Computational Science Technical Note **CSTN-041**

Boundary Conditions and Locality in an Agent-Based Predator-Prey Model

C. J. Scogings and K. A. Hawick and H. A. James

2008

Agent-based simulation modeling provides important insights into emergent behaviours in complex systems. We report on the effects of introducing boundary conditions into a sophisticated predator-prey model in which animal agents or "animats" usually exist in an unbounded plain. We find that, due to the localisation of the animats, the introduction of the boundary conditions has no noticeable effect on emergent behaviour. We provide a summary of the factors that have created this highly localised simulation behaviour.

Keywords: agent model; animat simulation; boundary conditions; spatial emergence

BiBTeX reference:

```
@INPROCEEDINGS{CSTN-041,
        author = {C. J. Scogings and K. A. Hawick and H. A. James},
        title = {Boundary Conditions and Locality in an Agent-Based Predator-Prey
                Model},
        booktitle = {Proc. 17th IASTED Int. Conf on Applied Simulation and Modelling},
        year = {2008},
        number = {609-016},
        pages = {1-6},
        address = {Corfu, Greece},
        month = {23-25 June},
        publisher = {IASTED},
        institution = {Massey University},
        keywords = {agent model; animat simulation; boundary conditions; spatial emergence},
        owner = {kahawick},
        timestamp = {2008.01.14}
}
```

Boundary Conditions and Locality in an Agent-Based Predator-Prey Model

C.J. Scogings, K.A. Hawick
Computer Science, Institute for Information and Mathematical Sciences, Massey University,
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
and
H.A. James
The South Australian Partnership for Advanced Computing, Adelaide, SA 2005, Australia
heath.james@sapac.edu.au

Abstract

Agent-based simulation modeling provides important insights into emergent behaviours in complex systems. We report on the effects of introducing boundary conditions into a sophisticated predator-prey model in which animal agents or "animats" usually exist in an unbounded plain. We find that, due to the localisation of the animats, the introduction of the boundary conditions has no noticeable effect on emergent behaviour. We provide a summary of the factors that have created this highly localised simulation behaviour.

Keywords: agent model; animat simulation; boundary conditions; spatial emergence.

1 Introduction

Modeling emergence in complex systems using spatial agents or "animats" [1] continues to yield new insights into collective phenomena in physical, biological and sociological simulation settings [2–5]. The importance of locality and the role that boundary conditions play in animat simulations are issues that have been unclear in previous work.

We have refined our predator prey model over a number of years and it has been introduced and discussed in several previous publications, for example [6] and [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 control variables of the model and then analysing any new animat behaviours. In particular we have documented fascinating emergent behaviours such as the defensive spirals and other features discussed in [8]. An example of animats forming interesting clusters is shown in Figure 1. This pattern is reminiscent of the well-known Belousov-Zhabotinsky reaction [9, 10] which itself displays significant spatial emergence from localised microscopic properties.

Our simulation usually positions animats in an unbounded plain and in each timestep every animat executes one of an internal set of rules regulating its movement, eating and/or breeding. In a typical simulation, the population expands rapidly across the plain and various macro-behaviours emerge in the process. One problem is that the animat population becomes too large for effective computation, thus limiting the total number of time-steps that can be executed. For this reason we decided to introduce boundaries to limit animat expansion. Our initial conjecture was that the newly imposed boundaries would alter animat behaviour, perhaps dramatically. However, this was not the case and, in fact, the imposition of boundaries had no effect on emergent behaviour but did succeed in limiting the total population.

The reason why boundaries have no observable effects is that our simulation is highly localised and that agent behaviour is driven primarily by local conditions. We believe

