

Altruism Amongst Spatial Predator-Prey Animats

Chris Scogings and Ken Hawick

Computer Science, Institute of 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

Understanding the emergence or suppression of altruism is an important step towards understanding real-life many-agent systems. We explore the relative survival traits of spatial animats in our predator-prey model and find some quantifiable emergent advantages of altruistic behaviour on the part of individual animats.

Keywords: altruism, animats, predator-prey.

Introduction

“Although a high standard of morality gives but a slight or no advantage to each individual man and his children over the other men of the same tribe... an advancement in the standard of morality will certainly give an immense advantage to one tribe over another.”

Charles Darwin, The Descent of Man, 1871

The evolution of altruism has been debated for several decades and is still a source of some argument. On the one hand we have those who aspire to the “selfish gene” theory of Dawkins (1976) and regard altruism as (at best) negligible. On the other hand we have the approach of Sober and Sloan Wilson (1998) who state that it is possible for “pure altruism” (helping unrelated individuals and receiving no payback) to evolve. In between lie a range of possibilities including kin selection and group selection (Maynard Smith, 1964). Of particular interest is the recent work on the evolution of strong altruism in randomly formed groups (Fletcher and Zwick, 2004) and this paper also provides an excellent review of the current altruism debate. Many of the theories mentioned above use models to strengthen their position. However, most of these models are analytical and none use the “animat” (Wilson, 1991) approach as we have done here.

It appears that Darwin, quoted above in (Wilson and Wilson, 2007), believed in unadulterated group selection. This is the idea that individuals within a group will behave altruistically towards others so that the group as a whole will prosper. More on Darwin’s position can be found in (Sober

and Sloan Wilson, 1998). The research we present in this paper uses an animat model to explore whether this approach is viable within an Artificial Life simulation. We believe that our findings add strength to Darwin’s original suggestion – namely that an altruistic group will have an advantage over a selfish group even though selfish individuals have an advantage over altruistic individuals within the group. However, this will only occur under certain conditions.

Complex systems using spatial agents or animats have existed for some time – see for example (Tyrrell and Mayhew, 1991; Adami, 1994; Holland, 1994) – and have yielded rich insights into emergent group behaviour in physical, biological and sociological simulation settings. Experimentation in this area is ongoing as illustrated in (Ronkko, 2007).

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 (Hawick et al., 2005b; Scogings et al., 2006). 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 (Hawick et al., 2004). An example of these features is provided in Figure 1. This shows some typical (and highly robust) wave front and proto-spiral pattern generation behaviours in our model. We have been able to study these in a quantifiable manner by applying automatic feature detection techniques to the spatial animat patterns (Hawick et al., 2005a).

This paper consists of the following sections: A brief overview of our predator-prey simulation; a discussion of how we introduced altruistic behaviour into the model; experimental runs of the model simulating “good times” (i.e. a grass value conducive to high prey population growth); experimental runs of the model during “bad times” (with a much lower grass value); and finally, a brief summary and conclusion.

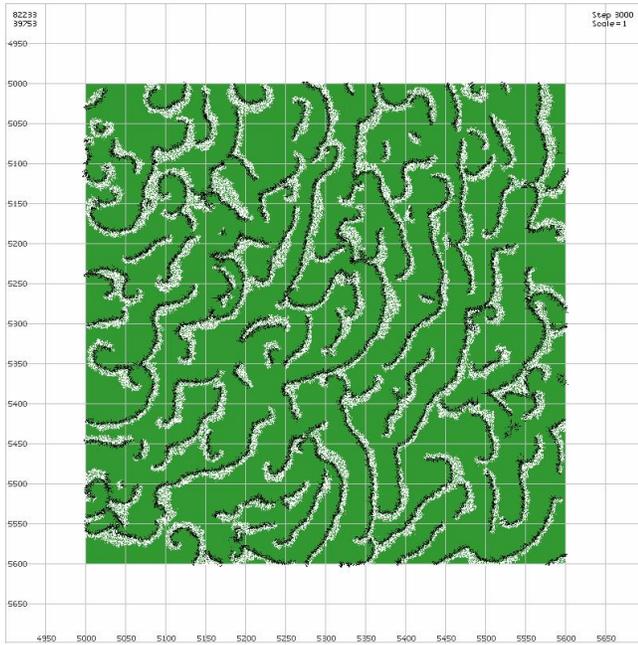


Figure 1: A typical run at step 3000. Predators are black and prey are white and all animats are selfish (the original model). The animats inhabit a square “grassy area” with a grass value of 60 which ensures healthy animat populations. Note the typical emergent clusterings, including spirals.

