

EcoPS - A Particle Swarm Algorithm to Model Group-Foraging

Cecilia Di Chio
Department of
Computer Science
University of Essex, UK
cdichi@essex.ac.uk

Riccardo Poli
Department of
Computer Science
University of Essex, UK
rpoli@essex.ac.uk

Paolo Di Chio
Dipartimento di Sistemi e
Istituzioni per l'Economia
University of L'Aquila, Italy
pdc@ec.univaq.it

ABSTRACT

Recent work has introduced a simulation model of ecological processes in terms of a very simple Particle Swarm algorithm. This abstract model produced qualitatively realistic behaviours, but do these results hold up in a model constrained by more plausible biological assumptions? The objective of this paper is to answer this question.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search; I.6.3 [Simulation and Modeling]: Applications

General Terms

Performance

Keywords

Particle Swarm, Swarm Intelligence, Artificial Life, Adaptive Behaviour, Biological Application

1. INTRODUCTION

Despite its socio-biological background, the field of applications for Particle Swarm (PS) systems [6] has mainly been the optimisation of nonlinear functions. In [3] and [4], a new way to use PSOs for the simulation of ecological processes has been introduced. This focuses on the *group-foraging* problem, which so far has been largely restricted to the field of individual-based modelling in ecology (for a review, see [2]).

Behavioural ecology is the branch of evolutionary biology which studies the ecological and evolutionary basis for animal behaviour (i.e., what are the “historical” reasons for certain animal behaviour and what role does an animal’s behaviour play in allowing it to adapt to its environment).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO’07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

Being such a widespread phenomenon in the animal kingdom, group-living is certainly one of the most studied complex animal behaviours [7]. Two general requirements for grouping behaviour are:

1. individuals have to be close to each other in space and time, and
2. animals must show *social attraction* (i.e., they have to “actively” seek to be close to each other, instead of simply meeting at a certain point because of the attraction to environmental conditions at that point).

Amongst the variety of grouping behaviours, *group-foraging* is one of the most interesting. Some of the most likely hypotheses for *group-foraging behaviour* [7, 9] are:

1. aggregations find more food more quickly than individuals do, and so animals in a group feed more effectively;
2. animals in bigger groups can allocate more time to feed and less to look for predators;
3. by observing the behaviour of other members of a group, animals can gain useful information.

Foraging efficiency is usually a matter of *trade-off* between competing priorities [8]. Theoretical models predict that, while joining a group will not increase an individual’s ability to find food, the time spent to obtain food is reduced. The trade-off for the reduction in searching time is that a smaller share of food will be available, the only exception being when the resource is so abundant that consumption by one individual does not decrease the availability for others.

The Food Particle Swarm (FPS) model presented in [3] and [4] was a simple abstraction of an ecological system, in which the animals, the environment and their interaction had only limited affinity with real animals and their behaviour. There we argued that elements of the standard PS algorithm such as *social attraction* and *communication amongst individuals* map fairly strongly to the behavioural ecology problem. Some minor changes to the basic algorithm allowed us to use the PS as a simulation model of the system. The objectives of those works were to study the emergence of grouping behaviour among particles, and at the same time to understand which are the most successful designs and the best parameter values for the particle swarm

algorithm to allow this emergence. Even though the model was extremely simple and unrealistic (in the sense intended in [1]), we obtained interesting results.

With this paper, we propose a more biologically plausible alternative to that model. The additions come from discussions with specialists in the field of fish behaviour and group-foraging (see sec. 5).

The main goal of this work is to check whether the qualitatively realistic behaviours produced in [3] and [4] hold up in a model constrained by more plausible biological assumptions.

The motivation for our work can be found in [5]. There, the authors argue that simulations can be valuable tools to verify the consistency of a theoretical position, although they cannot substitute empirical data collection. They define simulation models as *opaque thought experiments* in which the consequences follow the premises in a non-obvious manner (i.e., the internal working of a simulation model is complex). The “opacity” of the models implies long developing and testing of the theory, before relating it to empirical investigation of the world. Although still being in the development phase, and not yet linked to any real empirical data, the aim of the present work is to model a natural phenomenon (group-foraging) which is complex, emergent and self-organising. We do so through artificial life simulations.

The remainder of this document is structured as follows. In section 2 we describe the general approach to the problem of group-foraging: in sections 2.1 and 2.2 we present a brief summary of the FPS algorithm as it was introduced in [3] and [4] respectively, we then explain in detail the biological factors added to the model which lead to the Ecological Particle Swarm (EcoPS) (2.3). In section 3 we present the experiments settings (3.1) and simulation results (3.2). We conclude in section 4.

2. ECOLOGICAL PARTICLE SWARM ALGORITHM

In [3], we modelled a very basic abstraction of the group-foraging problem in which:

- there are no predators;
- the population size is constant;
- animals can communicate with each other.

