

Modelling Group-Foraging Behaviour with Particle Swarms

Cecilia Di Chio¹, Riccardo Poli¹, and Paolo Di Chio²

¹ Department of Computer Science,
University of Essex, UK
{cdichi, rpoli}@essex.ac.uk

² Dipartimento di Sistemi e Istituzioni per l'Economia,
University of L'Aquila, Italy
pdc@ec.univaq.it

Abstract. Despite the many features that the behaviour of the standard particle swarm algorithm shares with grouping behaviour in animals (e.g. social attraction and communication between individuals), this biologically inspired technique has been mainly used in classical optimisation problems (i.e. finding the optimal value in a fitness landscape). We present here a novel application for particle swarms: the simulation of *group-foraging* in animals. Animals looking for food sources are modelled as particles in a swarm moving over an abstract food landscape. The particles are guided to the food by a smell (or *aura*), which surrounds it and whose intensity is proportional to the amount of food available. The results show that this new extended version of the algorithm produces qualitatively realistic behaviour. For example, the simulation shows the emergence of group-foraging behaviour among particles.

1 Introduction

The particle swarm optimisation (PSO) algorithm is a function optimisation tool based on the simulation of a simplified social model, the *particle swarm* [4]. The algorithm uses a population (*swarm*) of candidate solutions (*particles*) that fly over the fitness landscape looking for optima. In the original version, the particles are driven by two forces which attract them to the best location encountered both by any other member of the swarm and by themselves.

Despite its socio-biological background, the field of applications for PSOs has mainly been the optimisation of nonlinear functions. In [3], we introduced a new way to use PSOs for the simulation of ecological processes, so far largely restricted to the field of individual-based modelling in ecology (for a review, see [2]). In particular, we looked at a classical problem in behavioural ecology: the *group-foraging* problem. With this paper, we extend our previous work by introducing more realistic features for the food sources. In particular, food sources are surrounded by an *aura*, which can be interpreted as the smell the food releases, and which attracts the particles to the source. Our objective is not only to study the emergence of grouping behaviour among particles, but also to understand which is the most successful design and the best parameter settings for

the particle swarm algorithm to allow this emergence. Our research is part of a multidisciplinary project called XPS¹. One of the goals of XPS is to extend the particle swarm algorithm with strategies from biology. Modelling group-foraging is just a first step towards the simulation of more general and complex group-behaviour in animals. We necessarily have to start simulating simple behaviours, but we intend to progressively add complexity, and eventually obtain a model as complete as possible of animal social behaviour.

The paper is structured as follows. In section 2 we give a brief introduction to the field of behavioural ecology and the problem of group-foraging. Section 3 describes our approach to the problem: the Food Particle Swarm (FPS) algorithm. In section 4 we summarise the settings for the experiments and present and discuss the results of our simulations. We conclude in section 5.

2 Behavioural Ecology and the Group-Foraging Problem

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 is the role an animal’s behaviour plays in allowing it to adapt to its environment. Among the more complex and intriguing animal behaviours, group-living is certainly one of the most studied, being such a widespread phenomenon in the animal kingdom [5]. Two general requirements for grouping behaviour are: (1) individuals have to be close to each other in space and time (e.g. the elective group size concept developed by Pitcher (see [5]) requires that the animals are close enough to each other in order to allow a continuous exchange of information), 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).

Animals show grouping behaviour for different reasons. One of these is *foraging* (i.e. in behavioural ecology, all those interactions that occur between a predator and its prey, being it animal or plant) [5]. Some of the most likely hypotheses for *group-foraging behaviour* are: (1) aggregations find more food (e.g. bigger preys, larger patches) 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 [7]; (3) by observing the behaviour of other members of a group, animals can gain useful information (e.g. individuals use information on the position of others to obtain food from sources that are otherwise difficult to find) [5].

Foraging efficiency is usually a matter of *trade-off* between competing priorities, e.g. energy gained versus energy spent, energy gained versus risk of predation, energy gained versus losses to rivals, etc.[6]. 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

¹ XPS stands for eXtended Particle Swarm. Details of the project can be found at <http://xps-swarm.essex.ac.uk>.