The Model

Our model consists of two species of interacting animats – the predators and the prey. Animats have a very simple state: a health variable; an age; and a spatial position in the 2-dimensional square mesh world. The food chain is therefore very simple as shown in Figure 2. Our system is an open one as far as energy is concerned. Prey consume “grass” which is assumed to be continually replenished although we can adjust this rate and can also adjust the spatial pattern of grass and hence the underpinning geometry of the flat world. Predators consume only prey and other things being equal we can reproduce the well known boom-bust limit cycles predicted by predator-prey models such as the Lotka-Volterra coupled differential equations (Lotka, 1925; Volterra, 1926) and their spatial variants (Gallego, 2003).

The simulation runs as a sequence of discrete and synchronous time-steps and in each time-step the following operations are performed for every animat:

- (Phase 1) health check
- (Phase 1) age check
- (Phase 1) locate neighbours
- (Phase 1) execute one rule
- (Phase 2) update all variables

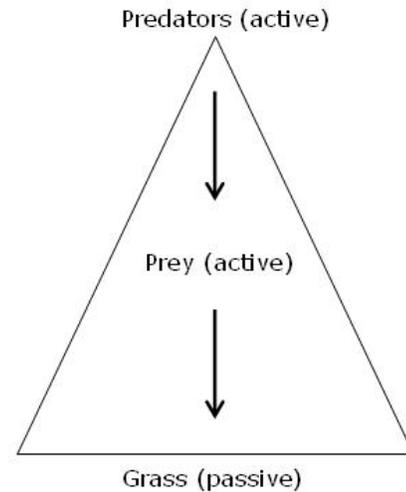


Figure 2: Model participants. Predators and Prey are spatial animat agents, whereas grass is a passive (continuously replenished) resource.

Although the model is synchronous animats are updated in a random order, which we found adequate to remove any spatial artefacts from sweep order. The process is a two-phase system in which the variables for all animats are updated after all checks have been made and all rules have been executed. The two-phase system was developed in order to ensure fairness across all the animats in the model and a full discussion of alternative updating systems is available in (James et al., 2004).

Every animat carries a small set of microscopic rules that govern its behaviour and this rule set is passed on unchanged to any offspring. It is possible to allow mutations to rules and to introduce genetic algorithms into the model but an important feature of our work is to make small, well-defined changes to the microscopic model and measure the effects of those changes. We have experimented with changing the order (priorities) of the rules and have investigated which rule sets generate the most successful animat groups (Hawick et al., 2005b). This approach is a good way of quantitatively investigating the microscopic rule space without getting lost in the combinatoric explosion that hinders a more simple minded evolutionary “suck it and see” investigation.

In this series of experiments, the rule set for predators is:

1. if well fed – breed with an adjacent predator
2. if hungry – eat an adjacent prey animat
3. if well fed – move towards another predator
4. if hungry – move towards prey
5. move randomly

and the rule set for prey is:

1. if well fed – breed with an adjacent prey animat
2. if hungry and not crowded – eat grass (if available)
3. if well fed – move towards another prey animat
4. if hungry and crowded – move away from other prey
5. move away from an adjacent predator
6. move randomly

Breeding only has a certain chance of success. This is a simple alternative to factoring in a host of complicated parameters including birth defects, nutrition, adequate shelter and so on. For these experiments the chance of a successful birth was set to 15% for predators and 40% for prey. The prey conditions involving crowding were introduced to prevent prey forming enormous clusters in any area of the grid that happened to be temporarily free of predators. If a prey animat has k or more adjacent neighbours, it is deemed to be “crowded” and can not eat grass (an abstract simulation of “over grazing”). For these experiments k was set to 10.