These simplifications are still valid in the EcoPS model, which also shares the following behavioural features with [3] and [4].

Particles in the swarm represent foraging animals. The food is in the form of circular *patches*, which are scattered over the (2-dimensional) landscape according to different configurations, mirroring different situations that may happen in nature.

When a particle lands on a patch, it stops and starts feeding. Since particles move according to PS equations, they attract each other, causing other particles to gather on the same patch. When the food on a patch finishes, the particles start foraging again. The fitness of each particle is calculated as the amount of food available for the particle to eat (i.e., the amount of food left on the patch), weighted by the amount of food that a particle can eat (i.e., the particle’s own intake factor).

2.1 FPS without aura

A feature which uniquely characterises the model in [3] is that animals can neither see nor smell the food: the only way they can find food is relying on other animals finding it before them and then attracting them towards the food.

2.2 FPS with aura

In the version of the model presented in [4], the restriction on the perception of the patches is relaxed by introducing a smell (*aura*) surrounding the food and letting the animals being guided by it. Particles are attracted to a patch by its aura, which has an intensity that (1) is proportional to the amount of food available on the patch, (2) is an exponentially decreasing function of the distance from the source, and (3) has a spread proportional to the size of the patch. Particles follow the “aura gradient” until they reach the surface of the patch, where they stop and start feeding.

The amount of food eaten by each particle is its fitness F , which can be interpreted as the energy gained while feeding. Therefore, when a particle reaches a patch and there is food available, the fitness increases. Unlike in nature, the fitness of these abstract animals also increases when they approach food patches (i.e., the intensity of the aura contributes to the fitness). The first situation mirrored the biological nature of the problem (an animal gaining energy while feeding on a patch of food), and was therefore biologically plausible. However, the fact that the distance of the particle from the patch contributed to the fitness was entirely unbiological (the “smell” of food does not increase an animal’s energy), though intentional. More formally,

$$F(t) = \begin{cases} F(t-1) + FE & \text{if particle is on patch} \\ FA * e^{-(dis - \frac{FS}{2})} & \text{otherwise} \end{cases} \quad (1)$$

where: $F(t-1)$ is the particle’s fitness at the previous time-step; FE is the amount of food eaten by the particle; FA is the amount of food available on the patch; dis is the distance between the particle and the surface of the patch; FS is the size of the patch.

The particles are driven by PS equations for force f

$$f_i = \phi_1 R_1(x_{s_i} - x_i) + \phi_2 R_2(x_{p_i} - x_i) \quad (2)$$

with the first component (*social interaction*) causing the particles to be attracted with random magnitude towards the best position found by the swarm, and the second (*individual learning*) causing the particles to be attracted with random magnitude to the best position found by themselves.

Velocities are updated according to the formula

$$v_i(t) = \begin{cases} 0 & \text{if food is on patch} \\ Random & \text{if food is finished} \\ \kappa((\omega v_i(t-1)) + \Delta t f_i) & \text{otherwise} \end{cases} \quad (3)$$

which differs from the standard PS velocity update equation since it allows the particles to stop on the food patches; when a particle leaves the patch, its velocity is re-initialised at random.

Finally, the new particle’s position is computed as usual, i.e.,

$$x_i(t) = x_i(t-1) + \Delta t v_i(t) \quad (4)$$

where Δt is a factor used to decrease the step the particles take when they move, and has been introduced to promote a smoother movement and more refined search.

2.3 EcoPS

As was mentioned in the introduction, this paper presents a more biologically plausible alternative to the FPS model. This EcoPS model includes the following biological factors:

Local neighbourhood - in FPS, we used a global neighbourhood topology, in which each particle was able to exchange information with everyone else. In real biological systems, the information the animals share is often very limited and they have to rely only on their nearest neighbours. We have therefore substituted the global topology neighbourhood with a local one: each particle will only exchange information with a limited number of other particles, namely the ones which are within a certain distance. (We consider that a particle is in its own list of neighbours: an animal relies on its neighbours' information, as well as its own.)

Forgetful particles - the key biological issue that inspires this addition is that a place that was good for feeding some time ago might no longer be good; we have therefore incorporated forgetting as well as learning into the model. We have created a "sliding temporal window" of constant size within which the particles remember their previous best position. As the time progresses, this "memory window" slides forward, allowing the particles to forget about previous locations and to learn about new ones. Whatever the value of the fitness, at each time step an old location is deleted from the particle's memory and a new one is introduced. The new best fitness value for the particle is then chosen amongst the values in memory.