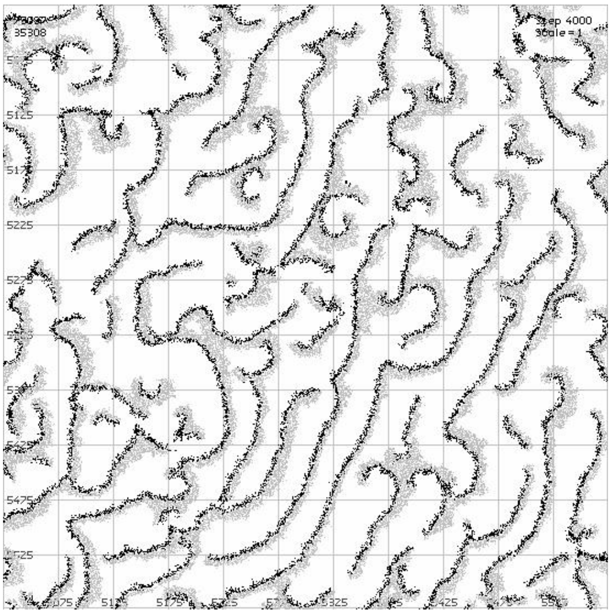

Figure 1: Clusters of animats at step 4000 (predators are black and prey are grey) with a prey crowd value of 10 (more likely to breed). These animats exist in an unbounded plain and the full picture is much larger than that depicted here but the scale has been set to match the other diagrams in this article. 73,097 prey and 35,308 predators are visible in this view. The total numbers at this stage are 570,221 prey and 286,946 predators. Compare this with Figure 7 in which the prey crowd value is 5, forcing prey to noticeably spread out.

that it is important to maximise localisation as this allows animats to interact naturally without constraints and thereby lays the foundations for interesting emergent phenomena which may otherwise be stifled or warped by imposed global conditions such as "reaper functions" which are commonly used in some genetic algorithm based simulations. In practice however, there is always a need for a certain number of global variables and we also discuss these briefly.

A brief overview of the model is provided in section 2. We then discuss the effects of introducing boundary conditions in section 3 and note that the boundaries have no observable effect on animat behaviour. In section 4 we discuss the localisation factors that are the main reason why the introduction of boundaries does not affect the model. We summarise the remaining global factors in the model in section 5. Finally, we suggest some areas for further study in section 6.

2 The Model

The model consists of two species of animat – the predators and the prey – normally existing in a featureless, unbounded plain. Predators need to eat prey to survive and prey need to "graze" to survive. At each step of the model each animat will execute one of a list of rules. These rules often have conditions that must be satisfied before the rule can be executed. Some examples of rules are:

- Move away IF a predator is adjacent. (Prey rule)
- Eat prey IF a prey animat is adjacent AND current health is below a certain value. (Predator rule)

The rules are consulted in an order of priority. The animat always executes the first rule in its list for which the conditions are satisfied. When a new animat is "born" it inherits the rules of its parents. This inheritance could include mutation operators to produce genetic effects. However, in keeping with our philosophy of making very small changes to the model in order to be able to measure the effects, such mutations have not yet been studied. We have experimented with changing the order of priority of the rules and thus produced different sub-groups of animats where each sub-group has the same set of rules but with a different priority order. This work was first published in [6] and different aspects are still under investigation. The model has evolved over the last two years with rules being edited, inserted or deleted in order to achieve maximum localisation. Generally we have found a very rich set of behaviours with a very short (4-7) set of possible rules for each animat.

We have developed a range of quantitative metrics to assist in the analysis of the model and in relating the microscopic input parameters of the animats to the emergent macroscopic properties exhibited. These include measurements of the spatial extent and density of animats; the relative populations of different animats and sub-species; how often animats execute a particular microscopic rule; and the number and type of spatial clusters formed under various conditions.

3 Introducing Boundary Conditions

The early versions of the model placed the animats on an unbounded plain. This ensured that the movement of both individual animats and animat clusters could proceed without artificial constraints. It also ensured that the population could expand freely. However, the unbounded plain with (potentially) infinite population expansion creates problems.

Most simulations involve more than 200,000 animats and on occasion a simulation has reached a million animats. A model with a population of this size takes up an excessive amount of disc space and also becomes unacceptably slow. In previous work [7] we introduced measures to minimize the computation required at each time step. Although this work was successful in speeding up the simulation, it simply delayed the point at which an unacceptable slow down occurs.

We then decided to experiment with introducing boundaries into the model. The boundaries would limit the space into which the population could expand, thus keeping the population to levels at which an unacceptable slow down would not occur. The boundaries were successful in limiting the animat population as shown in the graph in Figure 2.

