Computational Science Technical Note **CSTN-050**

Energy Flow and Conservation in an Artificial Life Agent Model

C. J. Scogings and K. A. Hawick

2008

Unlike microscopic physics models, agent-based artificial life models are commonly constructed without a strict energy conservation rule. We report on the effects of incorporating energy conservation in our predator-prey based agent model. We discuss the role of microscopic model parameters that give rise to macroscopic emergent effects that can be measured. Arguably non-conserving models like ours can be described as mesoscopic. We show that the effects of an energy conservation law for reproducing animats are no different from having a breeding success regulation rule in our model.

Keywords: animats; artificial life; agents; simulation, energy conservation of energy

BiBTeX reference:

```
@INPROCEEDINGS{CSTN-050,
        author = {C. J. Scogings and K. A. Hawick},
        title = {Energy Flow and Conservation in an Artificial Life Agent Model},
        booktitle = {Proc. 2008 International Conference on Artificial Intelligence (ICAI'08),
                Las Vegas, USA},
        year = {2008},
        pages = {113-119},
        address = {Las Vegas, USA},
        month = {14-17 July},
        publisher = {CSREA},
        institution = {Computer Science, Massey University},
        keywords = {animats; artificial life; agents; simulation, energy conservation
                of energy},
        owner = {kahawick},
        timestamp = {2008.03.12},
        url = {www.massey.ac.nz/~kahawick/cstn/050}
}
```


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

Technical Report CSTN-050

Energy Flow and Conservation in an Artificial Life Agent Model

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

Abstract

Unlike microscopic physics models, agent-based artificial life models are commonly constructed without a strict energy conservation rule. We report on the effects of incorporating energy conservation in our predator-prey based agent model. We discuss the role of microscopic model parameters that give rise to macroscopic emergent effects that can be measured. Arguably non-conserving models like ours can be described as mesoscopic. We show that the effects of an energy conservation law for reproducing animats are no different from having a breeding success regulation rule in our model.

Keywords: animats; artificial life; agents; simulation; conservation of energy.

1 Introduction

The use of "animat" [1] simulations to model emergent behaviour in artificial life systems [2] is now well established – see for example [3–5]. Our predator-prey model has been introduced and discussed in several previous publications, for example [6, 7]. The model consists of two species of "animats" that interact to produce fascinating emergent formations. Figure 1 shows a typical situation after 3000 steps where the animats have formed into clusters and various patterns have emerged. The most interesting of these are the spiral formations discussed in [8]. Some spirals are clearly visible in Figure 1.

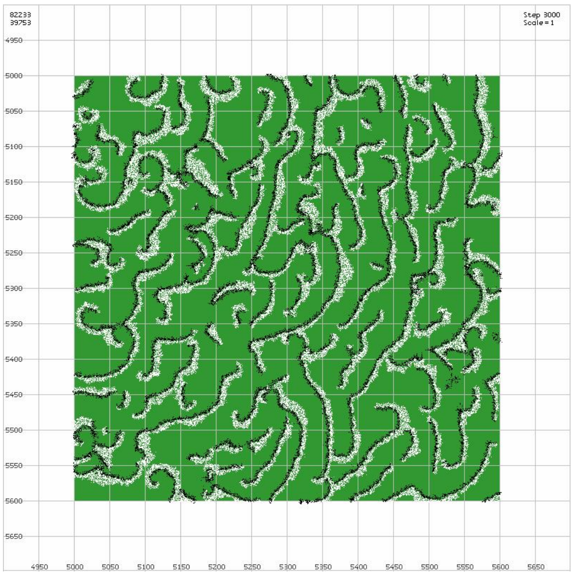

Figure 1: A typical run at step 3000. Predators are black and prey are white and this is the original model (i.e. no animats are losing health when giving birth). The animats inhabit a square "grassy area" with a grass value of 60 which ensures healthy animat populations. Note the typical emergent clusterings, including spirals. The crowd limit for predators is 6 and the crowd limit for prey is 10.