Fitness as energy - we re-define the fitness of a particle as the long-term total amount of food consumed minus the energy spent in movement, which is proportional to the velocity (for most animals, it is linearly related to their current speed). In [10], the rate of acquiring energy when swimming at speed v is given by

$$F(v) = PQ - (1 - P\tau)D_0v^2 \quad (5)$$

where: P is the mean rate at which food is encountered, Q is the energy gained by eating one unit of food, τ is the feeding time, and D_0 is the swimming cost coefficient (D_0v^2 is the rate at which energy is used by the animal swimming at speed v^2 , with $D_0 = 6\pi\mu l$, where μ is the viscosity of the fluid and l is the length of the animal). We use this equation to derive the fitness of each particle as

$$F(t) = \begin{cases} F(t-1) + CQ & \text{if particle is on patch} \\ F(t-1) - D_0v^2 & \text{otherwise} \end{cases} \quad (6)$$

where: C is the amount of food consumed, and Q and D_0v^2 are as before.

Environmental attraction - as we mentioned in section 2.2, the aura's contribution to the fitness has no biological analogue. Instead, we have added a third term to the acceleration equation which represents the strength of attraction (related to the aura) to the nearby patch. Being interpreted as the smell that the food emits, the aura actively attracts the particles towards the patches, affecting the dynamics of the sys-

tem. The new equation for the acceleration is now composed of three terms:

1. social interaction,
2. individual learning, and
3. environmental attraction.

$$f_i = \underbrace{\phi_1 R_1(x_{s_i} - x_i)}_1 + \underbrace{\phi_2 R_2(x_{p_i} - x_i)}_2 + \underbrace{\phi_3 R_3 \times \text{aura} \frac{(x_{A_i} - x_i)}{|x_{A_i} - x_i|}}_3 \quad (7)$$

where: ϕ_1 , ϕ_2 and ϕ_3 are the *social interaction*, the *individual learning* and *environmental attraction* rate respectively; R_1 , R_2 and R_3 are random variables; x_i is the current position of the particle; x_{s_i} is the best position found by the swarm; x_{p_i} is the best position found by the particle; x_{A_i} is the centre of the patch of food closest to the particle; *aura* is the intensity of the aura at the position of the particle.

In EcoPS we adopt the update equations for position and velocity proposed in FPS (equations 3 and 4).

3. EXPERIMENTS

From the results in [3] and particularly in [4], we know that FPS is a promising model for the simulation of group-foraging which produces qualitatively realistic behaviours. Therefore, the main objective of this work is to verify that the four biological assumptions (introduced in sec. 2.3) are advantageous for the model.

In order to do this, we have run experiments for four scenarios, relating to the choice of particle type, neighbourhood topology and presence of food auras:

scenario 1. forgetful particles with a global topology and food patches with auras;

scenario 2. forgetful particles with a local topology and food patches with auras;

scenario 3. standard particles with a global topology and food patches without auras (this situation corresponds to the one in [3]);

scenario 4. standard particles with a local topology and food patches without auras.

3.1 Settings

For each scenario, the same set of parameter value combinations were tested, as detailed in Table 1 (the tests included all possible combinations of the values stated).

Each one of these combinations has been run in three different food configurations: with one, three and ten patches of food, randomly scattered on the landscape, and with different sizes and amounts of food available (see Table 2). Since we interpret the food as energy, we define the food amount factor (*FAF*) as the food energy per unit of patch area, and the food consumption factor (*FCF*) as the rate of consumption of food from each patch.

Given the large number of parameter combinations, it is impossible to present all the results obtained in the limited

Table 1: General experiment parameters

Parameter	Value
Number of iterations	250
Number of particles (N)	10, 25, 50
Δt	0.25, 0.5, 1.0
World size	(25×25), (50×50)
Movement cost	0.02, 0.0003, 1.5
Maximum velocity	0.0, 1.0, 5.0
Neighbourhood size	2.0, 5.0
Memory threshold	10, 50, 100
Minimum aura	0.0, 0.1

Table 2: Food patches configurations

	1	3	10
Size	7.0	6.0, 4.0, 2.0	2.0
Position	(-5.0,-5.0)	random	random
FAF	2.0	fixed/variable	1.0
FCF	0.1	0.1	0.1

space available. We therefore consider only the most significant ones. All the presented experiments were run in a (25×25) world, with 10 particles, 3 patches, movement cost 0.02 and all three choices of Δt . The other parameters were problem specific and have been chosen according to the specific case study. In particular, we focus on the following biological assumption:

Forgetful particles - we tested whether it is an advantage (in terms of energy gain) for a particle to have a limited amount of memory (i.e. the particle only remembers a certain number of past iterations) or to have an unlimited memory (i.e., the particle remembers its whole past “life”). We checked scenario 2 against 4, with parameter settings as in Table 3, to see whether different neighbourhood topologies have different effects on the memory of the particles.