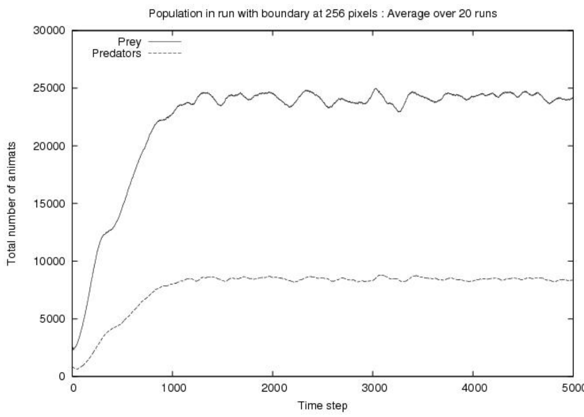

Figure 2: Graph showing how population growth is curtailed in an enclosed circular area with a boundary at a distance of 256 pixels from the origin. The prey population is the larger of the two values at each step.

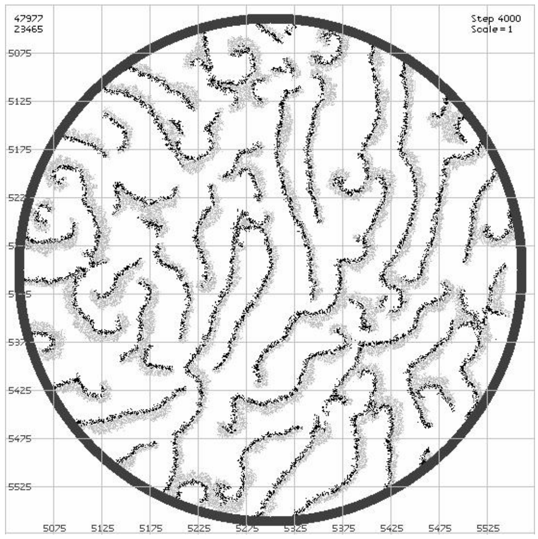

Figure 3: Clusters of animats at step 4000 (predators are black and prey are grey) with a prey crowding factor of 10 (more likely to breed). This model has a boundary at a radius of 256 pixels and the scale matches the other diagrams in this article. Note that the clustering behaviour of the animats appears unaffected by the boundary and remains essentially the same as that depicted in Figure 1.

Prior to the implementation of the boundary conditions, it was conjectured that the boundaries would impact on animat behaviour and cause different macro-behaviours to emerge or possibly stifle emergent behaviour altogether. However, there is no observable change to animat behaviour. For example, the types of clusters formed in Figure 3 (where animats are contained in an enclosed area with a boundary at 256 pixels from the origin) appear to be the same as those appearing in Figure 1 (an unbounded plain with the same prey crowd value of 10). Similarly the behaviour depicted in Figure 4 appears very similar to that depicted in Figure 7, where both simulations have a prey crowd value of 5.

The imposition of boundaries does not alter the general type of clustering that occurs but the boundaries do prevent animats from moving through them. Consequently a model with boundaries in it will never produce exactly the same clustering as a model (with the same parameters) with no boundaries. However, what is clear is that there are no new effects caused by the boundaries. We measured the density (number of animats per unit area) in bands moving outwards from the origin as it was assumed that animats would reach the boundary and then fill up in the outermost ring, possibly rebounding back into inner rings. In fact no significant change was found as shown in Figures 5 and 6.

Thus boundaries can be inserted into a model, thereby significantly reducing the overall population, with no significant difference in animat behaviour or in the type of macro-clusters that emerge. The reason for this is localisation. Every animat constantly reacts to its immediate neighbours and is unaware of the boundaries of the model. Animats that are even a relatively short distance away from the boundary are therefore completely unaffected by it. Animats that are very close to the boundary are prevented from moving in a particular direction, but otherwise continue to interact with their neighbours in the usual way. There is no increase in density at the boundary because the local crowding conditions continue to regulate the population. It is possible that at the moment an animat reaches the boundary and is prevented from moving further, there is a small increase in density. However this means that the animats at the boundary are slightly more crowded and thus will breed less and thereby ensure that the