The recent publication of [9] has re-ignited the debate of whether animat simulations should model the conservation of energy during the process of prey consuming grass, predators consuming prey, and so forth. The transfer of energy could be simulated in great detail or could be incorporated into the abstract control parameters which are usually required anyway. Our model does not implicitly track the conservation or transfer of energy but we decided it would be worthwhile to conduct some experiments to see how important this factor is in an animat simulation.

Models based on microscopic physics notions will typically conserve various properties. Statistics models based on the so-called grand canonical ensemble are held in equilibrium with some reservoir both in terms of energy and entity number (e.g. particles). A canonical ensemble model will exchange energy but not particles, and models in the so-called micro-canonical ensemble will exchange neither [10]. It is not trivial to fit typical animat agent based models into this schema. Models such as ours have rules for the number of entities to change (animats are born and die) and for energy to flow through the system allowing local entropy to decrease – as the animats organise themselves into information-rich patterns.

Our model (like many others) is formulated in terms of *plausible* microscopic rules for individual animats, but is arguably more of a mesoscopic model – intermediate between microscopic and macroscopic since the physical space of our system supports multiple animat occupancy. We also model our predator-prey animats as based on a food hierarchy with "grass" as a resource that is continually replenished. In this sense Figure 2 illustrates the role of energy in our model. Energy flows through the system, supporting localised entropy increases.

We decided to investigate the effects of transferring energy from "mother to child" during the birth process and we therefore modified the model such that when a parent animat produced a new animat, the parent animat would lose a certain number of "health points". This can be viewed as a transfer of energy from the parent to the new animat. Experimentation would show exactly how many health points should be lost in this process.

A brief overview of the model is provided in section 2. Section 3 describes how the model was modified so that animats lost health when producing offspring. Section 4 describes a model in which the crowd limit has been adjusted, thus changing the way the animats behave. This model is compared to that in which animats lose health during the birth process. Our conclusions appear in section 5.

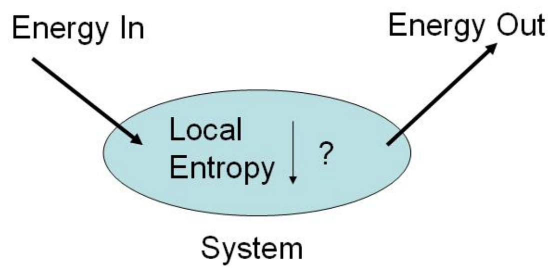

Figure 2: Energy flow through the model system. Energy is not conserved (locally) and thus (information rich) complex patterns can emerge and (local) entropy can decrease in accordance with the second law of thermodynamics.

2 The Model

Our model consists of two species of interacting animats – the predators and the prey. The predators need to eat the prey to survive and the prey need to eat "grass" to survive. The early versions of the model did not require prey animats to eat anything and the concept of "grass" has been recently introduced in [11]. Grass also carries a specific "grass value" and 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 – these results are discussed in [11]. Grass also has a useful side effect in that animats can not exist without it and so the population is limited to the grassy area and can not increase to the point where it becomes unmanageably large.

Figure 3 illustrates the effects of a low grass value on animat populations and should be compared with Figure 1 which has a high grass value. It is interesting to note that despite the low populations, clustering behaviour and spiral formation continue to emerge. Note that Figure 1 provides a snapshot at step 3000 whereas Figure 3 only shows step 1000 – in fact, with the low grass value, all animats were extinct by step 3000.