Environmental attraction - we checked whether and how the aura surrounding the patches affects the foraging behaviour of the particles. We compared scenario 2 against 4, with settings as in Table 4. These tests might seem superfluous, since the difference between [3] and [4] was the addition of the aura to the model. In [4], however, the aura was used in the fitness calculation in a fairly non-biological way. In our EcoPS model, the aura is used in a more biological fashion, as the rate of attraction to the environment that an animal experiences when, for example, it smells some food.

We should point out that we do not consider any alternative for the fitness interpretation other than the one presented in sec. 2.3 (i.e., the fitness is interpreted as energy).

In all the experiments presented, we consider a local neighbourhood topology (i.e., scenarios 2 and 4). It is worth noting that such a topology is not only a realistic biological hypothesis, but it is in fact “the best” choice of neighbourhood for this model. This means that, when using this topology, the overall fitness of the swarm (interpreted as the sum of the fitness of the particles) is higher than it is using an alternative topology (see Figure 1). Despite intuition suggesting that having global knowledge about the rest of

Table 3: Experiment parameters

Forgetful particles	
Memory threshold	10, 50, 100
Neighbourhood size	2.0, 5.0
FAF	fixed
Maximum velocity	1.0

Table 4: Experiment parameters

Environmental attraction	
Maximum velocity	0.0, 1.0, 5.0
Neighbourhood size	2.0, 5.0
FAF	variable
Minimum aura	0.1

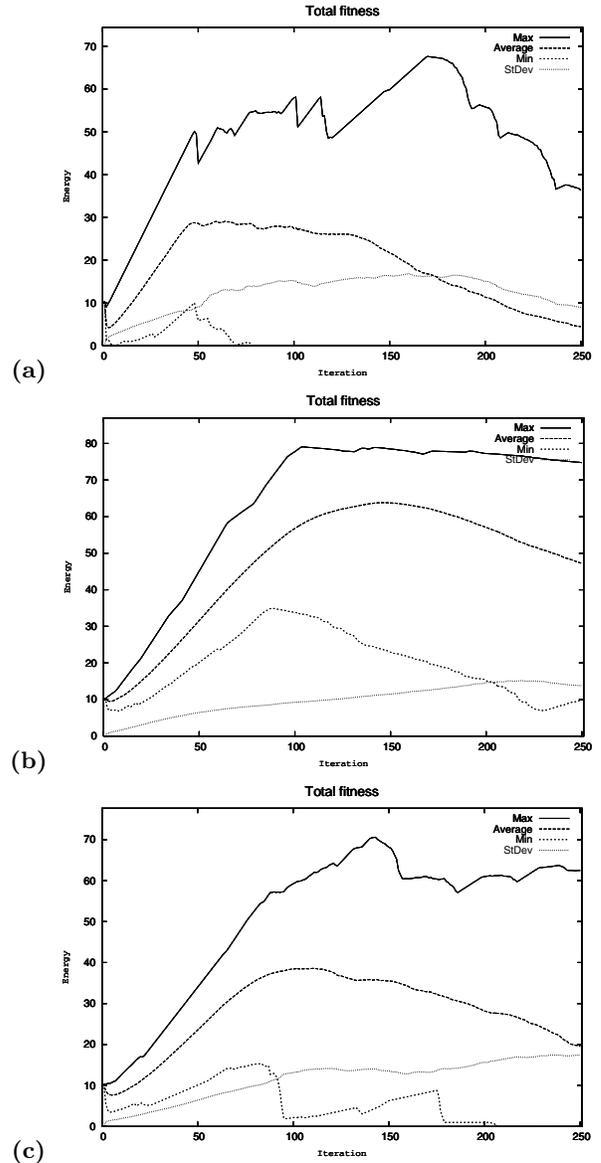


Figure 1: Swarm fitness with different neighbourhood topologies: (a) global, (b) local size=2.0, (c) local size=5.0. (X-axis: number of iterations; Y-axis: energy)

the swarm should be advantageous for a particle, the experiments show that the opposite is true. In fact, the smaller the neighbourhood, the better. One possible explanation for this phenomenon is that, when using a global neighbourhood, a higher number of particles initially reaches the patches to feed (peak in the fitness curve in Figure 1.(a)). As a consequence of the higher number of individuals feeding, the food finishes quickly, causing the fitness of the swarm to drop as quickly as it grows. This actually therefore results in a fairly low average fitness. On the other hand, with a local neighbourhood topology, fewer particles initially reach the patches (smoother curves in Figure 1 (b,c)), causing a more even distribution of food amongst the feeding particles, and a higher average fitness for the swarm. It appears that, even in such a simple model, adopting a local (and possibly small) neighbourhood topology is an advantage for the swarm.

In order to ensure that the results were not biased by the initial distribution of particles, we repeated the experiments 50 times with different random number generator seeds.