Rules are considered in a strict priority order. Each time-step, every animat attempts to execute the first rule in its rule set. However, most rules have conditions so can often not be executed. For example, prey will only move away from a predator if a predator is actually adjacent. If the conditions for the first rule can not be satisfied, the animat attempts to execute the next rule in the set and so on. This Markov chain mechanism of rules is described in detail in (Hawick et al., 2007).

All animats in the model have a “current health” value. This value (in some ways analogous to “internal energy”) is reduced each time-step and if it reaches zero the animat “starves to death”. If an animat eats something (predators eat prey and prey eat “grass”) then the current health value will be increased by a certain amount, although it may never be increased past the maximum health value which is predetermined for each animat species. A “well fed” animat (see the conditions in the rules above) has a current health value of two-thirds or more of the maximum. A “hungry” animat has a current health value of less than one-third of the maximum. The concepts of health values and animats eating are discussed in (Scogings et al., 2007).

Early versions of the model did not require prey animats to eat anything and the concept of “grass” has been recently introduced. Grass can be placed in specific locations on the map and each grassy area carries a specific “grass value”. When a prey animat eats the grass, its current health is increased by the grass value. This means that animats will do well on grass with a higher value and will struggle to survive on grass with a lower grass value. Grass therefore has a useful side effect in limiting the animat populations to the

grassy area and preventing them becoming unmanageably large. This is important as our model boundaries are effectively open ones with no periodicity or hard reflecting walls. The wilderness is in fact a good way to limit the population without introducing spurious spatial artifacts due to boundaries. Animats can (and do) diffuse out into the wilderness and quietly starve which is less perturbative to the model’s remaining and ongoing population than if they were corralled or transported spatially across a periodic boundary.

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. Note that we try to keep all microscopic parameters very simple and usually express them where possible as percentages. 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. 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 (Hawick et al., 2006).

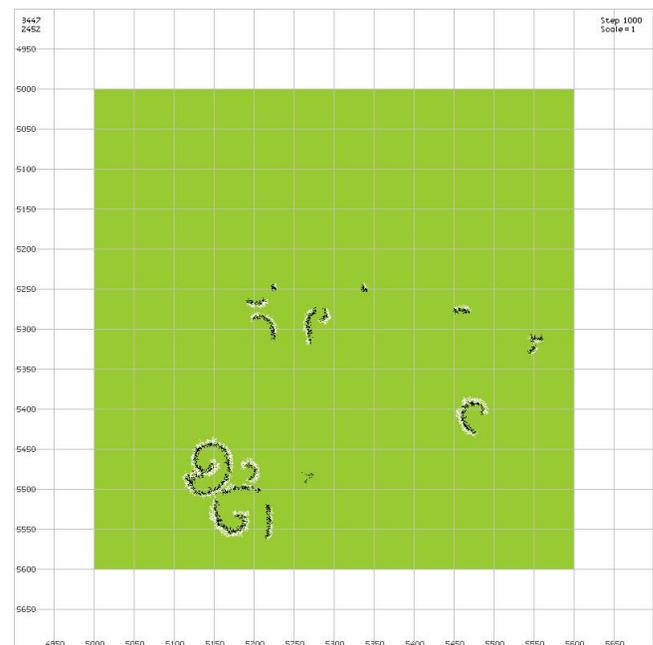


Figure 3: A run at step 1000 with a (uniform) grass value of 30. Predators are black and prey are white. This “low” grass value means low populations of both prey and predators, although emergent clustering behaviour (especially spiral formation) is still apparent. This situation should be compared with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

Introducing Altruism into the Model

How to define altruistic behaviour is a central part of the general debate on altruism and a good review of this can be found in (Fletcher and Zwick, 2004). Most models use a variant of Hamilton's equation (Hamilton, 1964) but our model is not analytical so requires a form of altruism represented by a behavioural rule rather than an analytical equation. We decided that this should take the form of sharing food resources as these are most vital to animats' ongoing existence. (In fact, the sharing of food resources is probably a reasonable way to measure altruism among humans in the world today.) We also decided to only study altruistic predators as the "higher life form" in the model. As prey always have equal access to "grass" we felt that there was not much to be gained in studying the behaviour of altruistic prey. Therefore, in order to introduce altruism into the model we decided to regard predators' current health value as "transferable currency". Some predators are marked as altruistic whereas others are marked as selfish – i.e. they remain unchanged from the original model.