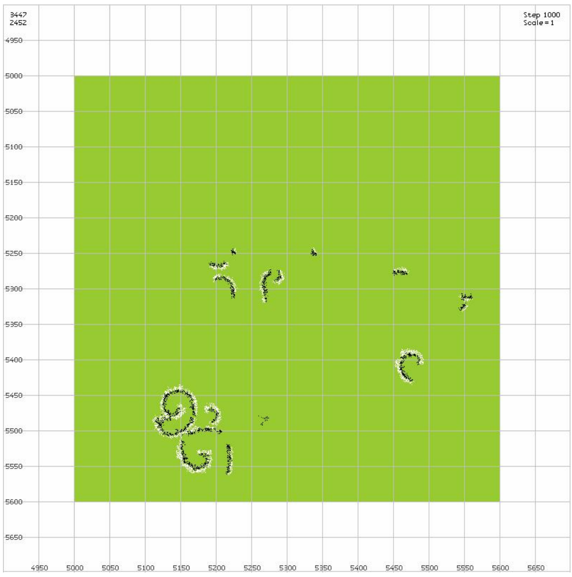

Figure 3: A run at step 1000 with a grass value of 30. Predators are black and prey are white. The low grass value means low populations of both prey and predators, although emergent clustering behaviour (especially spiral formation) is still apparent. The crowd limit remains 6 for predators and 10 for prey. This situation should be compared with Figure 1 in which the grass value is 60.

The simulations discussed here all take place on a large, square“grassed area” with a uniform grass value of 60. A value of 60 is above average and thus ensures that a shortage of food will not affect the population changes recorded in these experiments. The grassed area (and the fact that the animats stay within the area) is clearly visible in the figures – see for example Figure 1.

Each species is allocated a set of control parameters as follows:

Prey:
max age = 20; max health = 100;
birth rate = 40%; crowd limit = 10

Predator:
max age = 50; max health = 200;
birth rate = 15%; crowd limit = 6

Each animat maintains a set of variables including current health, current age, location of neighbours and so on. In each time step, the current health is decreased and the current age is increased. If the health reaches zero the animat “starves” and if the age reaches the maximum age for that species the animat dies of old age. The current health can be increased by eating. Whenever an “eating” rule is successfully executed, the current health is increased but may never exceed the maximum health for that species. The concepts of animats eating and health values have been discussed in [12].

The birth rate is used whenever an animat attempts to execute the “breed” rule and provides the chance that a successful birth will occur. For example, if a prey animat executes the “breed” rule there is a 40% chance that a new prey animat will be produced. The birth rate is a convenient abstract way of simulating a host of factors such as birth complications, disease, adequate shelter for young and so on. In general, prey animats breed more and can live in more crowded conditions than predators but do not live as long and are more prone to starvation.

The crowd limit is the number of animats of the same species that can be immediately adjacent before the animat becomes overcrowded. If animats are overcrowded then they can not easily eat and/or breed and will usually attempt to move away from neighbours in an effort to reduce the crowding. The crowd limit is a an abstract way of simulating factors such as overgrazing, food shortages, stress, fighting over resources and so on.

Thus a high crowd limit (say 15) will allow animats to comfortably coexist with up to 15 adjacent neighbours and will usually cause dense clusters of animats to form. In contrast, a low crowd limit (say 5) will force animats to spread out and will lower the total population in the model. The early versions of the model did not enforce a crowd limit and it was found that prey animats often formed very dense clusters in areas that were free of predators. The introduction of the crowd limit forced the prey to spread out and this, in turn, made it easier for predators to find prey. Thus the introduction of the crowd limit reduced the prey population (speeding up the simulation as less animats had to be managed) and also made the populations more stable over many time steps.

Every animat is also provided with a set of rules and in each time step, each animat executes one of its rules.

Every animat in a particular species has the same set of rules. When a new animat is "born" it inherits the rules of its parents. This rule inheritance process could include mutation operators that would enable rules to evolve over time. 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.

The rule set for prey is:

1. breed if health >30% and mate is adjacent
2. eat grass if health is <70%
3. move towards a prey animat if health >30%
4. move away from a prey animat if health <30%
5. move away from an adjacent predator
6. move randomly to an adjacent position

Rule 3 is included in order to assist animats to find potential mates (of the same species) and Rule 4 is included to reduce overcrowding. If animats are too densely packed they are not able to eat and need to move apart before being able to feed again.