3.2 Results

Forgetful particles - from Figure 2, it is evident how memory plays a crucial role in the particles' behaviour. The results, in this case, confirm the biological assumption that forgetting is an advantage over having an unlimited memory, as this might lead to remember old places which are no longer good for feeding. Even more interestingly, it seems that the shorter the memory the better, since the number of particles which visits each patch increases. A possible justification for this is the fact that, when the food on a patch is exhausted, the position which was previously the best one for the particle is no longer such. Therefore, there is no need for the particle to retain that position in its memory, and the faster it is forgotten, the better as the particle will be able to remember some potentially better new positions. This is true in the present context where, once the food on a patch is exhausted, it does not regenerate. If, instead, the food was able to regenerate after being eaten, having a short memory would probably not be as advantageous, as the animal would forget the information about a good spot in the world and would have to learn about it again. In future developments of this research, we will explore the differences between short and long term memory.

Environmental attraction - the presence of the aura as environmental attraction presents an advantage for the foraging behaviour of the particles in terms of (i) energy gain, (ii) number of patches reached, (iii) amount of food eaten (Figure 3), and (iv) ability to form clusters¹ (Figure 4). As we have already mentioned (see sec. 2.3), in this paper we consider the aura as being part of the environment (i.e., in the form of the "environmental attraction" in equations 8), as opposed to the interpretation of aura in [4], where it was used only in the fitness evaluation. Being interpreted as the smell that the food emits, and being proportional to the amount of food available on each patch, the aura actively attracts the particles towards the patches, caus-

ing the particles' behaviour to be strongly dependent on it, regardless of other parameters.

4. CONCLUSION

At the very beginning of this paper, we asked the following question: are the qualitatively realistic results of a simple, abstract simulation model for group-foraging behaviour "strong" enough to still be valid when the model is made more biologically plausible? We can now conclude that those results hold up to the biologically plausible constraints introduced here.

We have added four biological factors into the "unrealistic" model presented in [3] and [4]: local neighbourhood, forgetful particles, fitness interpreted as energy and environmental attraction. We have then tested this new EcoPS model to verify that these biological hypothesis (i) were an advantage for the behaviour of the particles, and (ii) represented a step forward in the direction of the simulation of a more realistic animal system.

The biological factors added to the model gave us a tool to better interpret the qualitatively realistic behaviours obtained as the results of the simulation, and will be a useful bridge to connect these simulation experiments to empirical ones.

Of course, the additions made were not the only ones possible; the few features chosen, after consultation with behavioural ecologists, were felt to be amongst the most key in addressing issues of biological plausibility. Future expansions will see the inclusion of further biological details to the model.

This work was intended to be a simulation model which aimed to help shed some light purely on the field of animal behaviours. Therefore, no attempt has been made to link the results found with the theory of Particle Swarm Optimisation (so far).