Just before the conclusion of phase 1 of each time-step an altruistic predator checks to see if it has an adjacent predator that has a lower current health value than its own. If this is the case, the current health values of the two predators are added together and each animat receives one half of this total value. In other words, an altruistic animat shares its current health with a neighbour (regardless of whether the neighbour is altruistic or selfish) such that the two animats receive equal shares and the current health of the altruistic animat will be reduced while the current health of the neighbouring animat will increase. Note that the altruistic animat will only share with an adjacent neighbour that is worse off than itself. No attempt is made to search for other neighbours to share with, and if several animats are both adjacent and worse off, one is chosen at random to share with.

When animats breed and produce offspring the new animats are precise clones of their "mothers". Two adjacent animats are required to breed but only one actually executes the rule to breed – this animat is known as the mother. This aspect of inheritance has existed in the model for some time and has been discussed in (Hawick et al., 2005b). Altruistic animats carry a "gene" that marks them as altruistic. Thus when altruistic animats breed they produce new altruistic animats (clones of themselves). This means that it is possible to track the success of the altruistic population versus the selfish population over time. We can then consider the question: is an altruistic society more successful than a selfish society?

Experiments 1, 2 and 3 : Good Times

In this section we describe the population variations that result from introducing altruistic predators into the system. Results are shown as animat populations plotted over time.

The populations show the well known "boom-bust" phase variation that is typically found in any predator prey system. A boom in prey is followed (with a suitable phase lag) by a boom in predators. The cycle continues when predators "boom" and cause a subsequent "bust" in the numbers of prey. Spatial variations complicate this situation considerably from the simple limit cycles encountered in dimensionless Lotka-Volterra-like systems, and our model generally displays a equilibrating epoch (usually less than 1000 time steps) during which the overall population can grow or shrink drastically. We typically find however, for a remarkably wide range of microscopic animat parameter values, that a long term stable epoch subsequently arises during which the variations are due to emergence and interaction or spatial patterns such as clumps, wave fronts and spirals. It seems valid therefore to draw robust conclusions about altruism in the long term.

The first three experiments simulated a square grassed area with a grass value of 60 (which is regarded as above average). Thus the prey had no problem finding food and consequently predators always had easy access to prey. In experiment 1 there were no altruistic animats. Thus this simulation matched previous normal conditions and could be used as the control. The situation at step 3000 is depicted in Figure 1 and shows typical densities of predators and prey and the usual emergent clustering behaviour. The animat populations for this experiment are shown in the graph in Figure 4 as the second line (prey) and the fourth line (predators).

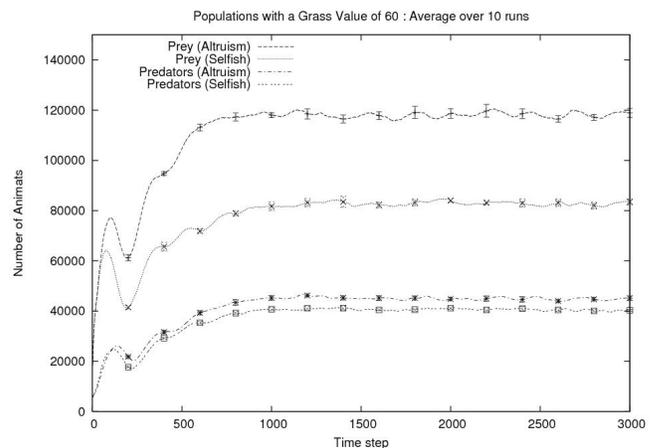


Figure 4: Plot showing total populations during experiments 1 and 2 when the grass value is 60. The populations from experiment 1 (all selfish) appear as the second line (prey) and the fourth line (predators). The populations from experiment 2 (altruistic predators) appear as the first line (prey) and the third line (predators). It is clear that the altruistic predators have succeeded in increasing their own population but more importantly are also "conserving" their prey food-source.