The rule set for predators is:

1. breed if health >50% and mate is adjacent
2. eat adjacent prey if health is <50%
3. move towards another predator if health >50%
4. move towards a prey animat if health <50%
5. move randomly to an adjacent position

The order of each rule set is very important. Every time step, each animat attempts to execute Rule 1. However, most rules have conditions such as "if health >50%", so it is quite possible that Rule 1 can not be executed. If this is the case, the animat attempts to execute Rule 2 and so on. Thus each animat works through the rule set and stops as soon as one rule has been executed. The rule set is therefore a priority list.

We have experimented with changing the priority of the rules within the rule set and have determined that the rule sets listed above are the most effective [13]. 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 set of possible rules for each animat.

The most interesting aspect of the model has been the emergence of macro-behaviours in the form of regular patterns of clusters comprising both species of animat. These clusters are persistent and recognisable across a range of conditions and control variables and are discussed in [14]. 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 Losing Health when Breeding

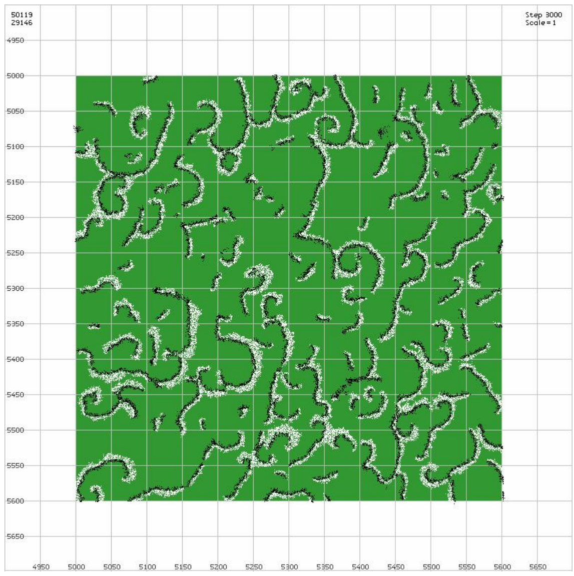

Figure 4: A run at step 3000 with a grass value of 60. Predators are black and prey are white. During this run animats lost one tenth of their maximum health each time they produced offspring. The emergent macro behaviours, including the formation of spirals, is essentially unchanged from the situation at step 3000 of the control – see Figure 1.

Previous versions of the model assumed that animats could breed and produce offspring with no effects on current health or energy. The experiments described in this paper penalise animats by removing a certain number of health points when they produce offspring. We repeated the experiment three times and each time removed a different amount of health. In each case, the amount removed was a proportion of the maximum health available. When we removed a quarter of the maximum health for each birth, the population could not be sustained. Then we removed one tenth of the maximum health for each birth and found that although this did reduce the total population, it made no noticeable difference to the emergent behaviours evident in the model – see Figure 4.

Finally we settled on removing one seventh of maximum health for each birth. Removing health has the predictable effect of lowering the total animat population and this is clearly shown in Figure 5. This graph compares the populations of the prey animats only but the comparison of the predators is similar.

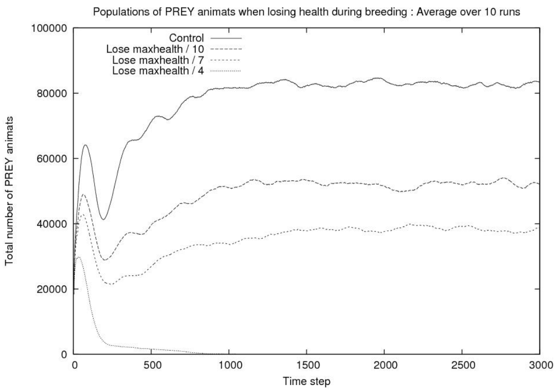

Figure 5: Graph showing how the prey animat population is affected by losing health when breeding. The top line shows the population of the control when no health is removed during breeding. The second line shows the population when losing one tenth of the maximum health per birth. The third line shows the effects of losing one seventh of the maximum health per birth and the bottom line shows that losing one quarter of the maximum health per birth is not sustainable. These populations are averages over ten runs using different random number seeds.