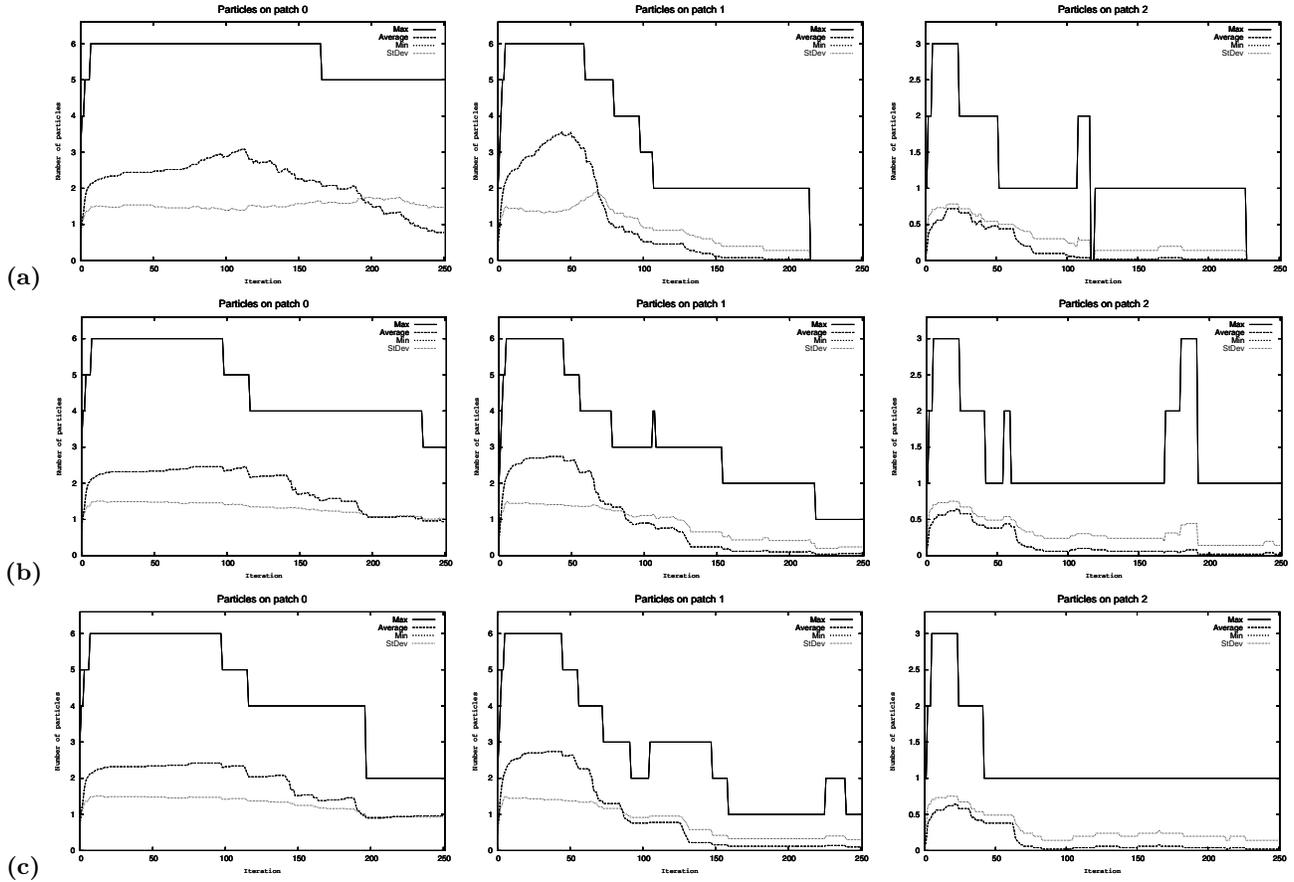
5. ACKNOWLEDGMENTS

We would like to thank Graeme Ruxton, Jens Krause and Darren Croft for their feedback, which helped clarify and validate the revised model in this paper.

This research is part of the EPSRC *XPS* project (GR/T11234/01).

¹We use the same definition of clusters used in [3] and [4].

scenario 2.



scenario 4.

