As a first step we arranged the existence of several "tribes" or sub-groups of each species. All animats of the same species used the same rules but each tribe used the rules in a different priority order. Although no evolution was taking place, this "survival of the fittest" process ensured that successful tribes dominated and tribes with inefficient rule sets became extinct. This work [11] enabled us to identify efficient rule sets which led to robust and stable animat populations and these rule sets could then be used as the "standard" rule sets in future work.

The standard rule set for predators was defined to be:

1. Breed (B)
2. Eat prey (E)
3. Seek mate (M)
4. Seek prey (P)
5. Altruistic (L)
6. Move randomly (R)

Using the codes given in Table 1 this rule set is written as BEMPLR.

The standard rule set for prey was defined to be:

1. Breed (B)
2. Eat grass (G)
3. Move randomly (R)
4. Seek mate (M)
5. Move away from other prey (A)
6. Flee from predator (F)

Using the codes given in Table 1 this rule set is written as BGRMAF.

The position of the rule "Flee from predator" was a surprise. We had initially assumed that it would be an important rule for prey and had placed it high up in the priority list. However in the "tribal" experiments described above, all successful tribes placed the rule very far down the list. This illuminates one of the many small but interesting facets of working with emergence.

In this section we describe the second stage of the evolutionary process in which we allow rule sets to evolve and mutate over time. This is not yet what some would regard as a full genetic algorithm because we are only shifting existing rules around and not changing the rules themselves. For example we do not allow prey to evolve the rule "Eat prey". The system is capable of implementing this type of "full" evolution but we prefer to introduce one small change at a time and attempt to measure the effects of that change. We prefer also to analyse higher-order behaviours such as hunting, trading or altruism within a predator-prey context whereas a model in which prey could eat other prey would no longer be a predator-prey model. All the experiments described below were performed with a predator birth rate of 15% and a prey birth rate of 55%.

5.1 Control

The first experiment acted as the control. Animats used the standard rule sets and every animat of each species was therefore identical to all the others. The grass value was set to a standardised level of 60.

5.2 Crossover only

Animat rule sets were allowed to evolve by using a 2-point crossover [18] applied to the rule sets of the parents. For example if parent predators had rule sets BEMPLR and PLEMRB then their offspring would possess the rule set BEMMRB where the first three "chromosomes" are drawn from the rule set of the main parent (defined to be the parent with the same gender as the offspring) and the last three are drawn from the rule set of the other parent. Note that this means that some rules may appear several times in a rule set and others may not appear at all. The grass value was again set to 60.

By step 3000, the following rule sets were found to be dominating for prey:

Prey Rule Sets (Male)
BGRGBR (21%) Red
GAMBGR (20%) Blue
GRMBGR (18%) Brown
MGBFGR (14%) Pink

Prey Rule Sets (Female)
BGRBGR (32%) Yellow
BGRGBR (19%) Olive
BGRFGR (17%) Cyan

This shows clearly that the most important rules for prey are B(breed) and G(graze) and both of these are repeated and placed near the top of the list (i.e. high priority). Rule M(seek mate) also features prominently but other rules, including A(move away to relieve overcrowding) have vanished completely. Note that the B(breed) rule is only executed by females and thus is the leading rule in all successful female rule sets. The B(breed) rule is not important to males and thus appears lower down the list in successful male rule sets. The F(flee) rule has almost entirely disappeared and is clearly not important for successful rule sets.

By step 3000, the following rule sets were dominating for predators:

Predator Rule Sets (Male)
LPMBER (14%)
EPMBLR (9%)
LEPPRM (8%)
PMBBER (7%)

Predator Rule Sets (Female)
BPMBER (11%)
PEBBLR (7%)
EPBPRM (7%)
BEPRLM (6%)

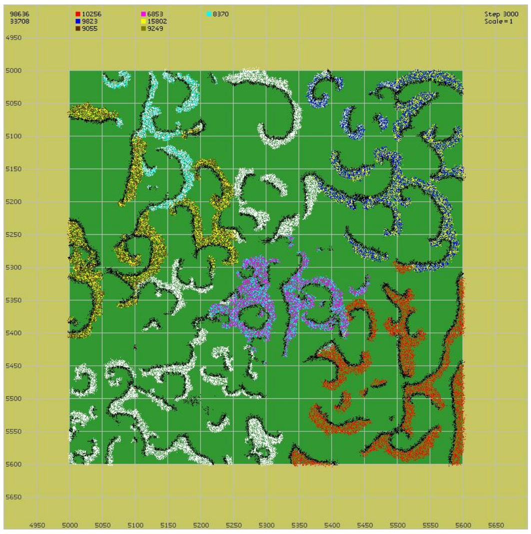

Figure 5: Clusters of animats at step 3000 of a run with crossover only. Predators are black and prey are white or coloured according to rule-set. Certain rule-sets have evolved more successfully than others (see section 5.2). Note the formation of spirals.