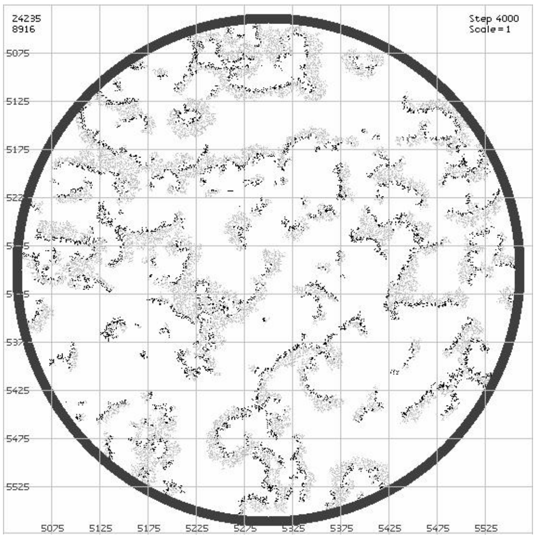

Figure 4: Clusters of animats at step 4000 (predators are black and prey are grey) with a prey crowding factor of 5 (less likely to breed). This model has a boundary at a radius of 256 pixels and the scale matches the other diagrams in this article. Note that the clustering behaviour of the animats appears unaffected by the boundary and remains essentially the same as that depicted in Figure 7.

density rapidly returns to normal. As all population measures are averaged over the total population of thousands of animats, these micro-effects are almost impossible to detect.

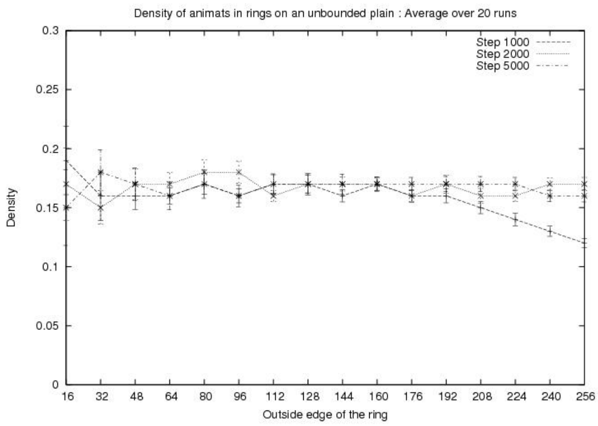

Figure 5: Graph showing the density of animats in successive bands around the origin for animats on an unbounded plain. The drop in the outermost band at time step 1000 is due to the fact that the animats have only just reached this band and it is not yet fully occupied.

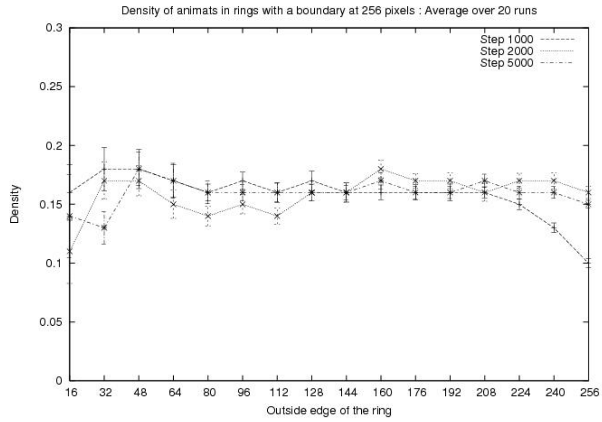

Figure 6: Graph showing the density of animats in successive bands around the origin for animats in an enclosed area with a boundary set at 256 pixels from the origin. The drop in the outermost band at time step 1000 is due to the fact that the animats have only just reached this band and it is not yet fully occupied.

4 Factors that Localise the Model

In this section we list the factors that are responsible for the localisation of the model. These factors are either inherent to each animat (such as its age or health) or are calculated in each time step for each animat (such as its nearest neighbours).

Every animat has a current **age**. When an animat first appears the current age is set to zero. At each time step the current age is increased by one. When the current age reaches the maximum allowable age, the animat is removed from the model. It is possible that age may play a greater role in future versions of the model, for example animats may only be allowed to breed if their current age is within a particular range.

Every animat has a current number of **health** points. Whenever an animat executes an "eat" rule the current number of health points is increased, although it may never become greater than a pre-defined maximum health value. At each time step the current number of health points is decreased by one. If it reaches zero the animat is removed from the model. Several rules have constraints relating to health that must be satisfied before these rules can be executed (see **Rules** below).