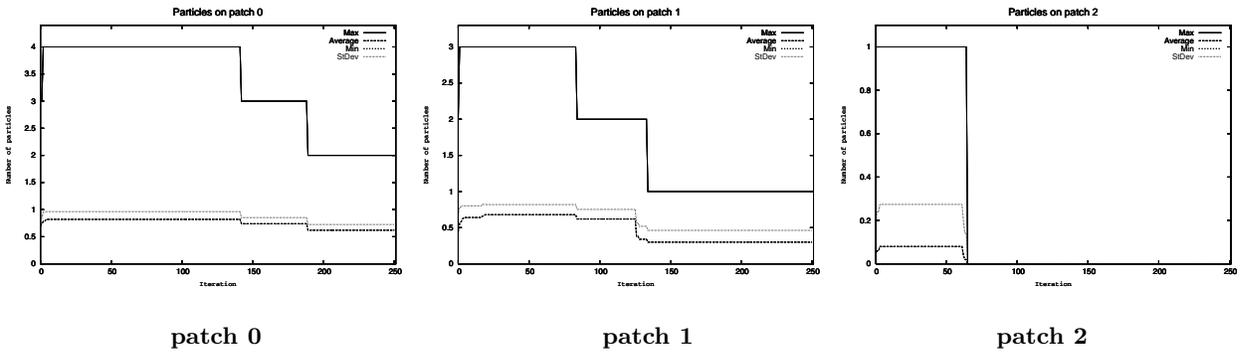
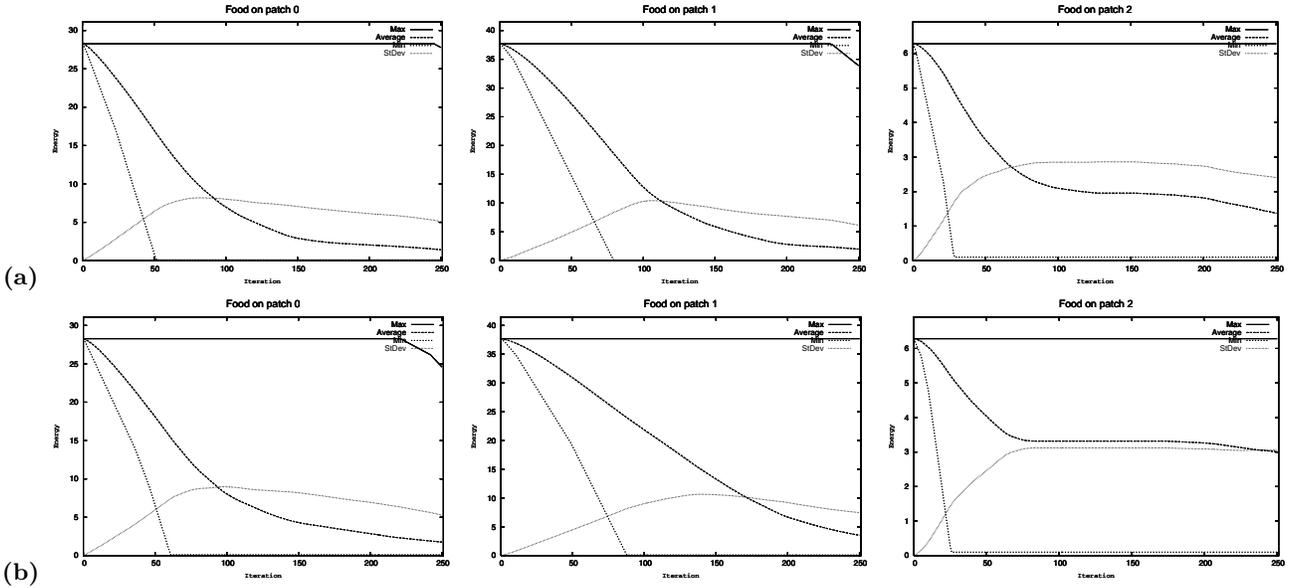
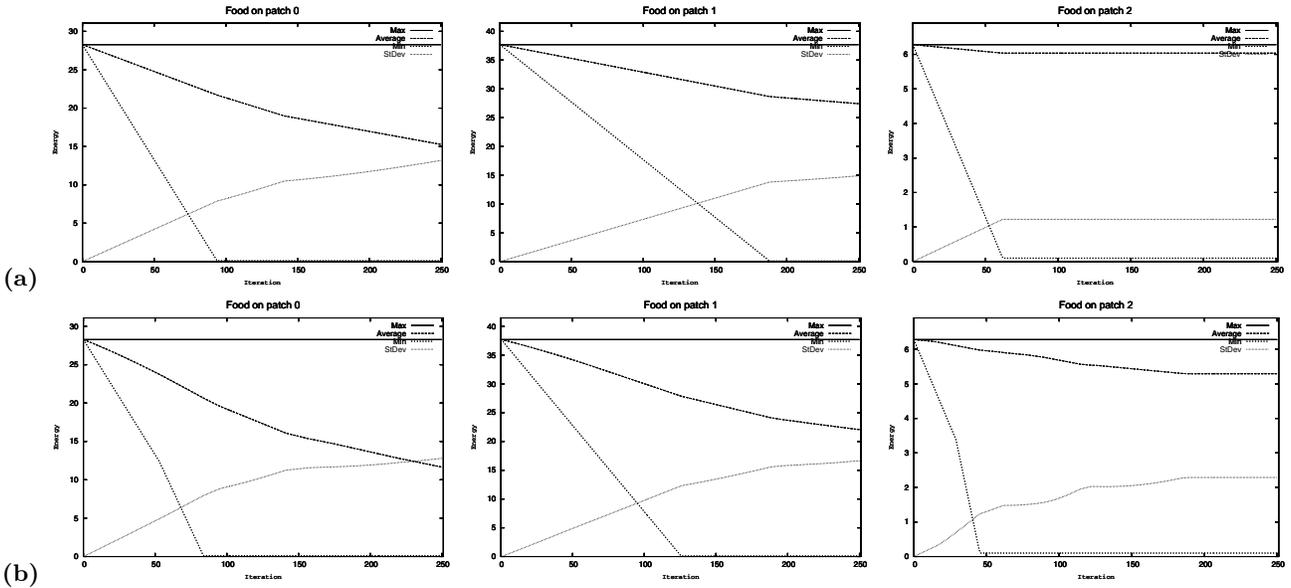


Figure 2: Difference in number of particles on each patch between scenario 2 (forgetful particles: (a) memory threshold=10, (b) memory threshold=50, (c) memory threshold=100) and 4 (particles with unlimited memory). (X-axis: number of iterations; Y-axis: number of particles)

scenario 2.



scenario 4.



patch 0

patch 1

patch 2

Figure 3: Difference in amount of food eaten on each patch between scenario 2 and 4 w.r.t. neighbourhood size: (a) neighbourhood size=2.0, (b) neighbourhood size=5.0. (X-axis: number of iterations; Y-axis: energy)