Figure 5 shows animats clustering when a crossover mechanism has been applied. The colours show the different subspecies or tribes amongst the prey. Spatial patterns such as spirals still form, as found in the original (non-evolutionary) model, but these are now superposed on the tribal segregatory patterns.

5.3 Crossover and Mutation

This experiment added mutation to crossover. When a new animat is created and a rule set is constructed for it (using crossover) there is a 5% chance that a mutation may take place. A mutation consists of a random rule being placed in a random position in the rule set.

By step 3000, the following rule sets were dominating for prey:

Prey Rule Sets (Male)
BMGAGRBGGR (897)
BMGAGRFGFR (658)
MGAMGRFFRR (629)

Prey Rule Sets (Female)
GBGBGRFGFR (1166)
BBGGRRBGGR (912)
GBABGRFFRR (676)

When comparing the results for crossover only (see section 5.2) and the results for crossover and mutation together (shown here) it is clear that the same rules are favoured, namely B(breed) and G(graze). However there are also significant differences as listed below:

- Mutation allows the list of rules to grow and/or shrink. In this case it is apparent that the replication of certain rules, such as G(graze), was advantageous and the rule sets have lengthened.
- The number of animats with the same rule set is lower. For example, a maximum of 15,802 female prey animats carried the same rule set with crossover only but only 1,106 female prey animats carried the same rule set after the introduction of mutation.
- There are more rule sets but some are very similar. For example the first two male prey rule sets are identical for the first six characters.

By step 3000, the following rule sets were dominating for prey:

Fox Rule Sets (Male)
MEBBPBBR (416)
EPMRBRB (350)
LPMEPRPR (292)
BMLLMERBP (267)
MPEPPRB (215)
next rule set (183)

Fox Rule Sets (Female)
BBEEPRPR (598)
MBEBPBBR (432)
BEEPMBRB (271)
next rule set (245)

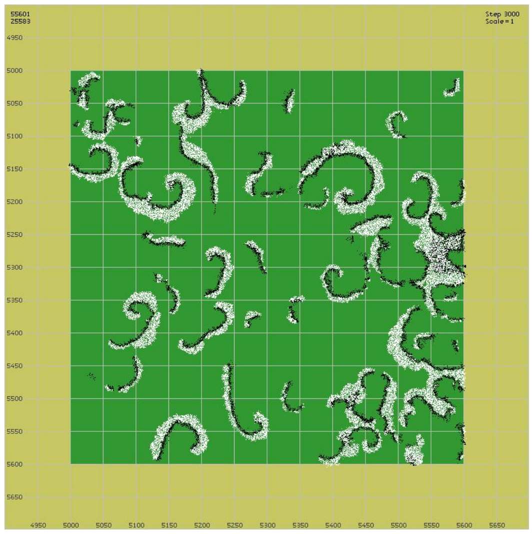

Figure 6: Clusters of animats at step 3000 of a run with crossover and mutation. Predators are black and prey are white. Individual rule sets are not shown as there are more of them and each set is much smaller. Spirals still form. Total population is found to be smaller than that with crossover alone.

Figure 6 shows a typical configuration after 3000 time-steps when both genetic crossover and mutation take place. In this system spirals again persist but the segregatory effect is much less pronounced. Many more different tribes survive but are smaller in number. As a result the overall population is less successful and there are fewer animats.

6 Evolution - varying conditions

Once the initial process of evolution had been established, a series of experiments were performed in order to analyse how well evolution assisted animats to adapt to a harsher environment. The nutritional value of the underlying grass area was modified in order to make it harder for prey (and hence predators) to survive.

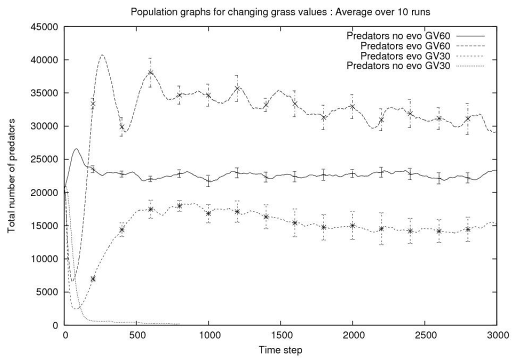

Figure 7: Total predator population as it varies s with time for the different algorithms - with and without evolution and with different grass richness values.