Early versions of the model used to initialise a new animat with the maximum number of health points. This often led to strange situations such as a group of predators which could sustain itself with no food supply. This occurred because each generation of (starving) predators would breed, creating a new generation of (maximally healthy) predators, before being removed due to a lack of food. This problem has been solved by changing the "birth process" so that a new animat receives the average number of health points from its two parents. This ensures that starving animats will produce starving children who are then far less likely to be able to breed and produce offspring of their own.

Neighbours are an extremely important part of the localisation of the model. The behaviour of an animat is strongly governed by the species, number and distance of neighbouring animats. In general, animats need to seek prey and/or mates from among their neighbours. Occasionally animats also seek to move away from immediate neighbours to relieve overcrowding, since animats that are crowded will not be able to breed and will have difficulty grazing (in the case of prey animats).

The **rules** that drive the behaviour of each animat are driven by local factors. For example, the Eat Rule for predators can be stated as:

Eat prey IF a prey animat is adjacent AND current health is below a certain value

This means that the predator will only eat (and increase its current health points) if there is an adjacent prey animat and its current health points are below the stated value. Thus the execution of the rule depends on a combination of the location (and species) of neighbouring animats and the animat's own health status.

The health of an animat is important in many of the rules, such as the Eat Rule above. Another example is the Breed Rule which states:

Breed IF another animat of the same species (a mate) is adjacent AND current health is above a certain value.

5 Global Factors

It is impossible to completely localise the model and there are certain factors which remain global. These may possibly be replaced by local factors in future versions of the model although it is likely that some can never be modelled at a local level. Global factors include the following:

When the current age of an animat reaches the **maximum age** it immediately dies due to "natural causes". Every animat has its current age increased by 1 in each time step of the model. Different maximum ages are supplied for each species of animat. It is possible that in future versions of the model, the global maximum age may be removed and replaced by a combination of local factors such as "longevity genes", current health and so on.

This is the **maximum number of health points** that an animat can accumulate. Every animat has its current health decreased by 1 in each time step of the model. If the current health reaches zero the animat immediately dies of "starvation". In any time step in which the animat executes an "eat" or "graze" rule the number of health points will increase, but will never become greater than the maximum health. Different maximum health values are supplied for each species of animat. It is possible that in future versions of the model, the global maximum health values may be removed and replaced by a combination of local factors. However this seems unlikely.

The **maximum vision** is the maximum range at which an animat can detect other animats. Animats need to know the locations of others to be able to seek prey and/or mates. In practice, this maximum is seldom used as most animats have immediate neighbours anyway, although it can occasionally assist isolated animats to move closer to a main cluster. Different maximum vision ranges are supplied for each species of animat in the model – generally predators have a considerably longer vision range than prey. It is possible that maximum vision could be removed and that animats could be allowed to evolve their own maximum vision although it is probable that an absolute maximum would be retained.

The **birth rate** is the chance of a successful birth. Note that all other conditions relating to breeding (adjacent mate, no overcrowding, etc) have to be satisfied before the birth rate check is applied. The birth rate is a percentage, typically in the range 10 to 20 percent. When an animat is about to breed a random number is generated in the range 1 to 100. If this number is less than or equal to the birth rate a new animat appears. Otherwise the breed rule fails. Different birth rates are supplied for each species of animat in the model. It is possible that some local factors (such as health and age) could be used in future to modify the birth rate for individual animats.

The **crowding value** is the number of animats that constitute a crowd. If an animat has this many (or more) immediate neighbours of its own species, it is regarded as overcrowded and can not breed. Different crowd values are supplied for each species of animat. The crowd value probably needs to be retained as a global value as it represents a combination of physical factors (housing materials, water supply, etc) that can not easily be simulated.

Changing the crowd value can make a dramatic difference to the number of animats in the model and also to the type of clustering that emerges. For example, the prey crowd value is 5 in Figure 7. Thus an animat with 5 or more neighbours