In experiment 2 all predators were designated as altruistic. This caused a slight increase in the total population of predators but an even greater increase in the population of prey, as shown by the first line (prey) and the third line (predators) in the graph in Figure 4. The situation at step 3000 is depicted in Figure 5 and shows much denser formations of prey animats because the altruistic predators are conserving their food source. It is interesting to note that the general pattern of emerging clusters, including spirals, has not been affected by the altruistic nature of the predators.

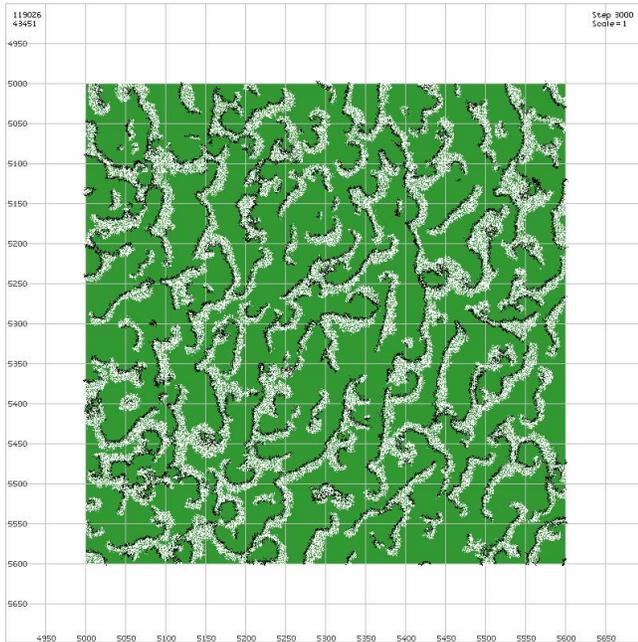


Figure 5: The situation at step 3000 during a typical run of experiment 2 with a grass value of 60. All predators (black) are altruistic and are conserving the prey population (white) which leads to formations of prey animats that are noticeably more dense than those in Figure 1 which depicts the situation for selfish predators at the same step and with the same grass value.

Remember that only predators are altruistic and there are no altruistic prey animats. Thus altruistic predators achieve two things: firstly they rescue some of their own number who may otherwise have starved to death; and secondly they also “conserve prey”. This happens because predators only eat when hungry. Thus if a predator receives some health points from an altruistic neighbour, it has no need to eat prey itself. So it appears that an “altruistic society” not only ensures that its own population will be greater than the selfish equivalent, but it also conserves food reserves for future generations.

Alas, while an altruistic society may be more successful than a selfish one, this is not obvious to the individuals within that society. Experiment 3 commenced with 50%

of the predators randomly designated as altruistic while the remaining 50% were designated as selfish. Thus two competing groups of predators were created. With the passage of time, it became clear that the altruistic predators could not compete with the selfish predators and the altruistic population died out completely. This is shown in Figure 6. Thus, although an altruistic society may be more successful as a group, altruistic individuals can not compete successfully with selfish individuals. However, perhaps this is only the case when times are good and food is plentiful.

All three experiments were run ten times with different random number seeds and the population graphs represent the averages from the ten simulations.

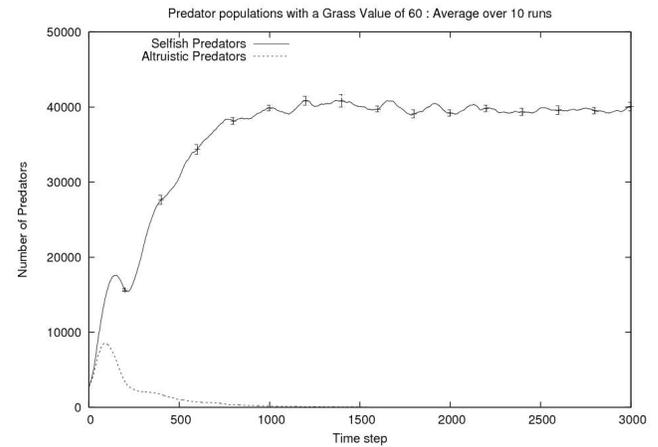


Figure 6: Graph showing predator population when the grass value is 60. The run starts with approximately equal numbers of altruistic and selfish predators but the selfish group soon comes to dominate the model.