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. Animals that perform better at increasing the benefits of foraging and decreasing its costs will propagate their genes more effectively than those whose foraging behaviour is less effective [1].

3 The Food Particle Swarm (FPS) Algorithm

As mentioned above, group-foraging is a fairly complex behaviour, and modelling it is a difficult task. Therefore, in this work we focus on an abstraction of the group-foraging problem, where: we do not take the presence of predators into account; food sources are surrounded by an aura attracting the animals to them; animals can neither reproduce nor die; animals can communicate with each other; they can “smell” food, and this is the only way that they can detect the presence of a food source (e.g. they cannot see the food).

A group of animals looking for food is modelled with a swarm of particles “flying” over a 2D landscape scattered with sources of food. Each source is represented as a circular *patch*. The intensity of the aura surrounding the patches is proportional to the amount of food available and is an exponentially decreasing function of the distance from the source, while its spread is proportional to the size of the patch. Therefore, the larger the amount of food and the closer the patch, the “taller” the aura. Also, the bigger the patch, the wider the aura. Food patches are distributed at random on the landscape. To reflect different conditions that may happen in nature, we will consider four configurations (Figure 1):

1. a single small source but with a large amount of food, whose tall aura covers a small portion of the landscape;
2. a single large source but with a small amount of food, whose short aura covers most of the landscape;
3. a small number of medium sources covering a restricted portion of the landscape;
4. a large number of small sources, sparsely scattered on the landscape.

The particles move over the food landscape according to the rules of the extended version of the particle swarm algorithm introduced in [3]. The equations controlling acceleration, velocity, and position of the particles are as in the standard PSO (equations (1)-(3)), but with an extra control which allows the particles to stop on the patch. Namely:

$$f_i = \phi_1 R_1(x_{s_i} - x_i) + \phi_2 R_2(x_{p_i} - x_i); \tag{1}$$

$$v_i(t) = \begin{cases} 0 & \text{if food is on patch} \\ \text{Random} & \text{if food is just finished} \\ k((\omega v_i(t - 1)) + \Delta t f_i) & \text{otherwise} \end{cases} \tag{2}$$

$$x_i(t) = x_i(t - 1) + \Delta t v_i(t). \tag{3}$$

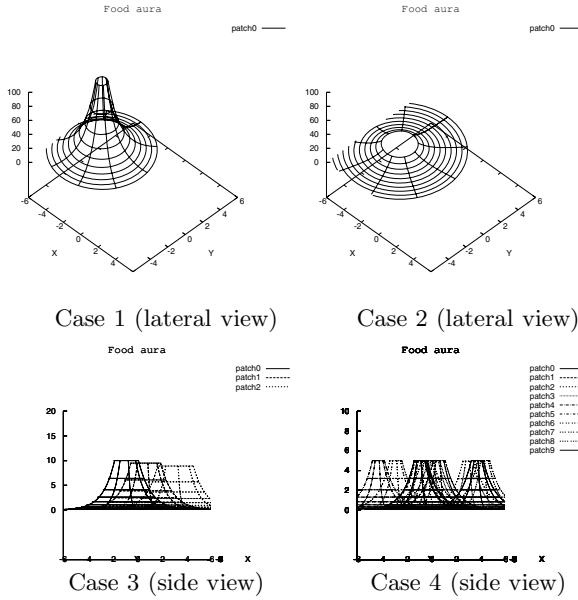


Fig. 1. Food patch distribution classes and aura intensity

where: ϕ_1 and ϕ_2 are the *social interaction* and the *individual learning rate* respectively; R_1 and R_2 are two random variables uniformly distributed in $[0, 1]$; x_i , x_{s_i} , and x_{p_i} are the current position of the particle, the best position found by the swarm, and the best position found by the particle respectively; k and ω are the constriction coefficient and the inertia weight respectively; Δt is a factor to decrease the step the particles take when they move to obtain “smoother” trajectories for the particles (i.e. a more refined search).