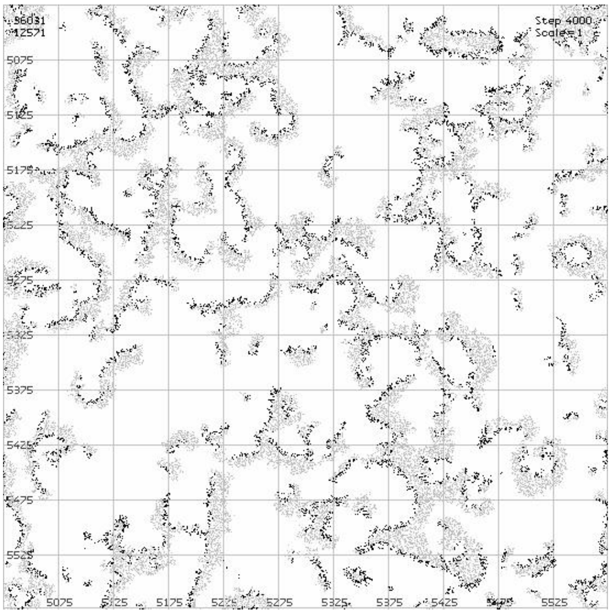

Figure 7: Clusters of animats at step 4000 with a prey crowd value of 5. These animats exist in an unbounded plain and the full picture is much larger than that depicted here but the scale has been set to match the other diagrams in this article. 36,031 prey and 12,571 predators are visible in this view. The total numbers at this stage are 169,194 prey and 59,236 predators.

in this simulation can not breed. This leads to a lower population (a total of 169,194 prey and 59,236 predator animats at step 4000) and also causes animats to spread out giving the diffuse look and "loose" clusters visible in Figure 7.

Figure 7 should be compared with Figure 1 in which a prey crowd value of 10 leads to a larger population (a total of 570,221 prey and 286,946 predator animats at step 4000) and much tighter clusters.

The factors relating to **grazing** apply only to prey as predators do not graze. In future versions of the model all of these factors could be removed and replaced by local "grass" conditions along the lines of the work in [5]. The global grazing factors in the model at the moment are: the grazing crowd value, used in a similar way to the crowd value for breeding, in that an animat with too many immediate neighbours is regarded as overcrowded and finds it increasingly difficult to graze; and the grazing success rate, similar to the birth rate, that stipulates the chance that the grazing will be successful – it can be thought of as representing grass conditions, water supply, etc.

6 Discussion and Further Work

When boundaries were introduced into the model, it was found that the localised nature of the simulation ensured that the boundaries had no observable effect on animat behaviour or on the formation of macro-clusters. This means that we will be able to run future experiments within boundaries and remain confident that any emergent properties will not be affected by the boundaries. This is of immediate and significant benefit because in a bounded model the total population remains constant instead of increasing exponentially. The smaller population will radically reduce the amount of computer space and time to run a model and this will allow simulations to be run for many thousands of time steps and possibly discover new emergent behaviour at a much later stage.

Maximising the local factors in a simulation model provides several advantages. Firstly, it removes anomalies such as predators surviving without a food source by producing healthy offspring, or prey animats forming vast clusters in areas where predators are absent. The removal of these anomalies helps to curtail excessive population growth and distributes animats evenly across the area of interest. A localised model also ensures that any emergent behaviours are not affected by artificial global constraints.

Finally, a localised model allows Artificial Life to more accurately simulate Real Life as real animals do not react to artificial global constants such as a limit on population numbers, since this information is not directly available to them. We conclude that the interplay between boundaries and macroscopic emergent behaviours is both surprising and interesting and may likely be relevant to ther animat agent-based models.

Further work is needed to increase the amount of localisation and to study its effects on the model. This work includes: using age to affect the success of breeding (for example, the very young or very old could be prevented from breeding); maximum health and/or age could be varied for individuals and possibly be passed as "genes" to new generations; and the introduction of "grass" such that health is transferred from the "grass" to prey animats when the graze rule is executed. This would remove the need for the global grazing control factors.

References

[1] 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 Sim-

ulation of Adaptive Behavior, Cambridge, MA: The MIT Press/Bradford Books (1991) 15–21

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

[3] 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

[4] 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

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

[6] Hawick, K.A., James, H.A., Scogings, C.J.: Roles of rule-priority evolution in animat models. In: Proc. Second Australian Conference on Artificial Life (ACAL 2005), Sydney, Asutralia (2005)

[7] 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) Paper 507–102 IASTED, ACTA Press.

[8] 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 (2004) 662–674

[9] Belousov, B.: A periodic reaction and its mechanism. Compilation of Abstracts on Radiation Medicine **147** (1959) 145

[10] Zhabotinsky, A.M.: Periodic processes of malonic acid oxidation in a liquid phase. Biofizika **9** (1964) 306–311