Experiments 4, 5 and 6 : Bad Times

The second set of experiments used exactly the same procedure as the first but this time was conducted with a grass value of 30 which is below average and means that both prey and predators struggle to survive. This was made apparent in experiment 4 (the control) in which populations of both prey and predators are markedly decreased as shown in Figure 3 and in fact most animats had starved to death by step 3000. No animats were altruistic in this run. The populations of experiment 4 appear as the third line (prey) and the fourth line (predators) in Figure 7.

In experiment 5 all predators were once again designated as altruistic. The results of this run were significant and are shown in Figure 7 where the first line represents the prey population and the second line is the altruistic predator population. The altruistic animats once again were able to “conserve” the prey resource and thus were able to maintain a stable population of both predators and prey, whereas the

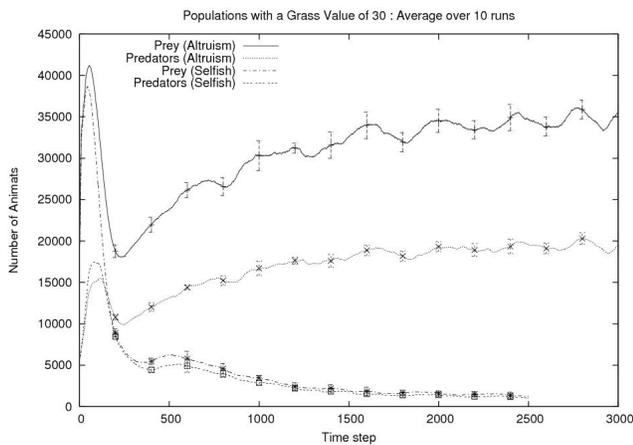


Figure 7: Graph showing total animat populations during experiments 4 and 5 with a grass value of 30. The lower two lines represent the prey and selfish predator populations in experiment 4. These lines are unreliable after step 2500 as the populations died out soon after that point in most of the runs. The top two lines represent prey and altruistic predator populations in experiment 5. This graph dramatically illustrates the success of altruistic predators at sustaining both their own population and that of their prey.

behaviour of the selfish predators in experiment 4 caused both prey and predator populations to crash. Thus in times of adversity when food resources are scarce, the tendency of the altruistic society to conserve the prey population allows it to survive indefinitely and gives it an enormous advantage over the selfish society. The success of the altruistic society can be seen by comparing Figure 3 with Figure 8 (both at step 1000 with a grass value of 30) which clearly shows the advantages of conserving the food-source.

However we are still faced with the problem that individuals within a society appear to tend towards selfishness. We therefore performed experiment 6 with 50% of the predators randomly designated as altruistic while the remaining 50% were designated selfish. This time, because of the low grass value and the scarcity of prey, the altruistic group prospered and came to dominate the model. The population graph for experiment 6 is shown in Figure 9.

Once again, all experiments were run ten times with different random number seeds and the population graphs represent the averages from the ten simulations.

Summary and Conclusions

We have described our spatial predator-prey model in terms of its microscopic parameters and some of the emergent spatial patterns it generates. We have presented some experimental results based on averaged population trends over time, linked to some typical configuration snapshots. These