6. REFERENCES

- [1] M.A. Bedau, *Can Unrealistic Computer Models Illuminate Theoretical Biology?*, GECCO1999 - Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program, 1999
- [2] D.L. DeAngelis and W.M. Mooij, *Individual-Based Modeling of Ecological and Evolutionary Processes*, Annual Reviews in Ecology, Evolution and Systematics, 2005
- [3] C. Di Chio, R. Poli and P. Di Chio, *Extending the Particle Swarm Algorithm to Model Animal Foraging Behaviour*, University of Essex, Technical Report, 2006.
- [4] C. Di Chio, R. Poli and P. Di Chio, *Modelling Group-Foraging Behaviour with Particle Swarms*, PPSN2006 - Ninth International Conference on Parallel Problem Solving from Nature, 2006.
- [5] E.A. Di Paolo, J. Noble and S. Bullock, *Simulation Models as Opaque Thought Experiments*, Artificial Life VII - Proceedings of the Seventh International Conference on the Simulation and Synthesis of Living Systems, 2000.
- [6] J. Kennedy and R.C. Eberhart, *Swarm Intelligence*, Morgan Kaufmann Publishers, 2001.
- [7] J. Krause and G.D. Ruxton, *Living in Groups*, Oxford University Press, 2002.
- [8] D. MacFarland, *Animal behaviour*, Longman, 1999.
- [9] J.K. Parrish and W.M. Hamner, *Animal Groups in Three Dimensions*, Cambridge University Press, 1997.
- [10] J.W. Pitchford, A. James and J. Brindley, *Optimal Foraging in Patchy Turbulent Environments*, Marine Ecology Progress Series, 2003

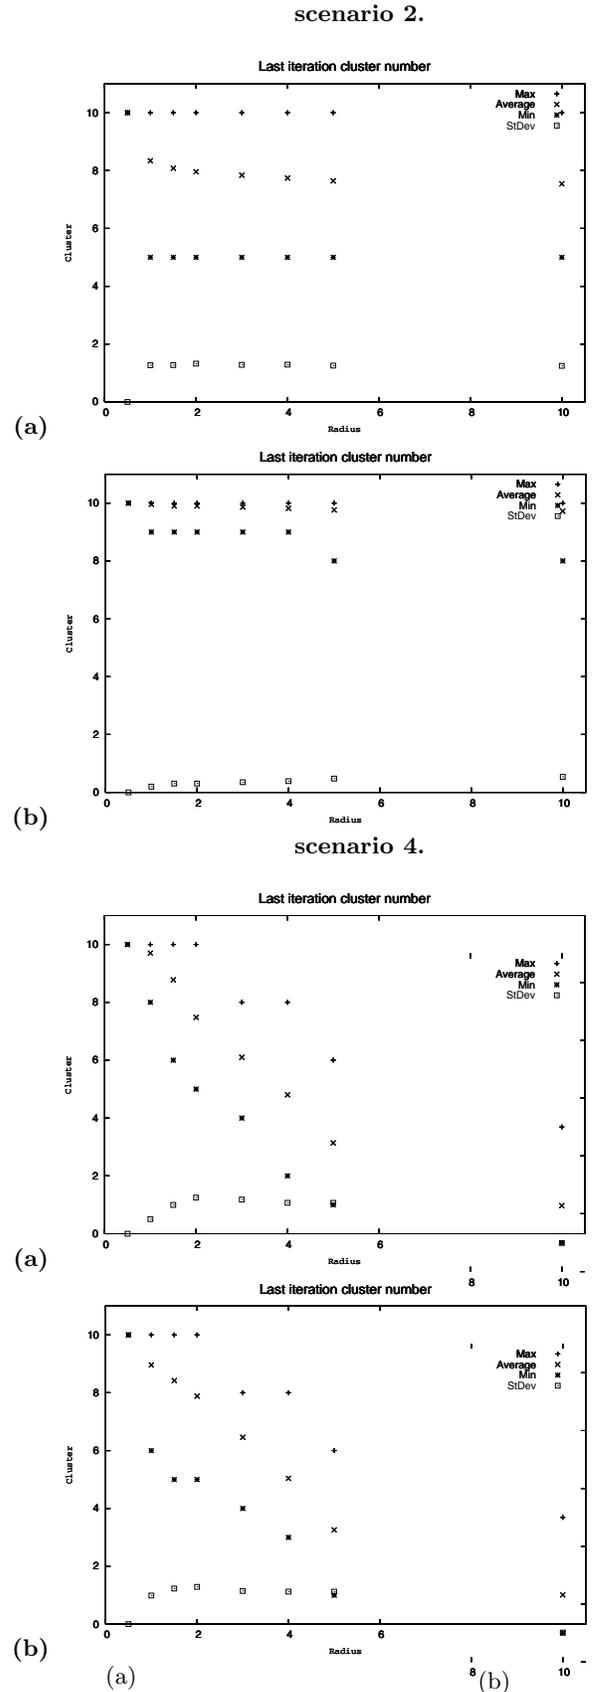


Figure 4: Difference in clustering between scenario 2 and 4 w.r.t. neighbourhood size: (a) neighbourhood size=2.0, (b) neighbourhood size=5.0. (X-axis: radius; Y-axis: number of cluster)