4 The Crowding Factor

The model has always incorporated a "crowd limit" for each species – an abstract value which is used to simulate the general effects of overcrowding on animats. If animats are overcrowded they will not easily eat or breed and will attempt to move away from neighbours to reduce the overcrowding. In the experiments described in section 3 the crowd limit was always set to 6 for predators and 10 for prey.

We decided to experiment by changing the crowd limit for prey animats and the effects of this change can be seen in Figure 6. The top line shows the higher prey population when the crowd limit is 12, the middle line shows the effect of a crowd limit of 11 and the bottom line shows the effect of a crowd limit of 10. In all three cases the loss of health when giving birth is set to one seventh of the maximum health. The bottom line (a crowding factor of 10) is a repeat of the third line in Figure 5. This graph compares the populations of the prey animats only but a comparison of the predators would be similar.

By adjusting the crowd limit we discovered that the loss of (one seventh of the maximum health) health per birth along with a crowd limit of 12 is very similar to no loss of health at all per birth along with a crowd limit of 10. By this we mean that the overall effect on the emergence of macro behaviours is negligible.

Figure 7 shows the situation in a simulation at step 3000 where animats are losing one seventh of maximum health per birth and the crowd limit for prey animats is 12. Emergent macro behaviours are clearly visible and are strikingly similar to the situation in Figure 1 in which animats do not lose any health during the breeding process but the crowd limit for prey is 10 (the control). In particular, the same spiral formations have occurred. Thus the overall effect of a model which attempts to conserve energy are no different from the original model after a single control parameter (the crowd limit) has been adjusted slightly.

5 Discussion and Conclusions

We have presented details of our animat model for predator-prey systems where energy influx is controllable through a global "grass" parameter. We have discussed the effects that conservation of energy (or health/food-points) has on the system.

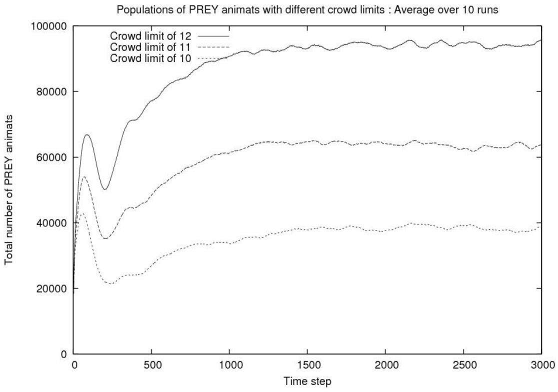

Figure 6: Graph showing how the crowding factor affects the total population. The top line shows the prey population with a crowd limit of 12, the middle line shows the effect of a crowd limit of 11 and the bottom line shows a crowd limit of 10. In all three cases animats are losing one seventh of the maximum health when giving birth. These populations are averages over ten runs with different random number seeds.