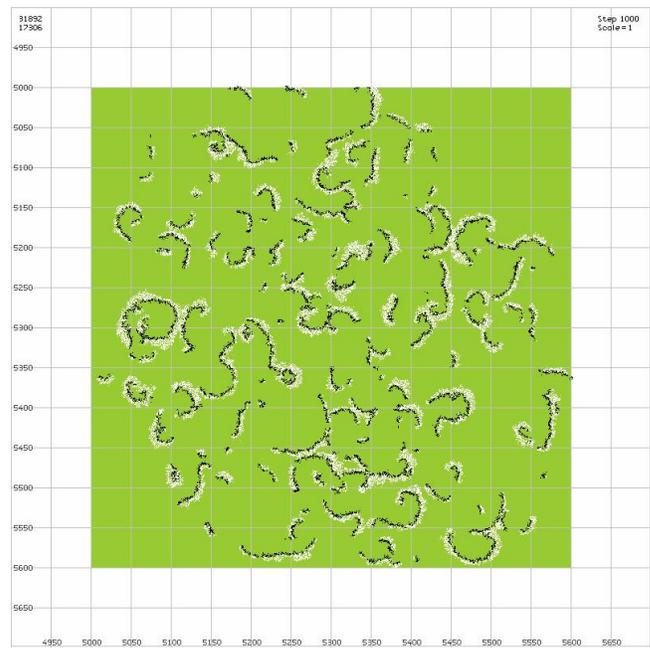


Figure 8: Experiment 5 at step 1000 with a grass value of 30. Predators are black and prey are white. The predators in this run are all altruistic and are doing remarkably better than the selfish predators (i.e. the original model) depicted in Figure 3 also at step 1000 with the same grass value.

suggest convergent and consistent results even over widely differing microscopic start conditions. Our model allows us to study the macroscopic properties of interacting “spatial herds” of individual animats. We are therefore able to study collective behaviour amongst the differing types of group or herd. We have tried in particular to explore the effects of altruism introduced at the microscopic level but manifesting itself as macroscopic success/failure levels.

We argue that this series of experiments supports quantitatively Darwin’s assertion that an altruistic society will always do better than a selfish one even though selfish individuals within the society will tend to dominate. The experiments indicate that the main reason for this might be that altruistic societies conserve their (food) resources. This behaviour is indicative of true group selection and not kin selection as altruistic predators have no way of recognising kin and secondly, they share resources with any other predator including selfish ones.

The experiments also suggest that altruistic societies are particularly successful in times of hardship when they can continue to survive while selfish communities may vanish completely. In addition, during these times of adversity the selfish individuals within the altruistic society will not succeed and dominance will be retained by altruistic individuals. There are a number of shared resources issues related to “The Tragedy of the Commons” (Hardin, 1968). Our an-

imats do not yet have a direct sense of altruism that governs their commons, but we hope to incorporate this into a future version of the predator animats.

The use of an Artificial Life model in a study of an abstract notion such as altruism highlights the benefits that research in this area can deliver. The model was not originally developed to analyse concepts such as altruism but we have demonstrated that it can easily be adapted to this, and (we expect) to other similar uses.

The concept of the altruistic society is itself an emergent macro-behaviour as individual altruistic animats do not attempt to assist their entire society but merely one adjacent neighbour and even then, only when that neighbour has a lower health value than themselves. Thus is the concept of a complex altruistic society based on very simple rules executed by the individuals within that society.

We are presently considering how prey animats could also exhibit a form of altruism - for example by individuals sacrificing themselves to prey to save their fellows. We suspect this effect may give rise to an interesting interplay and further rich spatial behaviours.

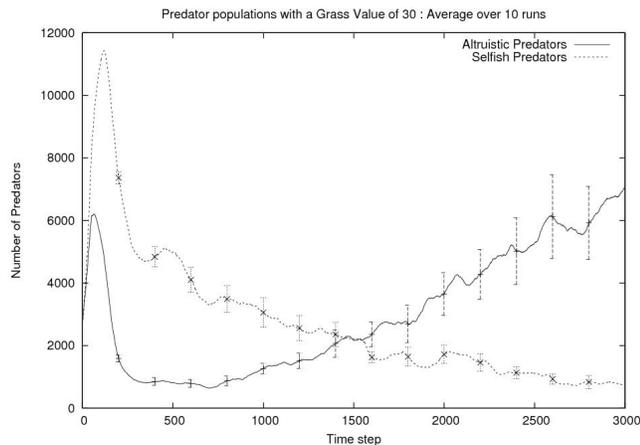


Figure 9: Graph for experiment 6 showing predator population when the grass value is 30. The run starts with approximately equal numbers of altruistic and selfish predators but eventually the altruistic group becomes dominant. The standard deviations (shown as error bars at regular intervals) are larger than usual for altruistic predators because the slope of the line after recovery (at about step 700) varied with each run. This does not change the general trends when comparing the two populations. This graph should be compared with that in Figure 6 where the selfish predators dominated because the grass value of 60 was significantly higher.

