Extending the Particle Swarm Algorithm to Model Animal Foraging Behaviour

Cecilia Di Chio[1], Riccardo Poli[1], and Paolo Di Chio[2]

[1] Department of Computer Science, University of Essex, Colchester, United Kingdom
{cdichi, rpoli}@essex.ac.uk
[2] Dipartimento di Sistemi e Istituzioni per l'Economia
University of L'Aquila, L'Aquila, Italy
pdc@ec.univaq.it

1 Introduction

The particle swarm algorithm [1] contains elements which map fairly strongly to the *group-foraging* problem in behavioural ecology: its continuous equations of motion include concepts of social attraction and communication between individuals, two of the general requirements for grouping behaviour [2]. Despite its socio-biological background, the particle swarm algorithm has rarely been applied to biological problems, largely remaining a technique used in classical optimisation problems. In this paper [3], we show how some simple adaptions to the standard algorithm can make it well suited for the foraging problem.

This work introduces a new way to look at the particle swarm algorithm, i.e. using it as a *simulation* tool in the biological field of behavioural ecology. Our research is part of the *XPS*[1] multidisciplinary project which aims, among other things, to explore extensions of the particle swarm algorithm by including strategies from biology. This work on foraging behaviour represents a first step in simulating more complex group behaviour in animals.

2 Approach

The simulation of animal grouping behaviour is extremely complex. Therefore, in order to make progress, we focus here on simulating an abstraction of the group-foraging problem, where: (a) there are no predators or any other source of risk or danger; (b) animals neither give birth nor die; (c) animals are "blind" and have no sense of smell, but they can communicate with everyone else, regardless of the size of the world; (d) the food sources do not deteriorate unless eaten and once a source is exhausted, it does not regenerate.

We propose two approaches to model foraging behaviour: the first uses a standard particle swarm algorithm, with the particles just slowing down in the proximity of food; the second approach modifies the basic algorithm in order to make the particles actually stop on the food source and remain there to eat.

[1] Details of the project can be found at http://xps-swarm.essex.ac.uk.

M. Dorigo et al. (Eds.): ANTS 2006, LNCS 4150, pp. 514–515, 2006.

The general idea behind our two approaches is the following. The particles in the swarm represent the animals looking for food sources. The sources are distributed over the 2-dimensional world and take the form of *patches* which contain a certain amount of food. To explore different situations that can happen in nature, there are three different configurations of food with respect to the number of patches, their size and the amount of food they contain.

The particles move over the food landscape according to the rules of the particle swarm algorithm. When a particle lands on a patch of food, it starts feeding on it, i.e. the amount of food available on that patch decreases (while the size remains the same). Because of the way the algorithm works, other particles are attracted to this patch, and start feeding as well. Eventually, the food on the patch will be exhausted (i.e. it will reach a minimum threshold), and the particles will start foraging again.

The food eaten by the particle is interpreted as the energy that the animal gains when feeding. The food intake, or equivalently the energy gained, represents the fitness of the particle. In the standard particle swarm algorithm, the particles represent the coordinates of points in the search space and their fitness is simply given by the value of the function we want to optimise. In our simulation, the fitness of the particles is evaluated 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).

3 Conclusion

We have shown how some simple adaptions to the standard particle swarm algorithm can make it well suited for the foraging problem.

The results (see [3]) show that the changes convert the standard algorithm into one which produces qualitatively realistic behaviour for a simplified model of abstract animals and their foraging environment. We have also highlighted sensible envelopes for parameter values and shown that it is important to keep the particles' trajectories "smooth".

With this simulation, we have begun with a simple abstraction of the group-foraging problem. In the near future, we will extend the analysis of this behaviour further, by introducing more realistic features for the food sources, and by refining the parallelism between particles and animals.

References

1. Kennedy, J., Eberhart, R. C.: *Swarm Intelligence*, Morgan Kaufmann Publishers (2001)
2. Krause, J., Ruxton, G. D.: *Living in Groups*, Oxford University Press (2002)
3. Di Chio, C., Poli, R., Di Chio, P.: *Extending the Particle Swarm Algorithm to Model Animal Foraging Behaviour*, University of Essex (2006)