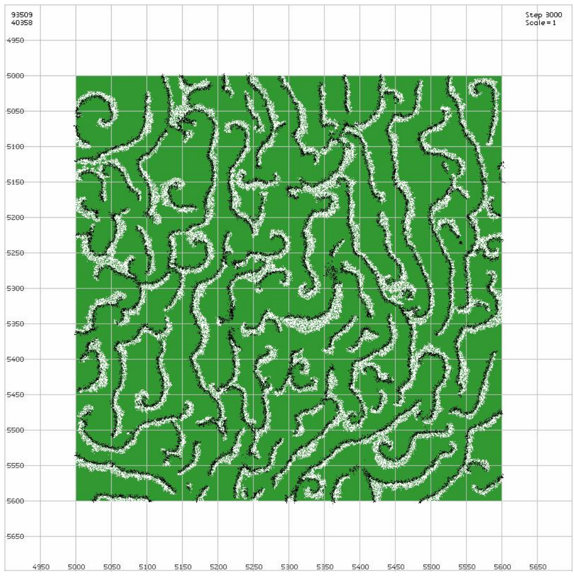

Figure 7: The position of animats at step 3000 in a simulation where animats lose one seventh of the maximum health per birth and the crowd limit for prey animats is 12. Predators are black and prey is white. The emergent macro behaviours, especially the formation of spiral clusters, is very similar to that in Figure 1 (repeated below) in which animats do not lose any health per birth and the crowd limit for prey is 10.

Penalising animats by removing health when they give birth certainly has an effect on the model. The greater the health removed, the lower the total population and, in extreme cases, the population dies off completely.

However, the use of crowd limits to simulate a host of smaller factors including climate, overgrazing, scarcity of nesting materials and the like, means that the loss of health per birth is unnecessary. The same overall effect can be achieved simply by adjusting the crowd limits.

In conclusion, it is unnecessary to attempt to introduce measures to simulate the conservation of energy in a model of this type. The effects of any such measures can be simulated by using simple abstract control parameters such as crowd limits.

We observe that in general it is not easy to make a direct comparison between a microscopic physics model formulation with explicit conservation laws and our animat oriented model which supports the definite notion of (non-conserved) energy flow through the system. We offer the suggestion that to avoid controversy such animat models might usefully be termed mesoscopic than microscopic, so as not to clash with the common usage of microscopic in the physics community.

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

[2] Ferber, J.: Multi-Agent Systems - An Introduction to Distributed Artificial Intelligence. Addison-Wesley (1999)

[3] 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

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

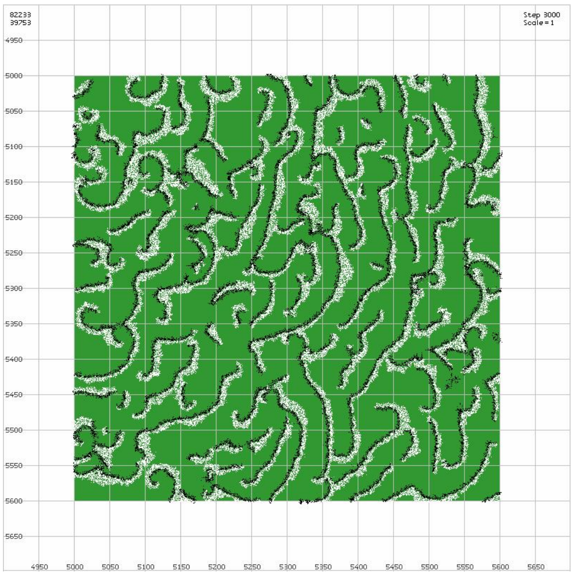

Figure 8: Figure 1 repeated to allow easy comparison with Figure 7. This is the control model in which animats do not lose any health per birth and the crowd limit for prey is 10.

[5] 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

[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, Australia (2005) 99–116

[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) 90–95

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

[10] Kadanoff, L.P.: Statistical Physics Statics, Dynamics and Renormalization. Number 981-02-3764. World Scientific (2000)

[11] Scogings, C.J., Hawick, K.A.: Global constraints and diffusion in a localised animat agent model. Technical report, Computer Science, Massey University (2007) CSTN-047.

[12] Scogings, C.J., Hawick, K.A., James, H.A.: Tuning growth stability in an animat agent model. In: Proceedings of the 16th IASTED International Conference in Applied Simulation and Modelling (ASM 2007). (2007) 312–317

[13] 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

[14] Hawick, K.A., James, H.A., Scogings, C.J.: A zoology of emergent patterns in a predator-prey simulation model. In Nyongesa, H., ed.: Proceedings of the Sixth IASTED International Conference on Modelling, Simulation, and Optimization, Gabarone, Botswana (2006) 84–89