References

- Adami, C. (1994). On modeling life. In Brooks, R. and Maes, P., editors, *Proc. Artificial Life IV*, pages 269–274. MIT Press.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford Univ. Press. ISBN 019857519X.
- Fletcher, J. A. and Zwick, M. (2004). Strong altruism can evolve in randomly formed groups. *Journal of Theoretical Biology*, 228:303–313.
- Gallego, S. (2003). *Modelling Population Dynamics of Elephants*. PhD thesis, Life and Environmental Sciences, University of Natal, Durban, South Africa.
- Hamilton, W. D. (1964). The evolution of social behaviour. *Journal of Theoretical Biology*, 1:295–311.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859):1243–1248.
- Hawick, K. A., James, H. A., and Scogings, C. J. (2005a). Manual and semi-automated classification in a microscopic artificial life model. In *Proc. Int. Conf. on Computational Intelligence (CI'05)*, pages 135–140, Calgary, Canada.
- Hawick, K. A., James, H. A., and Scogings, C. J. (2005b). Roles of rule-priority evolution in animat models. In *Proc. Second Australian Conference on Artificial Life (ACAL 2005)*, pages 99–116, Sydney, Australia.
- Hawick, K. A., James, H. A., and Scogings, C. J. (2006). A zoology of emergent patterns in a predator-prey simulation model. In Nyongesa, H., editor, *Proceedings of the Sixth IASTED International Conference on Modelling, Simulation, and Optimization*, pages 84–89, Gabarone, Botswana.
- Hawick, K. A., Scogings, C. J., and James, H. A. (2004). Defensive spiral emergence in a predator-prey model. In Stonier, R., Han, Q., and Li, W., editors, *Proc. Complexity 2004*, pages 662–674, Cairns, Australia.
- Hawick, K. A., Scogings, C. J., and James, H. A. (2007). Stochastic animats in complex spatial environments. Technical Report CSTN-040, Computer Science, Massey University.
- Holland, J. H. (1994). Echoing emergence: Objectives, rough definitions, and speculations for echo-class models. In Cowan, G. A., Pines, D., and Meltzer, D., editors, *Complexity: Metaphors, Models and Reality*, pages 309–342. Addison-Wesley, Reading, MA.

- James, H. A., Scogings, C. J., and Hawick, K. A. (2004). Parallel synchronization issues in simulating artificial life. In Gonzalez, T., editor, *Proc. 16th IASTED Int. Conf. on Parallel and Distributed Computing and Systems (PDCS)*, pages 815–820, Cambridge, MA, USA.
- Lotka, A. J. (1925). *Elements of Physical Biology*. Williams & Williams, Baltimore.
- Maynard Smith, J. (1964). Group selection and kin selection. *Nature*, 201:1145–1147.
- Ronkko, M. (2007). An artificial ecosystem: Emergent dynamics and lifelike properties. *J. ALife*, 13(2):159–187.
- Scogings, C. J., Hawick, K. A., and James, H. A. (2006). Tools and techniques for optimisation of microscopic artificial life simulation models. In Nyongesa, H., editor, *Proceedings of the Sixth IASTED International Conference on Modelling, Simulation, and Optimization*, pages 90–95, Gabarone, Botswana.
- Scogings, C. J., Hawick, K. A., and James, H. A. (2007). Tuning growth stability in an animat agent model. In *Proceedings of the 16th IASTED International Conference in Applied Simulation and Modelling (ASM 2007)*, pages 312–317.
- Sober, E. and Sloan Wilson, D. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behaviour*. Harvard University Press.
- Tyrrell, T. and Mayhew, J. E. W. (1991). Computer simulation of an animal environment. In Meyer, J.-A. and Wilson, S. W., editors, *From Animals to Animats, Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 263–272.
- Volterra, V. (1926). Variazioni e fluttuazioni del numero d'individui in specie animali conviventi. *Mem. R. Accad. Naz. dei Lincei, Ser VI*, 2.
- Wilson, D. S. and Wilson, E. O. (2007). Evolution: Survival of the selfless. *New Scientist*, 2628:42–46.
- Wilson, S. W. (1991). The animat path to AI. In Meyer, J.-A. and Wilson, S. W., editors, *From Animals to Animats 1: Proceedings of The First International Conference on Simulation of Adaptive Behavior*, pages 15–21. Cambridge, MA: The MIT Press/Bradford Books.