Figure 7 shows the typical boom-bust cycles in the population (of predators) superposed on the time varying envelop functions that differentiate typical population trends with and without evolutionary effects being incorporated. The most successful configuration - in terms of total population is when evolutionary effects are incorporated but the resources (grass) are throttled back – to encourage competition. Having too much grass available mean there are too many prey and the evolutionary effects amongst predators are not advantageous enough to show an effect.

7 Discussion and Conclusions

We have described our spatial animat model based on a predator-prey intermix of animat agents that co-exist on a simulated flat world. We have related how this model has lead to study and worthwhile interpretation of a number of interesting emergent phenomena, including: spiral battlefront patterns; tribal segregation; resource scarcity effects; pack formation and a limited degree of altruism.

These effects have been found in a number of real systems such as flocking and herding of birds or animals [19, 20] and ant colonies [21].

In this article we focused on the effects of introducing a limited form of evolutionary adaptation. Individual animats can give birth to new offspring that either combine features from their parents or can have small random mutations added.

In our system we found that a form of crossover had the most beneficial effect with some subspecies tribes forming and being generally more successful than the random mutated species.

We observe the interesting effect that while evolution can help the model population to greater success in terms of animats supported per unit area of space, the conditions need to be right to do this. In particular it required a competitive environment - with less resource available to prey and hence few prey available on average to predators, before this evolution of predators was seen to be quantitatively effective.

We believe models such as ours provide a valuable platform to experiment with evolutionary agent algorithms and effects, and that as well as giving general insights into emergent properties of a complex statistical system, there are also parallels with ecological and sociological systems in the real world.

References

[1] Turney, J.: Frankenstein's Footsteps - Science, Genetics and Popular Culture. Yale University Press (1998) ISBN 0-300-07417-4.

[2] Tyrrell, T., Mayhew, J.E.W.: Computer simulation of an animal environment. In Meyer, J.A., Wilson, S.W., eds.: From Animals to Animats, Proceedings of the First International Conference on Simulation of Adaptive Behavior. (1991) 263–272

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

[4] 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

[5] Hawick, K.A., Scogings, C.J., James, H.A.: Defensive spiral emergence in a predator-prey model. Complexity International (2008) 1–10

[6] Scogings, C.J., Hawick, K.A., James, H.A.: Tools and techniques for optimisation of microscopic artificial life simulation models. In Nyongesa, H., ed.: Proceedings of the Sixth IASTED International Conference on Modelling, Simulation, and Optimization, Gabarone, Botswana (2006) 90–95

[7] Hawick, K.A., Scogings, C.J., James, H.A.: Stochastic animats in complex spatial environments. Technical Report CSTN-040, Computer Science, Massey University (2007)

[8] 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 ISBN 0-89791-227-6.

[9] 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

[10] Wilson, S.W.: The animat path to AI. In Meyer, J.A., Wilson, S., 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

[11] 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 2007 (WSC2007), Washington DC, USA (2007) 1216–1222

[12] Ronkko, M.: An artificial ecosystem: Emergent dynamics and lifelike properties. J. ALife **13** (2007) 159–187

[13] Scogings, C.J., Hawick, K.A.: Global constraints and diffusion in a localised animat agent model. In: Proc. IASTED Int. Conf. on Applied Simulation and Modelling, Corfu, Greece (2008) CSTN-047.

[14] Scogings, C., Hawick, K.: Altruism amongst spatial predator-prey animats. In Bullock, S., Noble, J., Watson, R., Bedau, M., eds.: Proc. 11th Int. Conf. on the Simulation and Synthesis of Living Systems (ALife XI), Winchester, UK, MIT Press (2008) 537–544

[15] Scogings, C., Hawick, K.: Pack-hunting multi-agent animats. Technical Report CSTN-055, Massey University (2008)

[16] Holland, J.H.: Adaptation in natural and artificial systems. 2nd edn. MIT Press, Cambridge. MA (1992)

[17] Adami, C.: Introduction to Artificial Life. Springer-Verlag (1998) ISBN 0-387-94646-2.

[18] Lenski, R.E., Ofira, C., Collier, T.C., Adami, C.: Genome complexity, robustness, and genetic interactions in digital organisms. Nature (1999) 661–664

[19] Romanczuk, P., Couzin, I.D., Schimansky-Geier, L.: Collective motion due to individual escape and pursuit response. Phys. Rev. Lett. **102** (2009) 010602–1–4

[20] Hinchey, M.G., Sterritt, R., Rouff, C.: Swarms and swarm intelligence. Computer (2007) 111–113

[21] Bonabeau, E., Theraulaz, G., Fourcassie, V., Deneubourg, J.L.: Phase-ordering kinetics of cemetery organization in ants. Phys. Rev. E **57** (1998) 4568–4571