Particles are attracted to the food by its aura. They follow the gradient until they reach the surface of the patch, where they stop and start feeding. The food eaten by the particles increases their fitness. Unlike in nature, the fitness of our abstract animals is also increased when they approach food patches, i.e. the intensity of the aura contributes to the fitness. The first situation mirrors the biological nature of the problem, while the latter is inspired by the optimisation nature of the algorithm (where patches of food can be consider optima to be found). Formally:

$$fit_i(t) = \begin{cases} fit_i(t - 1) + FE & \text{if particle is on patch} \\ FA * e^{-(dis - \frac{FS}{2})} & \text{otherwise} \end{cases} \quad (4)$$

where: 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.

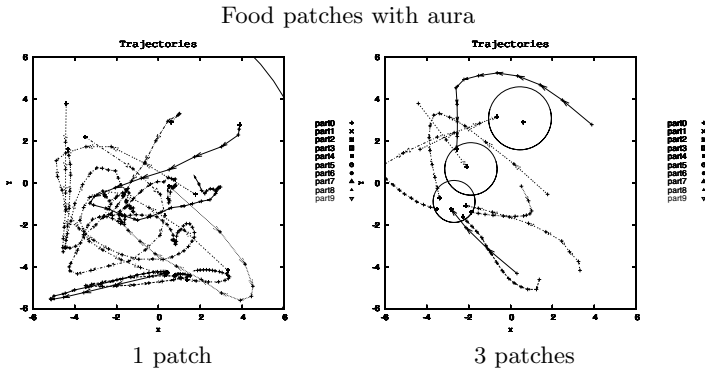


Fig. 2. Trajectories in FPS on different landscapes (200 iterations, $\Delta t = 0.25$)

The first particle that reaches a patch and starts eating will have the highest fitness, and therefore will become the best of the swarm. Other particles will then be attracted to the same patch not only by its aura, but also by the presence of the particle already feeding on it. When the food available is entirely consumed, the particles leave the patch with random velocity and start foraging again. While feeding, the amount of food available on the patch decreases, and so does its aura. Eventually, there will no longer be any food left on the patch, and the aura will be completely dissolved. If a particle is at the intersection of two or more auras, it will follow the one with highest intensity. If the particles are too far from any patch, or if the food is no longer available, the effect of the aura is considered irrelevant and their fitness is set to zero. This causes x_{s_i} to be equal to x_{p_i} and, as a consequence, the particles start oscillating close to these positions.

4 Results

The goal of this research is to produce a model of abstract animals and their foraging environment through which to observe the emergence of group-foraging behaviour. In terms of simulation, this means that the particles have to both gather together on the food sources (i.e. form *clusters* when feeding), and eat as much food as they can find (i.e. achieve a high fitness). Different foraging abilities and swarm sizes have been tested to investigate which settings produce a more “natural” behaviour, with various patch distributions² to observe how these behaviours change (Figure 2). Table 1 summarises both the algorithm and the landscape parameters.

All the parameters except Δt are related to the nature of the group-foraging problem. Δt is instead an algorithmic “artifact” which needs to be small for the continuous force equation to be discretised effectively (e.g. the default step of

² The four alternative patch distributions reflect the qualitative patterns explained in section 3.

Table 1. Independently varied parameter settings

| Parameter | Value |
|--------------------------------|---|
| Number of iterations | 200, 500 |
| Number of particles (N) | 10, 30 |
| Δt | 0.1, 0.25, 0.5, 1.0 |
| Number of food patches (F) | 1, 1, 3, 10 |
| Patch size & food amount | $F = 1$: size 1, food 100 |
| | $F = 1$: size 3, food 10 |
| | $F = 3$: size from 2 to 3, food from 9 to 10 |
| | $F = 10$: size 1, food 5 |




Fig. 3. Fitness in FPS and SPS on different landscapes (200 iterations, $\Delta t = 0.25$); means over 100 independent runs

the standard particle swarm algorithm ($\Delta t = 1$) makes the particles move with excessively large jumps which can cause them to miss the food patches).

Since the initial position of the particles on the landscape could influence the behaviour of the simulation, we repeated the experiments 100 times with different random number generator seeds. For completeness, we have run some comparative tests with the standard version of the particle swarm algorithm (SPS) on this newly defined landscape. We also report here some of the most significant results obtained in [3], to show how the introduction of the concept of




Fig. 6. Amount of food eaten in FPS and SPS (500 iterations)




Fig. 7. Clustering in FPS and SPS ($\Delta t = 0.25$, threshold = 0.5, 1.0, 1.5)

aura has improved the behaviour of the model. Given the large number of combinations of parameters we used in the experiments, it is impossible to present all the results obtained in the limited space available. We consider here only the more significant ones, run with 10 particles, 3 patches, and $\Delta t = 0.25$ and 1.0. We have checked differences between the FPS and the SPS versions, in both types of landscape (i.e. with or without aura), with respect to number of patches visited (i.e. the ability of the particles to find food), fitness value (i.e. the amount of food eaten by each particle), and clustering (i.e. the ability of the particles to gather together). The clustering technique uses the following definition.

Definition 1. *Two particles p_1 and p_2 are in the same cluster if there exists an ordered set of particles $\{p_{(0,1)}, p_{(0,2)}, \dots, p_{(0,n)}\}$, with $p_{(0,1)} = p_1$ and $p_{(0,n)} = p_2$, such that $d(p_{(0,k)}, p_{(0,k+1)}) \leq r$, where r is the cluster threshold.*

The results highlight that the introduction of the aura is beneficial both in terms of amount of food eaten (Figures 3 and 4) and of number of clusters (Figure 5).

The results confirm that, in FPS, there are a larger number of particles that are able to visit more patches, and therefore to eat a greater quantity of food (Figure 6). From the figures it is also evident how, with a larger value for Δt , particles in the SPS model succeed in finding all the patches, but they require more time than the particles moving according to the FPS algorithm. The experiments reveal that particles in the SPS model cluster more than they do in the FPS simulation. Our hypothesis is that this is due to the “stop-eat-restart” behaviour of particles in the FPS model: the random re-initialisation of velocities can cause trajectories to be directed differently, resulting in a lesser ability to aggregate. This phenomenon tends in any case to decrease with time (Figure 7). On average, despite the fact that particles in the SPS algorithm show a better grouping behaviour (for smaller iterations), the ones in the FPS model are able to find a greater number of patches and eat a larger quantity of food.

5 Conclusion

The standard particle swarm algorithm shares features like social attraction and communication between individuals with group-foraging in animals. It is therefore surprising that its main field of application has only been function optimisation. By using the particle swarm algorithm as a simulation tool, we want to change this and take advantage of this biologically inspired technique.

We have shown how it is possible to obtain a version of the particle swarm model well suited for the foraging problem by modifying the standard algorithm. Here the landscape is no longer static (food can be eaten), and the fitness is now related both to the proximity to food (the optimisation problem of finding food) and internal energy (the biological problem of eating enough). From the experiments, we have seen that these changes extend the standard particle swarm model into one which produces qualitatively realistic behaviour.

This work is part of an initial investigation, whose final goal is the simulation of more complex group behaviour in animals. With [3], we introduced a simulation for a simple abstraction of the group-foraging problem. In this paper,

we have further extended this study by refining the definition of the landscape, through the introduction of the concept of an aura surrounding the patches of food. Future investigations will focus on introducing other realistic features for the food sources (e.g. allowing the patches to deteriorate and regenerate, *ephemeral* patches) and for the particles (e.g. allowing them to reproduce and die). In fact, as stated in [8], “*it is possible that many of the different collective patterns are generated by small variations in the rules followed by individual group members*”.

References

1. J. Alcock, *Animal Behavior*, Sinauer Associates, 1998
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*, ANTS2006 - Fifth International Workshop on Ant Colony Optimization and Swarm Intelligence, 2006.
4. J. Kennedy and R.C. Eberhart, *Swarm Intelligence*, Morgan Kaufmann Publishers, 2001.
5. J. Krause and G.D. Ruxton, *Living in Groups*, Oxford University Press, 2002.
6. D. MacFarland, *Animal behaviour*, Longman, 1999.
7. J.K Parrish and W.M. Hamner, *Animal Groups in Three Dimensions*, Cambridge University Press, 1997.
8. D.J.T. Sumpter, *The principles of collective animal behaviour*, Philosophical Transactions of the Royal Society of London: Series B, 2006.